REPUBLIC OF UGANDA | TRANSPORT SECTOR

Uganda Transport Status Report 2026

Needs, network, gaps, costs and where resources should go across 135 districts, with the links to agriculture, minerals, health, education and industry: a spatial analysis with the AICTEM Transport Infrastructure Planning and Resource Allocation System (TIPRAS)

Road network modelled: 247,972 km | Rural Access Index today: 28.7 % | Candidate programme: UGX 92.8 tn | 135 districts

Reference run 014c9cd7 · norms UGA-2026.09-draft-f · prices UCSN+EV-UGA-2026.09b · base year 2026 · horizon 2026–2030 · exploratory

October 2026

About this report

This report sets out the status of Uganda's transport infrastructure — chiefly its roads — and what it would take to meet the country's access and movement needs. It answers the five planning questions of the AICTEM Transport Infrastructure Planning and Resource Allocation System (TIPRAS): what the needs are, what is available, where the gaps are, what it costs to close them and where resources should go, including what transport does for agriculture, minerals, health, education and industry. The chapters follow the engine's design manual (docs/00–23). Every figure is produced by the engine from its data layers and carries its quality tier.

Item	Value
Engine	TIPRAS (TransportPlanning.jl): modules TI-01 to TI-10, GIS core (CSR graph, multi-source Dijkstra, cut edges, Moran's I and LISA), JuMP/HiGHS optimiser
Reference run	014c9cd7-024b-4b62-8a40-e72bfe960683: geography UGA-2026.09-c, norms UGA-2026.09-draft-f, prices UCSN+EV-UGA-2026.09b; scenarios A–G and four optimiser specifications
Horizon	Base year 2026; allocation and outcomes to 2030
Geography	135 districts; 817,354 road links (247,972 km); 175,367 settlement cells; 24,952 sector nodes
Currency	Uganda shillings (UGX) at 2026 prices; evidence carried to 2026 with World Bank GDP deflators and exchange rates
Simulations	Seven one-at-a-time simulations (prices low/high, envelope ±30 %, unknown surface as earth, UCSN-first price map) in addition to scenarios A–G and the optimiser runs
Quality tiers	observed · measured · normed (evidence-based norm) · imputed · modelled · assumed; unit prices E5 (evidence candidate) and E7 (UCSN illustrative)
Rules	the seven gap dimensions are never combined into an index; capital is always shown with its induced maintenance; no scenario is called recommended; network attributes are never inferred
Status	Exploratory: draft norms and candidate prices. Figures are planning estimates to be confirmed by the road agencies' data (docs/requests).
Reproduction	julia scripts/run_uga.jl; julia scripts/run_simulations.jl; python report/build/collect.py; python report/build/cache_layers.py; python report/build/figures.py; python report/build/render_equations.py; node report/build/build_docx.js; report/build/finalise_word.ps1

Contents

About this report 2
Contents 3
List of figures 11
List of tables 15
Executive summary 19
The needs 19
What the network has 19
The gaps — seven dimensions, never combined 19
What it costs 20
Where resources go 20
Simulations 20
What the evidence is, and is not 21
1 Introduction 22
1.1 Why a needs-based planning engine 22
1.2 Uganda in 2026: the planning context 23
1.3 How the roads sector is organised 26
1.4 The five planning questions 27
1.5 Principles that shape every result 28
1.6 How to read this report 28
1.7 Structure of the report 29
1.8 What this report can and cannot say 30
1.9 Implications for planning 30
2 Architecture of the planning engine 31
2.1 The planning chain 31
2.2 How the system is put together 33
2.3 The run model 35
2.4 The fiscal loop 36
2.5 Performance 36
2.6 Where the build departs from the design 37
2.7 Implications for planning 37
3 Data model, quality tiers and provenance 38
3.1 Why the data model matters for planners 38
3.2 Schemas 38
3.3 Keys and spatial keys 39

3.4 The geography tier ... 39
3.5 The norms tier ... 40
3.6 Quality and basis: two tags, never merged ... 42
3.7 Observations ... 43
3.8 Results and the shared GIS layer ... 43
3.9 Provenance: the manifest, the trace and `explain()` ... 43
3.10 Units, currency and time ... 44
3.11 The missing-data policy ... 44
3.12 Implications for planning ... 45
4 TI-01 Spatial base: people, places and the network ... 46
4.1 What the module does ... 46
4.2 Method ... 46
4.3 Results: the population base ... 48
4.4 Results: the network ... 49
4.5 Results: the places that must be reached ... 56
4.6 Data quality and limitations ... 60
4.7 Implications for planning ... 60
5 TI-02 Needs: who must be connected, and to what ... 61
5.1 What the module does ... 61
5.2 Method ... 61
5.3 Results: the access requirement ... 64
5.4 Results: node connection standards ... 70
5.5 Results: movement generation ... 70
5.6 Data quality and limitations ... 71
5.7 Implications for planning ... 71
6 TI-03 Demand and the required standard of every link ... 72
6.1 What the module does ... 72
6.2 Method ... 72
6.3 Results: the distribution of movement ... 75
6.4 Results: the required standard ... 80
6.5 Data quality and limitations ... 82
6.6 Implications for planning ... 82
7 The network today: supply and access ... 84
7.1 What the module does ... 84
7.2 Method ... 85
7.3 Results: what the network is made of ... 86
7.4 Results: links closed in the wet season ... 87

7.5 Results: the Rural Access Index ... 89
7.6 Results: travel time to services ... 92
7.7 Results: condition from agency statistics ... 100
7.8 Data quality and limitations ... 101
7.9 Implications for planning ... 101
8 Gaps: seven dimensions, never combined ... 103
8.1 What the module does ... 103
8.2 Why the seven gaps are never added ... 103
8.3 Method ... 104
8.4 Results: the gap register ... 105
8.5 G1 Access: who is unserved, and why ... 106
8.6 G2 Condition: a gap known only by network ... 114
8.7 G3 Capacity ... 114
8.8 G4 Standard ... 115
8.9 G5 Connectivity ... 118
8.10 G6 Sector connection ... 120
8.11 G7 Resilience ... 121
8.12 The evidence gap: survey before works ... 123
8.13 Data quality and limitations ... 124
8.14 Implications for planning ... 124
9 Costing: unit prices, maintenance and the backlog ... 125
9.1 What the module does ... 125
9.2 Method ... 125
9.3 Results: what the existing network costs to maintain ... 133
9.4 Results: the condition backlog ... 134
9.5 Results: which prices carry the programme ... 135
9.6 Data quality and limitations ... 136
9.7 Implications for planning ... 137
10 The investment programme ... 138
10.1 What the module does ... 138
10.2 Method ... 138
10.3 Results: the candidate programme ... 140
10.4 Results: appraisal and cost-effectiveness ... 147
10.5 From gaps to works: what the programme closes and what it leaves ... 149
10.6 Data quality and limitations ... 149
10.7 Implications for planning ... 150
11 Resource allocation: where the money goes ... 151

11.1 What the module does 151
11.2 Method 152
11.3 Results: the envelope and what is left of it 155
11.4 Results: preserve or expand? 159
11.5 Results: how many projects, and when 160
11.6 Results: where the money goes 161
11.7 Results: the optimiser 164
11.8 Results: the trade-off between access and condition 165
11.9 Results: capital against access 166
11.10 Data quality and limitations 167
11.11 Implications for planning 167
12 Outcomes: what each allocation achieves 169
12.1 What the module does 169
12.2 Method 169
12.3 Results: rural access at the horizon 171
12.4 Results: who gains 173
12.5 Results: access to services and markets 174
12.6 Results: where access improves 176
12.7 Results: the network left in poor condition 178
12.8 Results: why project gains do not add up 178
12.9 Results: what each allocation buys and gives up 179
12.10 Simulations: what changes the result 180
12.11 Data quality and limitations 187
12.12 Implications for planning 188
13 Budget and the medium-term expenditure framework 189
13.1 What the module does 189
13.2 The evidence: Uganda's approved roads budget 189
13.3 Method 192
13.4 Results: scenario B as a budget 193
13.5 Results: three views, one total 194
13.6 Results: can the network be maintained? 195
13.7 Results: where the budget and the maintenance fall 196
13.8 Results: who finances scenario B, and who maintains it 197
13.9 Outbound to the fiscal models 197
13.10 Data quality and limitations 197
13.11 Implications for planning 198
14 Integration with the sector and fiscal engines 199

14.1 The rules every exchange follows ... 199
14.2 What comes in ... 199
14.3 What goes out ... 201
14.4 The unit-cost link: evidence against UCSN ... 202
14.5 What integration means for the results ... 203
14.6 Implications for planning ... 203
15 The planning interface ... 204
15.1 The principle: five questions, every screen a map and a table ... 204
15.2 The pages ... 206
15.3 The costing studio ... 211
15.4 The road network on every map ... 212
15.5 Running the engine from the interface ... 213
15.6 District transport profiles ... 214
15.7 The rules enforced on screen ... 214
15.8 What is still missing ... 215
15.9 Implications for planning ... 215
16 Calibration and validation ... 216
16.1 Pre-run gates ... 216
16.2 The test suite ... 216
16.3 Validation checks on the Uganda run ... 218
16.4 Calibration: what has been compared with Uganda ... 219
16.5 What is and is not calibrated ... 220
16.6 Implications for planning ... 221
17 Reference catalogues: norms and unit costs ... 222
17.1 How norms are held ... 222
17.2 The norm set behind run 014c9cd7 ... 223
17.3 The tables one by one ... 224
17.4 The envelope: roads share and pool split ... 227
17.5 The unit-cost map ... 229
17.6 The catalogues (tpl_ref) ... 230
17.7 Data quality and limitations ... 231
17.8 Implications for planning ... 231
18 Build status against the build plan ... 232
18.1 What the build plan asks for ... 232
18.2 Status by phase ... 233
18.3 Deviations from the specification ... 235
18.4 The test suite ... 236

18.5 The Uganda run history 237
18.6 What is missing 239
18.7 Data quality and limitations 240
18.8 Implications for planning 240
19 The database 241
19.1 What the specification asks for 241
19.2 What is built: an SQLite store 241
19.3 Migrations 242
19.4 The audit tier 243
19.5 The norm tier 244
19.6 Observations 245
19.7 Result tables 245
19.8 Run storage under outputs/ 246
19.9 The prune command 246
19.10 How the store is PostGIS-ready 247
19.11 What is built, tested and missing 247
19.12 Data quality and limitations 248
19.13 Implications for planning 248
20 Existing data assets and acquisition 250
20.1 The estate at a glance 250
20.2 Sources acquired for the Uganda run 252
20.3 What the acquisitions changed 259
20.4 The two estate folders in depth 259
20.5 What only agencies can supply 260
20.6 Licences carried into outputs 261
20.7 Data quality and limitations 262
20.8 Implications for planning 262
21 GIS analysis across modules 264
21.1 What the GIS subsystem does 264
21.2 Method 266
21.3 Results: which layers cluster 267
21.4 Results: where the hot spots and cold spots are 271
21.5 Results: how much of the country is in a cluster 273
21.6 Results: how the programme and the scenarios cluster 274
21.7 Data quality and limitations 275
21.8 Implications for planning 276
22 Sector linkages: agriculture, minerals, health, education, industry and trade 277

22.1 What the module does ... 277
22.2 Method ... 278
22.3 The state of the sector feeds ... 279
22.4 Results: sector nodes without their required connection ... 280
22.5 Results: the candidate programme by sector ... 281
22.6 Results: funding by sector in each scenario ... 283
22.7 Results: what each sector gains, in its own unit ... 285
22.8 Sector by sector ... 286
22.9 What TIPRAS sends back ... 293
22.10 Data quality and limitations ... 293
22.11 Implications for planning ... 294
23 TransportModelling reuse ... 295
23.1 What TransportModelling is ... 295
23.2 What was ported, component by component ... 296
23.3 Seed norms from TransportModelling configuration ... 298
23.4 Tests and benchmarks ... 300
23.5 Requests back to TransportModelling ... 301
23.6 Data quality and limitations ... 301
23.7 Implications for planning ... 301
24 Mathematical models ... 303
24.1 What the chapter covers ... 303
24.2 Method: population and the planning grid (TI-01) ... 304
24.3 Method: network, link time and seasons (TI-04) ... 304
24.4 Method: accessibility (TI-02, TI-04, TI-09) ... 305
24.5 Method: movement generation and distribution (TI-02, TI-03) ... 306
24.6 Method: traffic loading, AADT and axle loads (TI-03) ... 307
24.7 Method: the required standard (TI-03) ... 308
24.8 Method: pavement deterioration (TI-04, TI-06) ... 309
24.9 Method: capacity and queues (TI-04, TI-05) ... 309
24.10 Method: the seven gap dimensions (TI-05) ... 310
24.11 Method: costing (TI-06) ... 311
24.12 Method: road user costs and appraisal (TI-07) ... 312
24.13 Method: criticality and resilience (TI-05) ... 313
24.14 Method: spatial statistics (GIS) ... 314
24.15 Method: resource allocation (TI-08) ... 314
24.16 Method: outcomes, mode choice and budget (TI-09, TI-07, TI-10) ... 316
24.17 Implementation map ... 317

24.18 Data quality and limitations 318
24.19 Implications for planning 318
A Annex: district summary 320
B Annex: glossary and abbreviations 325

List of figures

Figure 1.1: Settlement cells and population (175,367 cells, WorldPop 2020 carried to 2026) 23
Figure 1.2: Household wealth (International Wealth Index) of settlement cells 24
Figure 1.3: Uganda's 135 districts and the national road network by class 25
Figure 2.1: The planning chain: ten modules from the spatial base to the budget 31
Figure 2.2: System architecture: data tiers, engine, store and interface 34
Figure 3.1: The norm set by quality tier: what is evidence and what is still assumed 41
Figure 4.1: Population density by district; right: local Moran (LISA) clusters, 999 permutations 48
Figure 4.2: The road network in the engine, by functional class (all 248,000 km) 50
Figure 4.3: Network km by functional class and known surface 51
Figure 4.4: How much of each road class has a known surface 51
Figure 4.5: Road surface as known to the engine: national and district roads (unknown surface in grey) 53
Figure 4.6: Terrain class of every road link (from the DEM slope) 54
Figure 4.7: Hazard exposure of the network: flood-exposed links (JRC 100-year) and seasonal river crossings 56
Figure 4.8: Sector nodes by kind 57
Figure 4.9: Sector nodes the network must serve: health, schools, markets, mines, industry and gateways 58
Figure 4.10: Distance from each settlement to the nearest mapped road 59
Figure 5.1: People who must reach each destination class (base year) 65
Figure 5.2: People requiring each destination class, by district (each map on its own scale) 66
Figure 5.3: People requiring access to: All-season road (2 km), by district; right: LISA clusters 67
Figure 5.4: People requiring access to: EmONC (2 h), by district; right: LISA clusters 68
Figure 5.5: People requiring access to: Secondary school, by district; right: LISA clusters 69
Figure 5.6: People requiring access to: Market, by district; right: LISA clusters 69
Figure 6.1: Freight flows on the network: agricultural surplus to markets 76
Figure 6.2: Distribution of required-state traffic over loaded links 77
Figure 6.3: Required-state traffic (AADT) on the 20,000 busiest links, from the all-or-nothing assignment 78
Figure 6.4: Pavement loading: equivalent standard axles per year on the most loaded links 79
Figure 6.5: Which test sets the required standard on each link 80
Figure 6.6: Km whose required standard is above the current surface, by district; right: LISA clusters 81
Figure 7.1: Links closed in the wet season (season rules on surface, terrain and river crossings) 88
Figure 7.2: Rural access today: plan figure, upper bound from unknown surface, and the published benchmark 89
Figure 7.3: Rural Access Index today by district; right: LISA clusters 90
Figure 7.4: Distribution of the Rural Access Index across districts 90
Figure 7.5: The 25 districts with the lowest rural access 91
Figure 7.6: Travel time to primary health, wet season, from every settlement (hexagon median) 93

Figure 7.7: Travel time to EmONC, wet season, from every settlement (hexagon median)..... 94

Figure 7.8: Travel time to a secondary school, wet season, from every settlement (hexagon median) 95

Figure 7.9: Travel time to a market, wet season, from every settlement (hexagon median)..... 96

Figure 7.10: Travel time to the district centre, wet season, from every settlement (hexagon median) 97

Figure 7.11: Distance to an all-season road, wet season, from every settlement (hexagon median)..... 98

Figure 7.12: Share of people reached within a given travel time, by destination class (wet season) 99

Figure 7.13: Road condition by network: agency-reported km in poor condition..... 100

Figure 8.1: Why people are unserved: the G1 access gap decomposed by cause..... 107

Figure 8.2: G1 Access: rural people beyond 2 km of an all-season road — by district; right: LISA clusters (this dimension only) 108

Figure 8.3: Top 20 districts: Rural people beyond 2 km of an all-season road (one dimension, one sort column) 109

Figure 8.4: G1 Access: people beyond the primary-health standard — by district; right: LISA clusters (this dimension only) 110

Figure 8.5: G1 Access: people beyond 2 h of EmONC — by district; right: LISA clusters (this dimension only) ... 110

Figure 8.6: Top 20 districts: People beyond 2 h of EmONC (one dimension, one sort column) 111

Figure 8.7: G1 Access: people beyond the primary-school standard — by district; right: LISA clusters (this dimension only) 112

Figure 8.8: G1 Access: people beyond the secondary-school standard — by district; right: LISA clusters (this dimension only) 112

Figure 8.9: G1 Access: people beyond the market standard — by district; right: LISA clusters (this dimension only) 113

Figure 8.10: G1 Access: people beyond the administrative-centre standard — by district; right: LISA clusters (this dimension only) 114

Figure 8.11: G3 Capacity: links over the volume/capacity threshold — by district; right: LISA clusters (this dimension only) 115

Figure 8.12: G4 Standard: km below the required surface — by district; right: LISA clusters (this dimension only) 116

Figure 8.13: G4 Standard: km not passable all season — by district; right: LISA clusters (this dimension only) .. 117

Figure 8.14: Top 20 districts: Km not passable all season (one dimension, one sort column) 117

Figure 8.15: G5 Connectivity: cut edges isolating a community — by district; right: LISA clusters (this dimension only)..... 118

Figure 8.16: G5 Connectivity: largest population behind a single cut edge — by district; right: LISA clusters (this dimension only) 119

Figure 8.17: Top 20 districts: Largest population behind a single cut edge (one dimension, one sort column) .. 119

Figure 8.18: G6: sector nodes lacking the connection standard their sector requires 121

Figure 8.19: G7 Resilience: people behind flood-exposed critical links — by district; right: LISA clusters (this dimension only) 122

Figure 8.20: Top 20 districts: People behind flood-exposed critical links (one dimension, one sort column) 122

Figure 8.21: The evidence gap: gap km with unknown surface and river crossings with unknown structure, by class 123

Figure 9.1: The price indices used to carry evidence to 2026 prices..... 127

Figure 9.2: Observations behind each evidence price, by evidence tier..... 128

Figure 9.3: Every evidence observation converted to UGX 2026, by price code (log scale) 129

Figure 9.4: Evidence unit prices (low · central · high) against the UCSN illustrative prices 130

Figure 9.5: Routine maintenance requirement of the existing network, by district; right: LISA clusters 134

Figure 9.6: The condition backlog: cost of treating the agency-reported poor km (low · central · high) 135

Figure 10.1: Capital by works type (low · central · high) beside its induced routine maintenance 141

Figure 10.2: Km of works in the candidate programme, by works type 141

Figure 10.3: The works the decision rule assigns to each link (spot improvement, rehabilitation, upgrading) ... 143

Figure 10.4: Candidate programme capital by district (left) and the routine maintenance it induces (right)..... 144

Figure 10.5: The 25 districts with the largest candidate programme, stacked by works type 144

Figure 10.6: Candidate programme capital per inhabitant, by district 146

Figure 10.7: Appraisal: net present value per project at a 12 % discount rate (log count)..... 147

Figure 10.8: Capital against rural-access gain, project by project..... 148

Figure 11.1: Capital allocated by funding pool in each scenario (horizon total)..... 156

Figure 11.2: Capital allocated by works type in each scenario: preserve or expand? 159

Figure 11.3: Projects funded in each scenario 160

Figure 11.4: Capital disbursed by year within the horizon, by scenario 161

Figure 11.5: Capital allocation per inhabitant by district in every scenario — one shared colour scale 162

Figure 11.6: Where scenario C (rural access) and scenario B (preserve first) differ: allocation per capita, C − B 163

Figure 11.7: The trade-off frontier: people served against poor network treated..... 165

Figure 11.8: Capital allocated against rural access achieved: every scenario..... 166

Figure 12.1: Rural Access Index at the horizon in every scenario (no scenario is recommended) 171

Figure 12.2: Who gains: rural access by household wealth quintile (IWI) in each scenario 173

Figure 12.3: The wealth gap in access: richest minus poorest quintile, by scenario (points) 174

Figure 12.4: People newly meeting each access standard against doing nothing, by scenario (each class in its own panel)..... 175

Figure 12.5: Rural Access Index at the horizon by district in every scenario — one shared colour scale..... 176

Figure 12.6: Gain in rural access against doing nothing, by district and scenario (shared scale) 177

Figure 12.7: Agency-reported poor km left untreated at the horizon, by scenario..... 178

Figure 12.8: Interaction adjustment: how far stand-alone project gains overstate the network result 179

Figure 12.9: Simulations: Rural Access Index at the horizon for every scenario under each simulation..... 183

Figure 12.10: Simulations × scenarios: RAI at the horizon (%) in every combination..... 183

Figure 12.11: Simulations: the candidate programme's capital (low · central · high) and induced maintenance 184

Figure 12.12: Simulations: capital allocated and projects funded under scenario B and scenario C..... 184

Figure 12.13: What moves the result: change in scenario C's RAI and in programme capital against the base simulation S0 185

Figure 12.14: Simulations: rural access of the poorest and richest quintiles under scenario C 185

Figure 12.15: Simulations: Rural Access Index at the horizon under scenario C, by district — one shared scale 186

Figure 12.16: Simulations: where a 30 % larger or smaller envelope changes scenario C's allocation per capita 187

Figure 13.1: The approved roads budget by funding pool, FY2024/25–FY2026/27 190

Figure 13.2: Scenario B in the national chart of accounts: capital by programme code and fiscal year 193

Figure 13.3: The three budget views of scenario B reconcile to the same total 194

Figure 13.4: Can the network be maintained? Routine maintenance requirement against maintenance-eligible resources (steady state) 195

Figure 13.5: Scenario B capital budget by district (left) and the routine maintenance requirement of the existing network (right) 196

Figure 15.1: The planning interface: the Overview page (KPIs with quality chips, access map, scenario charts) 205

Figure 15.2: Q5 Where resources go: scenarios side by side on one scale 207

Figure 15.3: The GIS explorer: any district layer, histogram, Moran scatter and Moran's I per layer 209

Figure 15.4: Q4 Cost: capital by district always beside its induced maintenance 210

Figure 15.5: The costing studio: cost roads from choices of scope, standard, assumptions and unit prices 211

Figure 15.6: Q2 What we have: rural access with the road network and wet-season closures 213

Figure 20.1: Sources assessed for the engine, by acquisition status 252

Figure 21.1: Global Moran's I of every district layer (blue: significant at 5 %) 268

Figure 21.2: Hot spots and cold spots of the key gaps (LISA, queen contiguity, 999 permutations) 271

Figure 21.3: How many districts sit in each LISA cluster, layer by layer (significant layers) 274

Figure 22.1: The candidate programme by sector served (a project can serve several sectors) 281

Figure 22.2: Capital funded per sector served, by scenario (bn UGX; rows are not additive) 283

Figure 22.3: Agriculture: markets over the market access gap (each sector's gap in its own unit) 287

Figure 22.4: Minerals: mines over km below the required standard (each sector's gap in its own unit) 288

Figure 22.5: Health: facilities over the primary-health access gap (each sector's gap in its own unit) 289

Figure 22.6: Education: schools over the secondary-school access gap (each sector's gap in its own unit) 291

Figure 22.7: Industry: industrial zones over km below the required standard (each sector's gap in its own unit) 292

List of tables

Table 1.1: How Uganda's roads sector is organised and financed, as reflected in the budget evidence 26
Table 1.2: The five planning questions, the modules that answer them and where this report reports them..... 27
Table 2.1: The ten modules of the planning chain 31
Table 2.2: The GIS subsystem as built (src/GIS) 32
Table 2.3: Technology stack, specified and as built 34
Table 2.4: Pins of the reference run 014c9cd7 35
Table 2.5: Performance targets and the reference run..... 36
Table 2.6: Deliberate departures from the design, as of September 2026..... 37
Table 3.1: The seven schemas, their mutability and the Uganda build 38
Table 3.2: Spatial keys and where their geometry lives..... 39
Table 3.3: The quality tier: how strong is the evidence? 42
Table 3.4: The basis: how did the value arise?..... 42
Table 3.5: Missing-data behaviour and its effect on the Uganda run..... 44
Table 4.1: Road network in the geography snapshot: km by functional class and surface 49
Table 4.2: Quality of the TI-01 spatial base in the reference run..... 60
Table 5.1: Access standards in norm version UGA-2026.09-draft-f 61
Table 5.2: Node connection standards in the reference run 63
Table 5.3: People who must be connected to each destination class, base year 64
Table 6.1: Traffic bands of the design standard (assumed norms, to be replaced by the national road design manual)..... 74
Table 6.2: Which test sets the required standard on each link (TI-03) 80
Table 6.3: Quality of the TI-03 results in the reference run..... 82
Table 7.1: Access standards used in the reference run (rural / urban)..... 86
Table 7.2: Rural Access Index today: plan figure, upper bound and published benchmark 89
Table 7.3: The 15 districts with the lowest Rural Access Index today..... 91
Table 7.4: Road condition by network, agency statistics 100
Table 8.1: The seven gap dimensions, their units and the instruments that close them 103
Table 8.2: The gap register: every dimension in its own unit (never combined)..... 105
Table 8.3: Access gap by cause (millions of people) 107
Table 8.4: Structures to survey (G5/G7 evidence gap) 120
Table 8.5: G6: sector nodes that lack their required connection standard 121
Table 8.6: Survey requirement: gap links whose surface is unknown 123
Table 9.1: How many documented observations stand behind each price 128
Table 9.2: Evidence unit prices, UGX 2026..... 129

Table 9.3: Unit cost catalogue applied in the reference run (map key: works, from-surface, to-surface, terrain) 131
Table 9.4: Works that remain unpriced (never funded) 132
Table 9.5: Condition backlog: the cost of treating the agency-reported poor km 135
Table 9.6: Programme capital by the price code that priced it, central values, UGX 2026 135
Table 10.1: The candidate programme: capital by works type, always with its induced routine maintenance 140
Table 10.2: Upgrading by target surface: capital beside induced routine maintenance, central, UGX 2026 142
Table 10.3: The 15 districts with the largest candidate programme 145
Table 10.4: How the candidate programme answers each gap dimension in the reference run 149
Table 11.1: Pools, their shares of the roads envelope and the works they may fund 152
Table 11.2: The seven rule-based scenarios 153
Table 11.3: Capital allocated by funding pool and scenario, horizon total 2027–2030 (bn UGX) 157
Table 11.4: Money left unallocated over 2027–2030, by scenario (bn UGX) 158
Table 11.5: The optimiser runs: one objective each, other measures as ε-floors 164
Table 11.6: Trade-off frontier: rural access against poor network treated 165
Table 12.1: Scenario comparison at the horizon (2030): each outcome in its own unit, compared, never scored 172
Table 12.2: Persons newly meeting each access standard against doing nothing, 2030 (thousand persons) 175
Table 12.3: What each allocation buys and gives up, 2030 (each column in its own unit) 179
Table 12.4: The simulation set: one change at a time from the evidence base case 180
Table 12.5: Simulation results: programme cost and access outcomes at the horizon 182
Table 12.6: Simulation results: what scenarios B and C fund 183
Table 13.1: Transport and roads in the approved budget estimates (bn UGX) 190
Table 13.2: Approved roads budget by funding pool, FY2024/25–FY2026/27 (bn UGX) 190
Table 13.3: Pool shares used for the envelope (three-year pooled, quality normed) 191
Table 13.4: Scenario B capital by programme code and fiscal year (bn UGX) 193
Table 13.5: Scenario B budget by fiscal year against the envelope (bn UGX) 193
Table 13.6: Scenario B budget: the three views reconcile 194
Table 13.7: Fiscal sustainability statement, scenario B (steady state) 195
Table 14.1: Inbound links and their state in the reference run 199
Table 14.2: Outbound files written by TI-10 201
Table 14.3: TIPRAS evidence prices against the UCSN illustrative coefficients (UGX 2026, central values) 202
Table 15.1: Pages of the planning interface 206
Table 15.2: The interface's hard rules and how they are met 214
Table 16.1: Pre-run gates and their behaviour in the reference run 216
Table 16.2: The test suite by layer 217
Table 16.3: Validation checks on the reference run 218

Table 16.4: Calibration status ... 220
Table 17.1: Norm versions in the store ... 222
Table 17.2: The norm set UGA-2026.09-draft-f: tables and quality ... 223
Table 17.3: Norm rows of UGA-2026.09-draft-f by origin and quality ... 224
Table 17.4: Access standards in UGA-2026.09-draft-f ... 224
Table 17.5: Traffic bands of the design standard (illustrative) ... 225
Table 17.6: Project overheads on base works cost (illustrative feasibility-stage rates) ... 227
Table 17.7: Roads share and pool split from approved budgets ... 228
Table 17.8: Unit-cost map of UGA-2026.09-draft-f and the prices it resolves to (UGX m, 2026) ... 229
Table 18.1: The nine build phases and their gates (docs/17) ... 232
Table 18.2: Build status by phase, with the gate result ... 233
Table 18.3: Deliberate deviations from the specification and how to close them ... 236
Table 18.4: The testsets in test/runtests.jl ... 236
Table 18.5: Uganda runs 1–8: inputs, run time and headline result ... 237
Table 18.6: Run 8 (014c9cd7) headline results and their quality ... 238
Table 18.7: Missing items, by kind, with what would close each ... 239
Table 19.1: The database specified in docs/18 Part A ... 241
Table 19.2: The database deviation as recorded ... 242
Table 19.3: Specified schemas and their form in the SQLite store ... 242
Table 19.4: Migration files ... 242
Table 19.5: Price vintages in the store ... 245
Table 19.6: Result tables ... 245
Table 19.7: Database status against docs/18 Part A ... 247
Table 20.1: The six parts of the estate and their Uganda state (verified 2026-09-23) ... 250
Table 20.2: Sources acquired for the Uganda run, by status ... 254
Table 20.3: Why sources could not be acquired by script ... 258
Table 20.4: Gaps of docs/19 §19.12 and what now fills them ... 259
Table 20.5: Verdicts of the gap assessment for the two estate folders ... 259
Table 20.6: Data still missing after the open-source search, and the holder ... 260
Table 20.7: Draft data-request letters (drafted, not sent) ... 261
Table 20.8: Licences and what they require of published outputs ... 262
Table 21.1: Where each kind of spatial work is done ... 264
Table 21.2: Core spatial operations in `src/GIS/` ... 265
Table 21.3: Global Moran's I of the district layers, with the number of districts in each LISA class ... 269
Table 22.1: What TIPRAS consumes from and returns to each sector (design, docs/21 §21.2) ... 278
Table 22.2: What each sector feed supplied to the reference run ... 280
Table 22.3: Candidate programme by sector served ... 281

Table 22.4: Candidate projects by the number of service sectors they serve 282

Table 22.5: Capital funded per sector served, by scenario (bn UGX; a project can serve several sectors, so rows are not additive) 284

Table 22.6: Induced routine maintenance of funded projects, by sector tag and scenario (bn UGX a year at steady state; rows overlap) 285

Table 22.7: Gain over the do-nothing path at the 2030 horizon, by destination class (each in its own unit) 286

Table 22.8: Outbound files written for the sector engines 293

Table 23.1: The five reuse verdicts of docs/22 §22.1 295

Table 23.2: TransportModelling components: verdict and state in TIPRAS 296

Table 23.3: Items marked "do not copy" and their state 298

Table 23.4: Norm rows seeded from TransportModelling configuration (draft-f) 299

Table 23.5: Requests back to TransportModelling and their state 301

Table 24.1: Main sets and quantities used in this chapter 303

Table 24.2: Implementation status of the mathematical models in the Uganda reference run 317

Table A.1: District summary: access, gaps and the candidate programme 321

Table B.1: Abbreviations 325

Table B.2: Glossary of technical terms 326

Executive summary

This report sets out the status of Uganda's road network — what the country needs from it, what it has, where the gaps are, what it would cost to close them and where the money could go — and what transport does for agriculture, minerals, health, education and industry. Every number is computed by the AICTEM Transport Infrastructure Planning and Resource Allocation System (TIPRAS) from its data layers: the SDI road network of 817,354 links (248,000 km), 175,367 settlement cells, 24,952 sector nodes, agency condition statistics, approved budget estimates and 275 documented unit-cost observations. The reference run (014c9cd7) is **exploratory**: its norms are a draft and its unit prices are evidence candidates awaiting review. Every figure carries its quality tier, and the reader should keep the tier in view.

The needs

- **Everyone needs an all-season road; many need more.** The planning population is 52.0 million in 2026. All of them should be within reach of an all-season road and a primary health facility; 12.2 million need emergency obstetric care (EmONC) within two hours, 10.9 million a primary school, 7.8 million a secondary school and 11.6 million a market (Chapter 5).
- **The standard each road must meet is set by its function, not only its traffic.** Of 817,382 links tested, 656,479 take their required standard from their function in the network, 160,189 from traffic and 686 from the sector node they serve (Chapter 6). Measured against those standards, 242,068 km sit below the surface they require — almost all of them earth or of unknown surface.

What the network has

- **A long network, little of it paved, most of it unknown.** Of 248,000 km, 7,269 km are paved, 1,972 km gravel and 69,388 km earth by their mapped surface. The surface of **68 % of the network is not recorded** in any source the engine can read (Chapter 4). This single fact bounds most results in the report.
- **Rural access is between 28.7 % and 93.5 %.** Counting unknown surface as earth, 28.7 % of rural people live within 2 km of an all-season road (the plan figure); counting it as gravel, 93.5 %. The published SDG 9.1.1 value is 53.1 % (2015). The width of this range is a **survey requirement, not a finding** (Chapter 7). District values run from 0 % (Buvuma, Kalangala) to 83 % (Manafwa, Kampala).
- **Condition is poor and known only by network.** The road agencies report 139,691 km in poor condition, including 70,656 of 75,404 km of community roads and 35,320 of 41,194 km of district roads, against 204 of 6,133 km of paved national roads. No link-level condition survey exists.
- **5,013 links close in the wet season**, 3,006 links (1,368 km) are exposed to the 100-year flood and 2,248 river crossings are seasonal.

The gaps — seven dimensions, never combined

Each gap is measured in its own unit and mapped on its own (Chapter 8):

- **G1 Access:** 33.3 million rural people are beyond 2 km of an all-season road — for 32.7 million the network reaches them but is below standard. 3.2 million are beyond the primary-health standard, 1.0 million beyond two hours of EmONC (Kotido the largest), 0.9 million beyond the secondary-school standard and 2.3 million beyond a market.
- **G2 Condition:** 139,691 agency-reported poor km (link-level condition unknown).
- **G4 Standard:** 242,068 km below the required surface.

- **G5 Connectivity:** single cut edges isolate communities across the network; the largest population behind one cut edge is reported by district.
- **G6 Sector service:** 103 of 117 health nodes with a connection standard, 21 of 23 industrial zones, 16 of 28 trade gateways and all 10 agricultural surplus zones (9.8 Mt a year) lack the connection their sector requires.
- **G7 Resilience:** people behind flood-exposed critical links, mapped by district.
- **The evidence gap:** 169,343 km of gap links have unknown surface and 1,958 river crossings an unknown structure. They are listed for survey, not for works.

What it costs

- **The identified programme costs UGX 92.8 tn (range 56.6–125.3 tn), and adds UGX 205 bn a year of routine maintenance.** 41,867 of 41,890 candidate projects are priced: 41,757 upgrades (72,725 km, 85.5 tn), 72 rehabilitation programmes (70,860 km, 6.2 tn) and 38 periodic-maintenance programmes (36,977 km, 1.1 tn) (Chapters 9–10).
- **Prices come from documents, carried to 2026 prices.** Fourteen evidence prices were built from 275 extracted observations (89 used), converted with World Bank deflators and exchange rates: engineered earth 42 M UGX/km, gravel 75 M, earth-road rehabilitation 66 M, gravel rehabilitation 59 M, low-volume seal 735 M and paved upgrading 4,269 M per km. They are E5 evidence candidates; bridges and spot improvements rest on one observation each.
- **One known overstatement.** The unit-cost map has no code for upgrading one paved surface to another, so about 2,500 km of surface dressing taken to asphalt is priced at the full paving rate — about 13 tn of the 92.8 tn. An overlay costs much less.
- **The condition backlog alone costs 7.3 tn**, 5.6 tn of it on community roads.
- **Maintenance is the binding constraint.** Keeping the existing network costs 955 bn a year; the plan adds 63 bn a year once built. Maintenance-eligible resources are 592 bn a year: a **funding ratio of 0.58** (Chapter 13).

Where resources go

- **Eight allocation rules, compared, none recommended** (Chapters 11–12). With the evidence envelope (roads 54.5 % of the MTEF transport ceiling), the rural-access rule (C) funds 8,611 projects for 11.3 tn and raises the Rural Access Index from 28.7 % to **71.5 %** by 2030; preserve-first (B) funds 40,486 projects for 17.1 tn and reaches 69.8 % while treating almost all the poor network; the current-trajectory fallback (A) reaches 37.4 %; productive sectors (D) 44.6 %; social services (E) 46.5 %; regional equity (F) 70.7 %; climate resilience (G) 40.8 %.
- **Access gains reach the poorest.** Under C the poorest wealth quintile rises from 10.0 % to 55.8 % and the richest from 52.1 % to 85.3 %.
- **Condition costs little in access.** The optimiser that maximises people served while treating at least half of the poor network serves 20.87 M against 20.90 M unconstrained.
- **Unspent money is a timing effect.** Under several rules part of the envelope is not spent in 2027 and 2030 because three-year works phase their cash and must finish by the horizon, not because works are unpriced.

Simulations

Seven simulations each change one input of the reference configuration (Chapter 12):

- **The access outcome is robust; the cost is not.** Prices at their low and high bounds move the programme from 56.6 tn to 125.3 tn, but the RAI under the rural-access rule only from 70.7 % to 72.6 %. An envelope 30 % smaller or larger moves it from 70.7 % to 72.1 %: 3.4 tn more buys 0.6 points of access and mostly buys condition and standard.
- **Maintenance affordability is the most price-sensitive result.** The funding ratio is 0.58 at central prices, 1.98 at low and 0.30 at high prices, because routine maintenance of unpaved roads is documented between 1.0 and 8.9 M UGX per km a year. A Road Fund unit-rate series would narrow this more than any other dataset.
- **The unknown surface is the largest uncertainty.** Treating the 169,343 km of gap links with unknown surface as earth grows the programme to 153.7 tn and lets the same 11.3 tn under C lift the RAI to 96 %: much of the access gap is a survey gap.
- **Pricing the works UCSN does not cover is worth 7 points of access.** With the UCSN-first price map, engineered-earth and earth-rehabilitation works are unpriced; C reaches 64.5 % instead of 71.5 % and B 59.2 % instead of 69.8 %.
- **The grouping of the rules is stable.** In every simulation the rural-access, regional-equity and preserve-first rules lie within 1–5 points of each other and well above the other four.

What the evidence is, and is not

- **Measured:** network geometry, settlements and population, sector nodes, terrain and hazards, the budget, agency condition statistics.
- **Modelled:** travel times, traffic, required standards, gaps, programme, allocations and outcomes.
- **Assumed:** 578 of 646 norm rows, including speeds and the value of time; the TransportModelling vehicle-operating-cost coefficients are uncalibrated.
- **Missing:** surface on 68 % of km, link-level condition, traffic counts, mine and industrial output, and Ugandan historical allocation shares. The data requests drafted in `docs/requests` would close most of these gaps; the first investment the evidence calls for is a **road inventory and condition survey**.

CHAPTER 1

1 Introduction

This report sets out the state of Uganda's transport infrastructure in 2026 as seen by TIPRAS, the Transport Infrastructure Planning and Resource Allocation System. It is written for the planners of the Ministry of Works and Transport (MoWT), the budget analysts of the Ministry of Finance, Planning and Economic Development (MoFPED), the local governments that run district and community roads, and the development partners who co-finance the sector. It does not report a plan that has been adopted. It reports what an engine built on the evidence available in September 2026 says about needs, supply, gaps, costs and the choices open to the Government, and how firm each of those statements is. The method follows the design manual (`docs/00_SCOPE_AND_PRINCIPLES.md` and the chapters that follow it). Every number comes from the reference run `014c9cd7` (geography `UGA-2026.09-c`, norms `UGA-2026.09-draft-f`, prices `UCSN+EV-UGA-2026.09b`, base year 2026, horizon 2030), from the agency statistics it reads, or from the evidence files listed in Chapter 20.

> The whole reference run is **exploratory**. Its norms are drafts, 578 of its 646 norm rows are assumed, its unit prices are evidence candidates rather than approved estimates, and the surface of 68 % of the network is unknown. The report shows what the engine can already say, and what better data would change.

1.1 Why a needs-based planning engine

Transport planning in Uganda, as in most countries, usually starts from one of two places. The first is money: the size of the road fund, the medium-term expenditure ceiling, or a share of gross domestic product. The second is a list: the pipeline of named roads that agencies and partners are already preparing. Both are legitimate inputs to a budget. Neither tells a ministry what the country needs.

TIPRAS starts somewhere else. It asks five questions in a fixed order, and it lets money enter only at the fourth and fifth. The design manual puts the question like this:

> Given where people live, where things are produced, where services are delivered, the network that exists and the money actually available — what transport infrastructure is required, where are the gaps, what will closing them cost, and where should resources go to connect people to services and producers to markets?

The consequence is a rule the engine keeps throughout: **never start from money, and never start from a project list**. Needs are computed from the population, the sector nodes that must be connected and the standard of connection each requires. The budget enters at the allocation module (TI-08) as a constraint. The ministry's pipeline enters at the investment module (TI-07), where it is netted out of the gap, not before.

TIPRAS sits beside two existing systems in the EDIP estate. The Africa Transport Data and GIS Repository (ATDGR) answers *what exists*: roads, bridges, rail, ports, airports, borders and logistics nodes, with their provenance. TransportModelling answers *how the network behaves*: deterioration, vehicle operating cost, traffic assignment, alignment and freight flows. TIPRAS consumes both. It does not harvest data and it does not re-implement a pavement research model or a four-step travel demand model. It does normative planning arithmetic on a network: needs, standards, gaps, costs and allocation, with a map behind every table.

1.2 Uganda in 2026: the planning context

1.2.1 People

The engine's planning population in the base year is **52.0 million persons**, spread over 175,367 settlement cells in 135 districts. This figure is not a census count. It is the UBOS district projection (CODPS 2022) carried forward at an assumed growth rate of 3 % a year, distributed within each district by the WorldPop 2020 constrained 100 m grid, which itself sums to 42.2 million. The difference between the two (a national raking factor of 1.23) is examined in Chapter 4. About 20 % of the planning population lives in cells the engine classes as urban.

Figure 1.1 shows where people live. Population per 10 km hexagon spans more than four orders of magnitude. The darkest cells are around Kampala and Wakiso, on the slopes of Mount Elgon in the east (Mbale, Bududa, Manafwa, Namisindwa) and in a band through the south-west. The sparsest areas are Karamoja in the north-east, the national parks and the forest reserves, which show as gaps in the hexagon field. Wakiso alone holds 3.9 million planning persons and Kampala 2.0 million. The ten most populous districts hold 24 % of the national total.

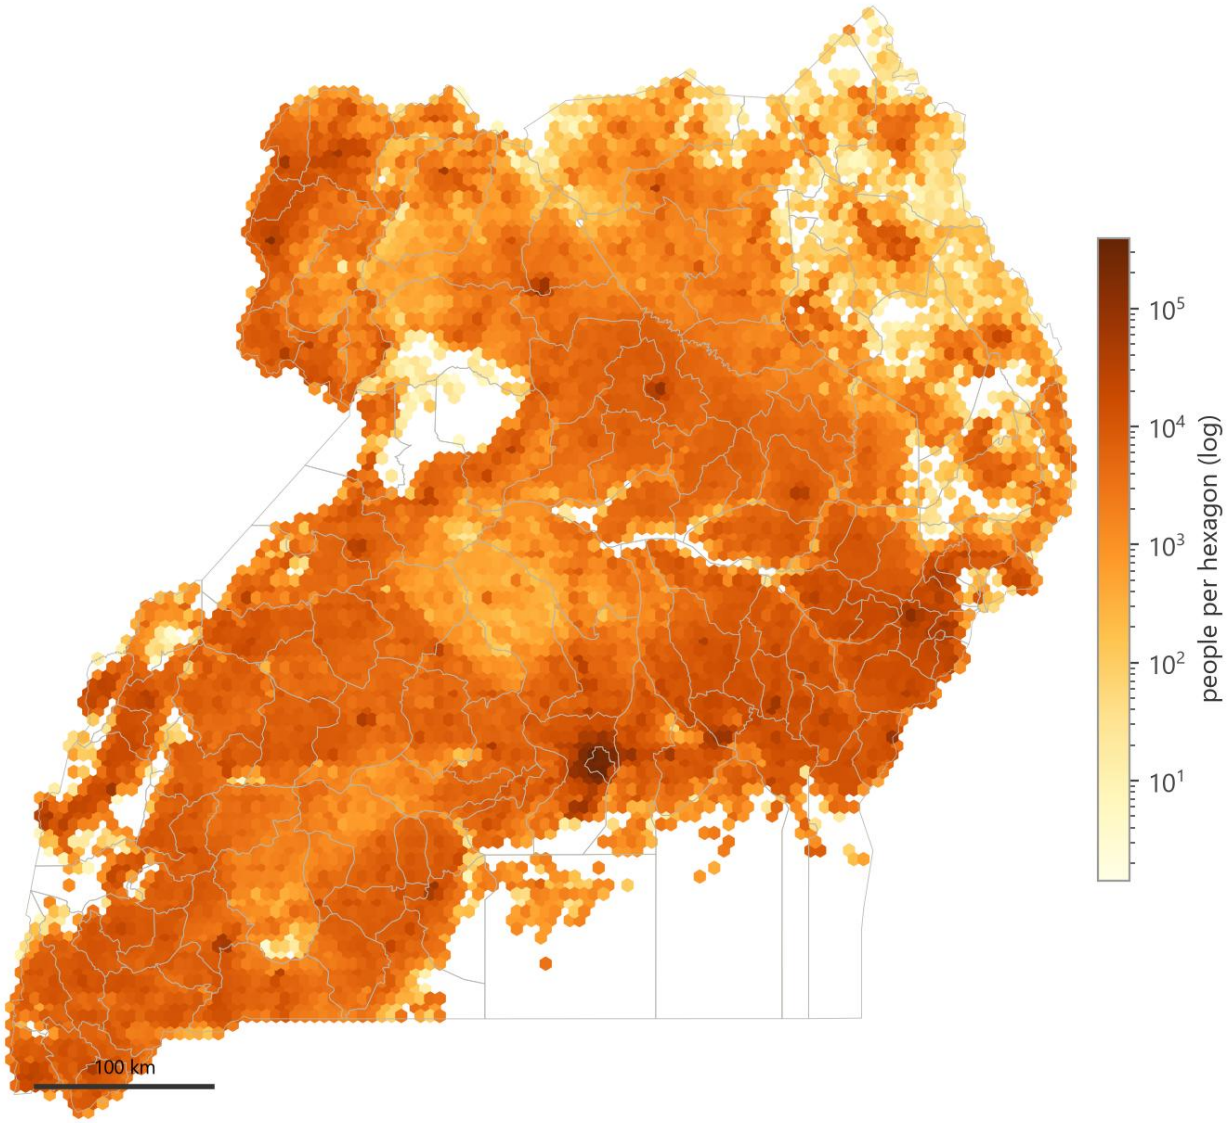

Figure 1.1: Settlement cells and population (175,367 cells, WorldPop 2020 carried to 2026)

Source: WorldPop 2020 constrained 100 m; UBOS projections

Figure 1.2 shows household wealth, measured by the International Wealth Index (IWI, 0–100) of each settlement cell. The district mean ranges from 13.9 in Otuke and 14.3 in Napak to 49.7 in Wakiso and 63.1 in Kampala. The

national median district scores 22.1. Wealth falls steeply with distance from Kampala and is lowest in Karamoja, in Acholi and Lango, and along the north-eastern border, where most hexagons sit below 15. The engine uses the IWI only as an equity weight: Chapter 12 reports access by IWI quintile. It is not an official poverty statistic, and the engine never combines it with a gap into a score.

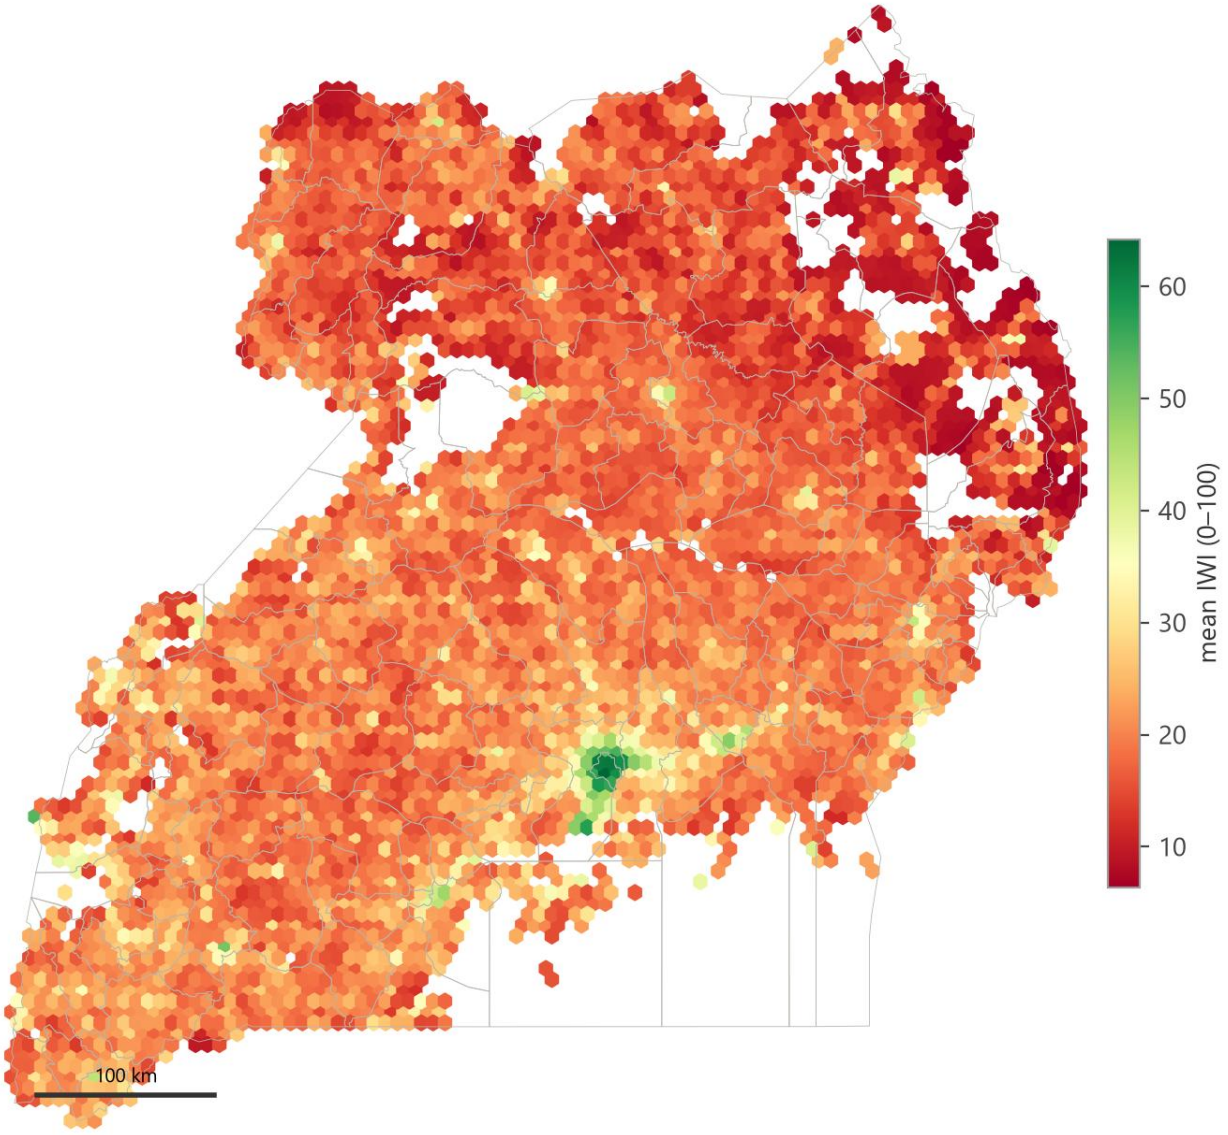

Figure 1.2: Household wealth (International Wealth Index) of settlement cells

Source: Global Data Lab IWI / Meta RWI (SDI)

1.2.2 The network

The engine's road network has **247,972 km of road** in 817,354 links, plus 28 ferry links. It is the SDI noded network, built from OpenStreetMap and classified into functional classes. Figure 1.3 shows its upper classes against the 135 district boundaries:

- 981 km of national trunk roads, the Northern Corridor through Kampala and the routes to the western and south-western borders;
- 3,415 km of national primary roads;
- 11,140 km of national secondary roads;
- 22,052 km of district roads.

Below these sit 160,842 km of community roads, 46,117 km of tracks, 2,634 km of service roads, 67 km of urban local roads and 723 km of unclassified links. The engine's network is therefore longer than the official network, because it contains every mapped path, but its national and district classes are shorter than the gazetted ones. Agency statistics for FY2022/23 report 21,200 km of national roads (6,133 km paved and 15,067 km unpaved) and 41,194 km of district roads. Chapter 4 explains the difference: the functional classes come from map tags, not from the gazette.

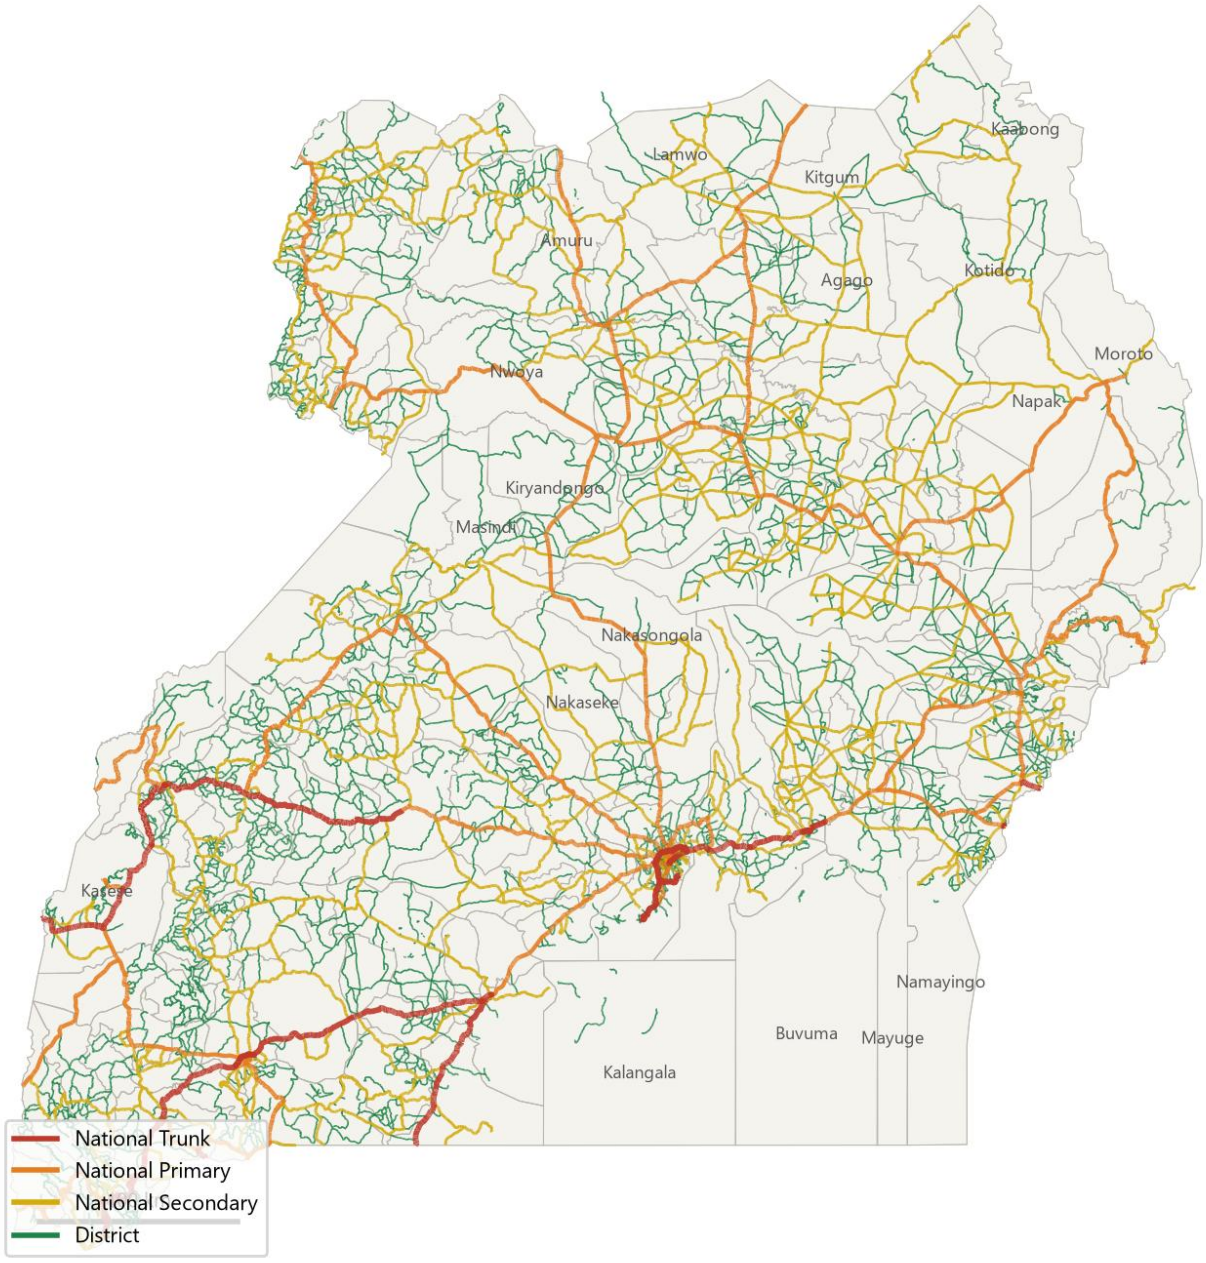

Figure 1.3: Uganda's 135 districts and the national road network by class

Source: Boundaries COD-AB; network: SDI noded network with OSM attributes

1.2.3 Surface, access and condition

What the engine knows about the network is uneven. Surface is recorded on only **31.7 %** of network km, mostly on national roads. Of the known km, 7,269 km are paved, 1,972 km are gravel and 69,388 km are earth. The remaining 169,343 km have no surface record at all. No link has a recorded condition.

These gaps drive the most important access figure in the report. The Rural Access Index (RAI), the share of rural people living within 2 km of an all-season road (SDG indicator 9.1.1), lies between **28.7 %** and **93.5 %** in the wet season. The lower bound treats every road of unknown surface as earth, which is closed in the wet season. The upper bound treats it as gravel, which is open. The World Bank's published estimate for Uganda is 53.1 % (2015). The engine cannot narrow this range without a surface survey. Chapter 7 shows the full distribution by district.

Condition is known only as agency statistics by network, not per link. Table 1.1 below lists the agencies; the statistics they report for FY2022/23 put **139,691 km** of public road in poor condition out of 173,657 km reported, or 80 %. The share differs widely between networks:

- national paved roads: 204 of 6,133 km (3 %);
- national unpaved roads: 1,657 of 15,067 km (11 %);
- district roads: 35,320 of 41,194 km (86 %);
- community access roads: 70,656 of 75,404 km (94 %);
- urban roads (KCCA, city, municipal and town council roads together): 31,854 of 35,858 km (89 %).

The national network, in other words, is in reasonable condition. The local networks that carry people to schools, clinics and markets are not.

1.2.4 Money

The approved budget estimates put programme 09, Integrated Transport Infrastructure and Services (ITIS), at UGX 4,968.2 bn in FY2024/25, 6,602.1 bn in FY2025/26 and 8,827.2 bn in FY2026/27. Of these, the lines the engine classes as roads are 3,005.7 bn, 3,572.1 bn and 4,813.6 bn. The roads share of transport in FY2026/27 is 0.545. External financing is more than half of the roads budget in the first two years (0.544 and 0.561 of the pool base). The routine maintenance requirement of the existing network, priced at evidence-candidate unit costs, is UGX 954.7 bn a year. With the induced maintenance of the assets the plan would add, it rises to 1,017.4 bn a year against 592.1 bn of maintenance-eligible resources in the final plan year: a maintenance funding ratio of **0.58**. Chapters 9 and 13 return to this.

1.3 How the roads sector is organised

The engine has to map every result onto the institutions that would act on it. Table 1.1 summarises them, as they appear in the approved budgets of FY2024/25 to FY2026/27.

Table 1.1: How Uganda's roads sector is organised and financed, as reflected in the budget evidence

Institution	Role in the roads sector	How it appears in the engine
Ministry of Works and Transport (MoWT, vote 016)	Policy, standards and regulation for the whole transport sector; since the RAPEX reform, also the national roads functions formerly run by UNRA	National trunk, primary and secondary roads; the government development and recurrent pools
Uganda National Roads Authority (UNRA, former vote 113)	Built and maintained the national road network; rationalised into MoWT under the RAPEX reform in 2024	A separate vote in FY2024/25 only; from FY2025/26 its lines appear as MoWT sub-subprogrammes "National Roads Maintenance & Construction"
Uganda Road Fund (URF, vote 118)	Finances road maintenance and transfers maintenance money to local road agencies	Road-fund maintenance pool; in the budget evidence, mainstreamed into MoWT from FY2025/26
Local governments: district, urban and community access roads (DUCAR, vote 609)	Maintain and improve district, urban and community access roads through grants and road-fund transfers	District-grant pool (DUCAR maintenance transfer of 176.0 bn, plus URF transfers of about 100.8 bn a year)
Kampala Capital City Authority (KCCA, vote 122)	Roads in the capital city	Urban network (2,103 km in the agency statistics)

Institution	Role in the roads sector	How it appears in the engine
Ministry of Finance, Planning and Economic Development (MoFPED)	Sets the medium-term expenditure framework (MTEF) ceilings, appropriates the budget by programme and vote	The envelope at TI-08 and the programme budget at TI-10
Development partners	Project loans and grants, mostly for national roads	The donor pool, which cannot fund routine maintenance

Two features of this structure matter for planning. First, the reform that rationalised UNRA and URF into MoWT changed the vote structure in the middle of the period the engine reads. The engine therefore builds its financing pools from budget lines (output by output) rather than from votes. The successor of the URF national maintenance transfer after FY2024/25 is identified by its output name, which is an inference, and the report flags it as one. Second, the networks in worst condition (district and community roads) are run by local governments with the smallest and least predictable resources. The three maintenance and grant pools are each between about 275 and 330 bn a year. Donor finance, by contrast, swings between 1,619.5 bn and 1,911.3 bn and cannot be spent on routine maintenance.

1.4 The five planning questions

Every module, every page of the interface and every chapter of this report answers one of five questions. Table 1.2 lists them with the modules that answer them and the chapters where the answers are reported.

Table 1.2: The five planning questions, the modules that answer them and where this report reports them

#	Question	Modules	Primary output	Chapters
Q1	What are the transport infrastructure **needs**?	TI-01 Spatial base, TI-02 Needs, TI-03 Demand	Access requirement by settlement and service; required standard per link and node	4, 5, 6
Q2	What is currently **available**?	TI-04 Supply	Network inventory, condition, capacity, baseline accessibility	7
Q3	Where are the **gaps**?	TI-05 Gaps	Seven separate gap dimensions, per link, node, settlement and district	8
Q4	What will it **cost** to fill the gaps?	TI-06 Costing, TI-07 Investment	Capital by intervention, lifecycle and induced maintenance, the investment programme	9, 10
Q5	Where should **resources go**?	TI-08 Allocation, TI-09 Outcomes, TI-10 Budget	Allocation under an envelope, what each scenario buys for each sector, the programme budget and MTEF	11, 12, 13, 22

A sixth question, *what will we achieve?*, is answered inside Q5 by TI-09. In transport an outcome only means something next to the allocation that produced it, so the report puts the two together.

GIS runs through all five questions. Every result row carries a spatial key (a link, a node, a settlement cell or a district), and every module writes map layers. National totals are sums of link and district results, never a national figure shared out top-down.

1.4.1 The two planning identities

The arithmetic behind the five questions rests on two identities. The first is about **access**: whether people can reach the services they need. For settlement cell *i*, destination class *k* (an all-season road, a health facility, a school, a market, an administrative centre) and year *t*:

$$A_{i,k,t} = \mathbf{1}\left[tt_{i,k,t} \leq \tau_{k,u(i)}\right] \tag{1.1}$$

$$U_{k,\,d,\,t} = \sum_{i \,\in\, I_d} P_{i,\,t}\,(1 - A_{i,\,k,\,t}) \tag{1.2}$$

where

- tt is the wet-season travel time (or distance) from cell *i* to the nearest destination of class *k* on the network;
- τ is the access standard for class *k* in an urban or rural settlement *u(i)*, a norm;
- A equals 1 when the standard is met and 0 otherwise;
- P is the population of the cell, and U is the population of district *d* that does not meet the standard.

The second identity is about **movement**: the goods and passengers a link must carry. Movement generated at sector nodes is distributed between origins and destinations, loaded onto links, converted into vehicles and axle loads, and turned into a required standard for each link. The required standard is then compared with the standard that exists, and the difference is priced. Chapters 5 and 6 set out the equations.

The two identities meet in the gap module (TI-05), which reports seven gap dimensions side by side: access (G1), condition (G2), capacity (G3), standard (G4), connectivity (G5), sector connection (G6) and resilience (G7).

1.5 Principles that shape every result

The design manual sets twelve principles. Six of them explain most of what a reader will notice in this report.

- **Needs first, money last (P1).** Required standards come from population, sector output and traffic. The budget enters only at TI-08.
- **Standards and costs are data, not code (P2).** Access thresholds, design standards, generation rates and unit costs are rows in the norm tables, each with a source, a year and a quality tier. No country coefficient is written into the engine's code. Chapter 3 explains how.
- **No composite index (P3).** The seven gap dimensions are always reported separately, each in its own unit. The engine never averages them into a priority score, and this report does not rank districts by any combination of them.
- **Every kilometre built carries its maintenance (P5).** Every capital figure is shown with its induced routine maintenance, the recurrent cost the new asset creates. Capital and maintenance are shown separately and never summed first.
- **Observed, estimated and modelled are never mixed silently (P10).** Every value carries a basis (observed, estimated, modelled or assumed) and a quality tier (measured, normed, imputed or assumed). The engine never generates a traffic count, a condition survey or a coordinate, and it never infers a road's surface from its class.
- **Scenarios are compared, not scored (P12).** The engine reports what each allocation buys for each sector side by side. It never marks a scenario as recommended. The optimiser runs are named by their objective and floors, for example "OPT, maximising people served within 2 km of an all-season road".

1.6 How to read this report

The report is organised by the five questions, preceded by three chapters on how the engine works and followed by chapters on its engineering, data and integration.

- **Read the quality tier with the number.** Where a result depends on an assumed norm, an estimated surface or a candidate price, the text says so. The tiers are explained in Chapter 3. Prices carry an evidence tier: E5 for document-evidence candidates, E7 for the illustrative coefficients of the unit-cost system (UCSN).

- **Read ranges, not points.** Where the engine cannot resolve a value it reports bounds. The RAI between 28.7 % and 93.5 % is the clearest case. Capital is reported as low, central and high.
- **Read each gap in its own unit.** People unserved, km in poor condition, links over capacity, km below standard, cut edges, nodes unconnected and people exposed to flood are seven different things. Maps show one dimension at a time.
- **Read the maps with their LISA panels.** Most district maps are paired with a local Moran (LISA) map. It marks districts that sit in significant clusters of high values (High–High), low values (Low–Low), or that differ from their neighbours (High–Low, Low–High), at 999 permutations. Chapter 21 explains the method.
- **Units.** Money is in constant 2026 Uganda shillings (UGX): bn is 10^9 and tn is 10^{12}. Lengths are in km, population in persons, fiscal years as FY2024/25.

1.7 Structure of the report

The report has an executive summary, 24 chapters and two annexes. Chapter *k*+1 implements design chapter *k* of the manual.

1. **Introduction** (this chapter).
2. **Architecture**: the ten modules, the GIS subsystem, the run model and the fiscal loop.
3. **Data model, quality tiers and provenance**: schemas, keys, norms versioning, quality and basis tags, and `explain()`.
4. **TI-01 Spatial base**: population, the network, terrain, hazards and the sector nodes.
5. **TI-02 Needs**: the access requirement and movement generation.
6. **TI-03 Demand and required standards**: distribution, loading, axle loads and the required standard of every link.
7. **TI-04 The network today: supply and access**: wet-season access, the RAI and condition.
8. **TI-05 Gaps**: the seven gap dimensions.
9. **TI-06 Costing**: unit prices, the condition backlog and the maintenance requirement.
10. **TI-07 Investment programme**: the decision rule, projects and appraisal.
11. **TI-08 Resource allocation**: financing pools, scenarios A–G and the optimiser.
12. **TI-09 Outcomes and simulations**: what each allocation buys.
13. **TI-10 Budget and MTEF**: the programme budget and fiscal sustainability.
14. **Integration with sector engines**: the exchange contracts.
15. **The planning interface**: the web application.
16. **Calibration and validation**: tests and gates.
17. **Reference catalogues: norms and unit costs**.
18. **Build status**.
19. **The database**.
20. **Existing data assets and acquisition**.
21. **GIS analysis across modules**: spatial autocorrelation and clusters.
22. **Sector linkages**: health, education, agriculture, minerals and industry.
23. **TransportModelling reuse**.
24. **Mathematical models**.

Annex A gives a one-line transport summary of every district. Annex B is a glossary of terms and abbreviations.

1.8 What this report can and cannot say

Three limits run through every chapter.

Surface and condition. The engine knows the surface of under a third of the network and the condition of none of it, link by link. Every access figure is therefore a range or carries a warning, and on 169,343 km of road the engine's answer to "what works are needed" is "survey first". The condition gap is reported only as network-level agency statistics.

Sector feeds. The minerals and trade exchanges in the estate are synthetic demonstration data and are not used. Mines and industrial parks enter as nodes without output. Agricultural surplus comes from MapSPAM 2010 production, not from a current sector model, and freight is grown at a flagged fallback rate of 5 % a year because no sector output trajectories were supplied.

Prices and norms. Unit costs are document-evidence candidates (tier E5) built from 89 of 275 extracted evidence rows. The norm set is a draft, and 578 of its 646 rows are assumed. No figure in this report should be read as an approved estimate.

These limits do not make the results useless. They make them exploratory, and they show precisely which data would change them. That list is itself a result of the engine, and Chapter 20 turns it into data requests.

1.9 Implications for planning

- The planning question should be put as the engine puts it: from population, services and producers to required standards, then to cost, and only then to the envelope. A plan that starts from the size of the road fund answers a different question.
- The local networks are where the condition problem is: 86 % of district roads and 94 % of community access roads were reported in poor condition in FY2022/23, against 3 % of national paved roads. Any allocation has to be read against that split.
- The institutional change that rationalised UNRA and URF into MoWT should be reflected in how financing pools are defined. Pools built from budget outputs survive a change of votes; pools built from votes do not.
- The single most valuable data investment for planning is a link-level surface and condition survey of district and community roads. Without it the Rural Access Index can only be stated as a range from 28.7 % to 93.5 %.
- Every capital figure in later chapters carries its induced maintenance. Read it alongside the maintenance funding ratio of 0.58, which shows that maintenance, not capital, is already the binding constraint.
- Results are exploratory. They are fit for comparing options and for directing surveys and data requests. They are not yet fit for appropriating money to named roads.

CHAPTER 2

2 Architecture of the planning engine

This chapter describes how TIPRAS is built: the ten modules and the order in which they run, the GIS subsystem they all share, the run model that makes every result reproducible, and the loop that ties the plan to the fiscal framework. It implements design Chapter 1 (`docs/01_ARCHITECTURE.md`) and reports where the Uganda build follows the design and where it deliberately departs from it. A reader who wants only the results can skip to Chapter 4. A reader who wants to know whether the results can be trusted, reproduced or extended should read this chapter first.

2.1 The planning chain

TIPRAS is a chain of ten computational modules that run in dependency order over one GIS subsystem. Figure 2.1 shows the chain as it runs for Uganda. The first three modules (dark blue) answer Q1, the needs. TI-04 answers Q2, what exists. TI-05 answers Q3, the gaps, in seven dimensions G1 to G7. TI-06 and TI-07 answer Q4, the cost. TI-08, TI-09 and TI-10 answer Q5, where the resources should go. Above the chain, the figure shows where each kind of outside evidence enters: the SDI geography, WorldPop population, the OpenStreetMap network and the hazard layers at the start; the norms and the agency condition statistics at TI-04 and TI-05; the unit costs (UCSN and document evidence) at TI-06; and the MTEF ceiling and the budget pools only at TI-08. The order is the point. Money enters at the eighth module, not the first.

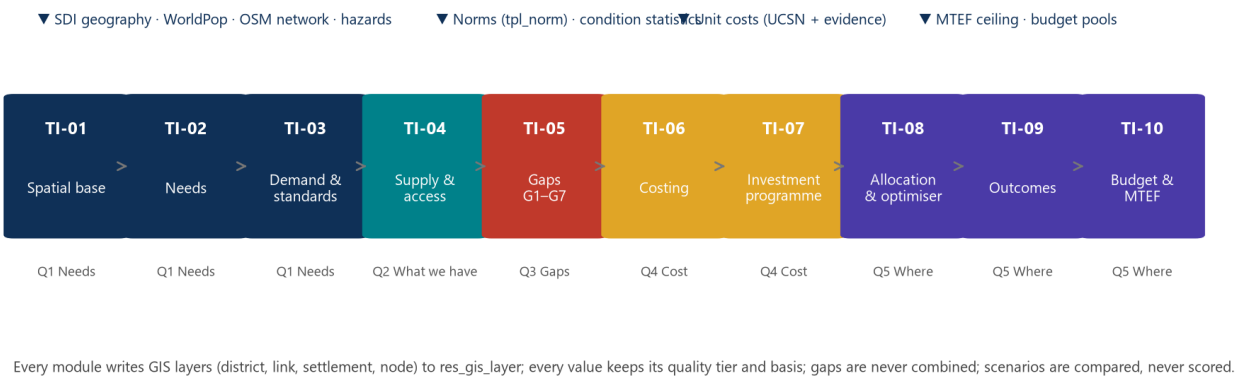

Figure 2.1: The planning chain: ten modules from the spatial base to the budget

Source: docs/01_ARCHITECTURE.md

Table 2.1 lists the ten modules with the question each answers and what it hands to the next.

Table 2.1: The ten modules of the planning chain

Module	Name	Answers	Question	Main output passed on
TI-01	Spatial base and demand generators	Who and what must be connected, and where	Q1	Raked settlement population, sector nodes, agricultural zones, destination sets, snap distances
TI-02	Transport needs	What access standard each place requires, and how much movement each node generates	Q1	Access requirement per settlement and class; node connection standards; freight and trip generation

Module	Name	Answers	Question	Main output passed on
TI-03	Transport demand	Which links carry that movement, and what standard each link therefore requires	Q1	Required-state AADT and ESA; the required standard of every link and its binding test
TI-04	Supply and condition	What network, structures and nodes exist, in what condition, with what access	Q2	Wet-season travel times, the Rural Access Index, condition by network
TI-05	Gaps	Where requirement exceeds supply	Q3	Seven gap dimensions per link, node, settlement and district, never combined
TI-06	Costing	What the required works and their maintenance cost	Q4	Capital by works type (low, central, high), routine and induced maintenance
TI-07	Investment programme	Which interventions, in which order, with what dependencies and appraisal	Q4	Projects with coverage sets, NPV, sector tags
TI-08	Allocation	Given a budget, where the money should go	Q5	Scenarios A–G and the optimiser runs
TI-09	Outcomes	What each allocation buys	Q5	Access, equity and sector outcomes at the horizon
TI-10	Budget and MTEF	The plan as a programme budget	Q5	Three reconciling budget views; outbound files to MTEF, FINEX and the sector engines

2.1.1 The dependency rule

A module may read the result tables of any earlier module, and the reference, norm and geography tables. It may not read the results of a later module. Written formally, for modules indexed by their position *n* in the chain:

$$\text{reads}(TI_n) \subseteq \{R(TI_m) : m < n\} \cup \text{Ref} \cup \text{Norm} \cup \text{Geo} \tag{2.1}$$

where R is the set of result tables a module writes, and Ref, Norm and Geo are the reference catalogues, the norm tier and the pinned geography. The rule has one exception, the fiscal loop between TI-08, TI-10 and the FINEX fiscal model, which is an explicit, bounded iteration (see below). The rule is what makes the chain auditable. Any figure in TI-05 can be traced back through TI-04 to TI-01 without ever meeting a result that was computed later.

2.1.2 The GIS subsystem is not a module

The GIS subsystem has no result tables of its own. It is a library that every module calls: it builds the routable graph, snaps points to it, computes travel times and catchments, finds cut edges, overlays hazards and computes spatial statistics. Each module then writes its own spatial outputs into one shared layer table. Table 2.2 lists what is built for Uganda.

Table 2.2: The GIS subsystem as built (src/GIS)

File	What it does	Used by
`Geometry.jl`	Stores coordinates in EPSG:4326 and measures in the projected CRS (UTM zone 36N, EPSG:32636)	all modules
`Snap.jl`	Snaps settlements and nodes to the nearest link segment with a KD-tree and stores the off-network distance	TI-01, TI-04
`Network.jl`	Builds the routable graph from the pinned snapshot in compressed sparse row (CSR) form; seasonal variants	TI-03, TI-04, TI-07, TI-09
`TravelTime.jl`	Multi-source Dijkstra from a destination set on the reversed graph, with an off-network walking leg	TI-02, TI-04, TI-07, TI-09
`Criticality.jl`	Cut edges (Tarjan's algorithm) and the population isolated behind each	TI-05, TI-07
`SpatialStats.jl`	Moran's I, local Moran (LISA), concentration index and Gini	every module's GIS layers

The boundary of the subsystem is deliberate. TIPRAS computes on the **vector network**: shortest paths, catchments, loads, cut edges and link overlays. It does not compute raster friction surfaces, least-cost alignments for new roads, flood models or slope from a digital elevation model. Those arrive as link attributes or candidate geometries from SDI or TransportModelling. One consequence is visible in Chapter 8: because no candidate alignments for new links have been supplied, the engine cannot yet separate "the network is missing" from "there is no destination" as a cause of an access gap.

The reason the engine computes shortest paths itself, rather than calling an outside engine, is the optimiser. To appraise a candidate the engine must recompute access on the modified network, and with tens of thousands of candidates a round trip to an outside engine per candidate is not feasible.

2.1.3 Travel time to a class of destinations

The most frequently used GIS operation is travel time from every settlement to the nearest destination of a class. Running one shortest-path search per destination would be prohibitive with 20,474 schools. The engine runs one search per class instead. It adds a super-source joined to every destination of the class and runs Dijkstra's algorithm from it on the reversed graph, so that the label at each vertex is the cost from that vertex to the nearest destination:

$$D_k(v) = \min_{j \in J_k(t)} \left[d(v, \mathrm{snap}(j)) + w_j \right] \tag{2.2}$$

$$tt_{i,\,k,\,t} = w_i + D_k(\mathrm{snap}(i)) \tag{2.3}$$

where

- J is the set of functional destinations of class *k* in year *t*;
- d is the shortest network cost between two vertices;
- w is the off-network walking leg at each end, from the snap distance;
- tt is the travel time from cell *i*; an unreachable cell is stored as missing, never as zero.

Running on the reversed graph makes the label correct on one-way links. The method is ported from TransportModelling's accessibility model and tested against brute force on random graphs.

2.2 How the system is put together

Figure 2.2 shows the system as built for Uganda in four layers.

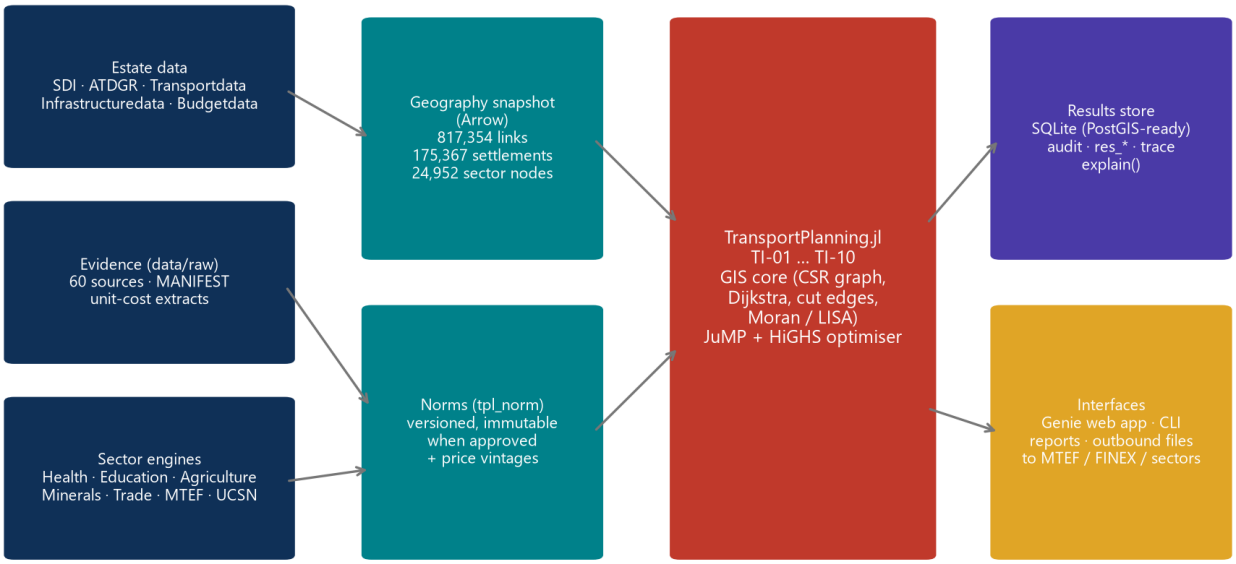

Figure 2.2: System architecture: data tiers, engine, store and interface

Source: docs/01_ARCHITECTURE.md; IMPLEMENTATION_STATUS.md

- **Inputs (left).** Three groups of input feed the engine: the estate data (SDI, ATDGR, Transportdata, Infrastructuredata and Budgetdata); the evidence collected for this build (60 registered sources with a manifest, and the unit-cost extracts); and the sector engines (health, education, agriculture, minerals, trade, MTEF and the UCSN unit-cost system).
- **Pinned inputs (centre left).** The inputs are frozen into two versioned objects before any run: a geography snapshot in Arrow format, holding 817,354 road links, 175,367 settlement cells and 24,952 sector nodes, and a norm set with its price vintage, immutable once approved.
- **Engine (centre).** `TransportPlanning.jl` runs TI-01 to TI-10 over the GIS core (CSR graph, Dijkstra, cut edges, Moran and LISA) and the JuMP and HiGHS optimiser.
- **Outputs (right).** The results go to a store (SQLite, with a schema ready to be ported to PostGIS) that holds the audit tables, the result tables and the factor traces behind `explain()`. They also go to the interfaces: the Genie web application, the command line, the reports and the outbound exchange files to MTEF, FINEX and the sector engines.

2.2.1 Technology

Table 2.3 compares the stack the design specifies with the stack as built.

Table 2.3: Technology stack, specified and as built

Layer	Specified	As built for Uganda
Language	Julia 1.10+	Julia, launched with `--pkgimages=no` because Windows Application Control blocks the precompile cache
Web and API	Genie.jl with HTTP.jl	Genie application on the EnergyModellingx pattern: 12 pages, maps, charts, exports and a run dialog
Database	PostgreSQL 16 with PostGIS 3.4	SQLite offline store (the specification's field mode) with forward-only migrations 0001–0003, plus Arrow geography snapshots
Network graph	Graphs.jl	Own CSR graph with super-source Dijkstra and Tarjan cut edges
Spatial index	NearestNeighbors.jl	KD-tree on projected coordinates for snapping
Optimisation	JuMP.jl with HiGHS	JuMP and HiGHS mixed-integer programme, money scaled to UGX bn

Layer	Specified	As built for Uganda
Maps	geopandas and matplotlib for the atlas; Leaflet in the web application	As specified; the renderers compute nothing
Tests	Test.jl	2,265 assertions passing in `test/runtests.jl`

The largest departure is the database. The local PostgreSQL server needed credentials the build did not have, so the engine uses the SQLite offline store the design allows for single-user work, with the same schema subset. The SDI database cluster was read only to export the noded network. The migrations are written to be ported to PostGIS when a TIPRAS database is provisioned.

2.2.2 Code ported, not linked

Several functions the engine needs already existed in TransportModelling: the network graph, speed profiles, multi-source accessibility, cut edges, deterioration and maintenance, appraisal and vehicle operating cost, the portfolio optimiser, the unit-cost reader and the MTEF reader. The design rule is to **port them, not depend on the package**, so that the two systems can be released independently. Each ported file names its origin and the date it was ported. Its literal defaults were removed and replaced by norm lookups: for example, the 25 km/h fallback speed in TransportModelling's speed profile was removed, so a missing speed now raises an error instead of quietly becoming 25 km/h. Chapter 23 lists the full inventory.

2.2.3 Command line and interface

The command-line tool, `tipras`, computes nothing itself. It dispatches to the engine. The commands used to build this report are:

- `migrate`, to create or update the store;
- `norms-import` and `norms-approve`, to load a norm workbook as a draft and approve it;
- `prices-load`, to load a price vintage;
- `run`, to run TI-01 to TI-10 with pinned geography, norms and prices, with an `--exploratory` flag;
- `explain`, to decompose any result cell into its factors;
- `prune`, to keep only the latest runs, because each national run adds about 2 GB to the store.

The web interface (Chapter 15) reads the same store and recomputes nothing. It can start a new run as a separate process, so the application stays usable during the hour or more a full run takes.

2.3 The run model

Every computation belongs to a run. A run is defined by its **pins**: the country, base year and horizon, the geography version, the norm version, the price vintage, the target set, the scenario, the envelope version, the season used for access tests and a flag that says whether the run is exploratory. The reference run for this report is pinned as follows.

Table 2.4: Pins of the reference run 014c9cd7

Pin	Value
Country, base year, horizon	Uganda, 2026, 2026–2030
Geography version	`UGA-2026.09-c` (SDI noded network with OSM attributes; links hashed)
Norm version	`UGA-2026.09-draft-f` (draft; 646 rows in 28 tables)
Price vintage	`UCSN+EV-UGA-2026.09b` (evidence candidates, tier E5)
Season for access tests	wet
Exploratory	yes
Code version	TransportPlanning 0.1.0
Started and finished	24 September 2026, 09:09 to 10:37

Three properties follow from the pins.

- **Reproducibility.** Two runs with identical pins produce identical numbers, including the permutation p-values of the spatial statistics, which use a fixed random seed.
- **Immutability of geography.** The geography snapshot is hashed. A network edit creates a new geography version and never changes an old one. The Uganda snapshot has gone through versions `a`, `b` and `c` as surface evidence, ferries, crossings and culverts were added; this report uses `c`, the only one still kept on disk.
- **Requirement and plan kept apart.** Every run keeps both the unconstrained requirement (TI-06) and the constrained plan (TI-08). The difference between them is the financing gap.

The **exploratory** flag is not a formality. It permits draft norms and candidate or illustrative prices, which a non-exploratory run refuses. Every Uganda run so far has been exploratory, and every page and table of this report carries that status. The run manifest also records every **fallback**: each place where an interface was missing and the engine used a documented substitute. The reference run records eight, including freight grown at a fallback rate of 5 % a year because no sector output trajectories were supplied, and unknown surface on 68.2 % of network km, which exceeds the 30 % threshold above which every access figure is flagged.

2.4 The fiscal loop

The fiscal loop is the one place where the chain turns back on itself. TI-10 sends the plan's requirement to FINEX, the fiscal model. FINEX returns an envelope, and TI-08 re-allocates within it. The loop stops when the total allocation changes by less than half a per cent between iterations, and in any case after five iterations:

$$\frac{|E^{(n+1)} - E^{(n)}|}{E^{(n)}} < 0.005, \quad n \leq 5 \tag{2.4}$$

where E is the total allocation in iteration *n*. If the loop does not converge, the engine reports the oscillation and stops rather than picking an iteration. In the reference run the loop is not closed: FINEX does not yet return an envelope, so TI-08 uses the MTEF transport ceiling with the roads share (0.545) and pool split taken from the approved budgets of FY2024/25 to FY2026/27. TI-10 writes the outbound files for MTEF, FINEX, HealthModellingx, EducationModellingx and the agriculture model, ready for the loop to be closed.

2.5 Performance

The design sets performance targets for a country of Uganda's size, about 820,000 noded links, 135 districts and five plan years. They require a simplified planning graph in which chains of degree-2 vertices are collapsed, while every junction, crossing and snapped destination is kept with a mapping back to the source links. Table 2.5 compares the targets with the reference run, which used eight threads.

Table 2.5: Performance targets and the reference run

Operation	Target	Reference run 014c9cd7
TI-01 to TI-06, unconstrained	< 10 min	6.0 min (TI-01 35.5 s, TI-02 13.9 s, TI-03 124.1 s, TI-04 67.7 s, TI-05 92.6 s, TI-06 27.9 s)
Candidate appraisal with network re-evaluation	2,000 candidates in < 20 min	41,757 priced projects appraised; TI-07 took 141.9 s
Rule-based scenarios A–G and the MILP	< 5 min for scenarios; MILP < 15 min at 1 % gap	TI-08 took 76 min: seven scenarios, four optimiser specifications, their attainable maxima and five frontier points; two solves reached the 900 s time limit
Outcomes and budget	not specified	TI-09 163.6 s; TI-10 9.3 s

Operation	Target	Reference run 014c9cd7
Whole run	not specified	88 min

The needs and supply side of the chain is well inside its target. The optimiser is not. Its run time grows with the number of specifications in the optimiser file and with the size of the candidate set: with 41,867 priced projects the pre-screen keeps at most 20,000 candidates plus the warm start, and two solves, the run with an emergency obstetric care floor and the 100 % point of the frontier, stopped at the time limit and returned their warm start. Chapter 11 treats those two results as findings, not as answers.

2.6 Where the build departs from the design

Table 2.6 lists the deliberate departures that are still open. None changes a result in this report silently. Each is recorded in the implementation status file and, where it affects a run, in the run manifest.

Table 2.6: Deliberate departures from the design, as of September 2026

Design	As built	Consequence and next step
PostgreSQL with PostGIS	SQLite offline store with Arrow snapshots	Same schema subset; port the migrations when a TIPRAS database is provisioned
Candidate alignments for new links from SDI or TransportModelling	Not supplied	Access gaps caused by a missing network cannot be separated from those caused by a missing destination
Rail, ports and waterways as links	Nodes only; 28 ferry links	Phase-2 scope in the design
Year-by-year outcome path	Horizon year only	TI-09 recomputes the network at 2030, not every year
Uncertainty propagated to appraisal	Low and high costs carried, not propagated to NPV	Corner runs to be added
Capacity-restrained incremental assignment	All-or-nothing loading only	Loads on the busiest corridors are overstated; see Chapter 6
Factor traces for every result	Traces written for TI-01 population and raking only	`explain()` works fully for population cells; other modules to follow

2.7 Implications for planning

- The chain is built so that money enters at TI-08. A planner who wants to know what the envelope buys must first accept the needs and gaps computed without it, and the engine keeps both the unconstrained requirement and the constrained plan for that reason.
- Every result is tied to a set of pins. When a figure from this report is quoted, the run identifier (`014c9cd7`) and its pins should be quoted with it, so that it can be reproduced and compared with later runs.
- The exploratory flag will stay on until the norms are approved by the agencies that own them and the prices are published by UCSN. That approval is an institutional step, not a technical one.
- The needs and supply modules already run within their performance targets on the full national network. The optimiser does not, and its run time should be managed by keeping the optimiser specifications to those the ministry actually wants to compare.
- The departures in Table 2.6 are where the next engineering effort pays: candidate alignments, capacity-restrained loading and full factor traces would each change what the engine can say, and none needs new money for works.
- The fiscal loop is designed but not closed. Until FINEX returns an envelope, the allocation chapters read the MTEF ceiling through budget-evidence pool shares, and say so.

CHAPTER 3

3 Data model, quality tiers and provenance

A planning engine is only as credible as its ability to say where each number came from and how much weight it can bear. This chapter describes how TIPRAS stores its inputs and results, how it keeps norms and prices as versioned data rather than code, how every value carries a tag for the strength of its evidence and a tag for how it arose, and how any result can be decomposed back to its sources. It implements design Chapter 2 (`docs/02_DATA_MODEL.md`), with the quality and basis rules of Chapter 16 §16.13, and reports the state of the Uganda norm set used by the reference run.

3.1 Why the data model matters for planners

Three questions recur whenever a planning figure reaches a budget discussion. Where does this number come from? Would it be the same if we ran it again? How much of it is evidence and how much is assumption? The data model is built so that the engine can answer all three for any figure, and so that it cannot quietly answer them wrongly. Its rules are simple to state:

- inputs are frozen into versions before a run, and a run records which versions it used;
- parameters that a country could dispute (standards, rates, costs) are rows of data with a source, a year and a quality tier, never constants in code;
- every observation, norm and result carries a **quality** tag and a **basis** tag, and the two are never merged;
- every result row carries exactly one spatial key, so that it can be mapped;
- nothing missing becomes a zero.

3.2 Schemas

The design specifies seven schemas, each with its own rule of mutability. Table 3.1 lists them with what the Uganda build holds in each. Because the build uses the SQLite offline store (Chapter 2), schemas appear as table-name prefixes rather than PostgreSQL schemas. The store holds 18 tables.

Table 3.1: The seven schemas, their mutability and the Uganda build

Schema	Holds	Mutability	In the Uganda store
`tpl_ref`	Reference catalogues: geography registry, road classes, surfaces, node kinds, destination classes, works types, vehicle classes, pools, budget programmes	Versioned; changed only by migration	Seed files and code-level catalogues
`tpl_norm`	The parameter tier: access and design standards, generation rates, speeds, season rules, capacity, maintenance, unit-cost map, appraisal parameters	Draft, proposed, approved (then immutable), superseded	`norm_version`, `norm_value`: 646 rows for `UGA-2026.09-draft-f`
`tpl_geo`	The pinned geography: boundaries, settlements, links, vertices, structures, sector nodes, hazard and terrain attributes	Immutable per geography version	`geo_version` plus Arrow files of snapshot `UGA-2026.09-c`
`tpl_obs`	Observations: condition statistics, prices, envelopes, sources	Append only; each row dated and sourced	`obs_price`, `obs_envelope`, `obs_source`

Schema	Holds	Mutability	In the Uganda store
tpl_result	Computed results keyed by run	Written once per run	res_access, res_gap, res_candidate, res_allocation, res_outcome, res_budget, res_gis_layer
tpl_audit	Run manifests, factor traces, events, exchange receipts	Append only; no update or delete for any role	audit_run_manifest, audit_factor_trace, audit_run_event, audit_exchange_receipt
tpl_stage	Raw extracts before harmonisation	Truncated freely; never read by a compute module	Kept outside the store, under data/raw

Migrations are forward-only and numbered. The Uganda store has three: 0001 (schemas and audit), 0002 (norms) and 0003 (reference, observations and results). A migration that has been applied outside development is never edited.

3.3 Keys and spatial keys

Every result row carries **exactly one** spatial key. This is how GIS reaches every module: a gap, a cost or an allocation can always be drawn because it always knows which link, node, settlement, cell or district it belongs to.

Table 3.2: Spatial keys and where their geometry lives

Key	Object	Uganda example
link_id	A road or ferry link in the planning graph	UGA_OSM111872938_10: the SDI road identifier (OSM way 111872938) plus a sub-segment number
structure_id	Bridge, culvert, drift or causeway	Carried as link attributes in this build (2,188 bridges, 5,039 culverts on 3,676 links)
node_id	A sector node or gateway	24,952 nodes: schools, health facilities, markets, mines, industrial zones, gateways, district headquarters
settlement_id	A settlement, including dispersed rural cells	UGA_C100_164: one of 175,367 populated 1 km cells
district_id	The planning district	UG3078 (Koboko); 135 districts, P-coded as in SDI and UBOS

Keys are text and stable across geography versions where the object has not changed. Integer keys are used inside the planning graph for speed but never appear in a result table. The district key matters more than it seems. HealthModellingx keys its facilities to 122 district names, while SDI, UBOS and EducationModellingx use 135 P-coded districts. TIPRAS assigns every node to a district by point-in-polygon against the pinned SDI boundaries and never uses a counterpart's district code.

3.4 The geography tier

The geography tier holds the pinned snapshot. Each version records its source exports, a content hash and the projected coordinate system used for measurement. For Uganda the projected system is UTM zone 36N (EPSG:32636), the same as TransportModelling, and every length and distance is measured in it. A length computed in geographic coordinates is a test failure (rule GIS-4).

A link in snapshot UGA-2026.09-c is a record of 27 fields. Some are geometric (length, terrain class from the elevation model, district by majority length). Some are inventory (functional class, surface and the basis of the surface, width, lanes, road reference, posted speed, smoothness). Some are structures and hazards (bridge, culverts, seasonal crossing, unrecorded river crossing, flood exposure). Two fields carry the network's largest

information gaps: **surface**, which is `unknown` on 169,343 km, and **condition**, which is `unknown` on every one of the 817,382 links. Neither is filled by inference from road class. Chapter 4 shows the snapshot in detail.

3.5 The norms tier

The norms tier is the critical path of the engine: every computation reads it. Every norm table shares one header:

- `norm_version`: the version the row belongs to;
- `geo_scope`: national, urban, rural, a region or a district;
- `source` and `source_year`: where the value comes from;
- `quality`: measured, normed, imputed or assumed;
- `note`.

3.5.1 Versions and approval

A norm version moves through a fixed life cycle: **draft → proposed → approved → superseded** (or rejected). Once a version is approved, a database trigger refuses any update or deletion of its rows. A correction becomes a new version that supersedes the old one. A run records the version it used, so a result can always be reproduced with the norms it was computed from.

The Uganda norm set has been rebuilt six times during the build, each as a new draft version: `UGA-2026.09-draft` (623 rows in 27 tables, seeded from TransportModelling's configuration and the manual's defaults), then `draft-b` to `draft-f` as evidence was added. The changes included evidence on surface and condition, the budget-evidence pool split and the switch from illustrative to evidence-candidate prices. Two further variants, `draft-f-env70` and `draft-f-env130`, scale the envelope for sensitivity runs. None of these versions has been approved.

3.5.2 Resolution order

The engine never reads a norm directly. It asks for the **most specific** applicable row, searching in a fixed order of geographic scope:

$$\mathrm{resolve}(T, \kappa, g) = \mathrm{first}(T[\kappa, \mathrm{district}(g)],\ T[\kappa, \mathrm{region}(g)],\ T[\kappa, u(g)],\ T[\kappa, \mathrm{national}]) \tag{3.1}$$

where T is a norm table, κ is the key (for example a destination class and settlement type) and g is the place the norm is needed for: its district, its region and whether it is urban or rural (u). If no row matches at any level, the engine raises `NormNotFound(table, key, geo)` and the run's gate aborts with a list of every missing norm. **It never defaults to zero or to a neighbouring category.** In the Uganda draft every row is national in scope, so the district and region levels are not yet used. They exist so that, for example, a mountainous district can carry its own unit cost or a city its own access standard without a code change.

3.5.3 What the Uganda norm set contains

Figure 3.1 shows the norm set of the reference run by quality tier. Of its 646 rows in 28 tables, **578 (89.5 %) are assumed** and **68 (10.5 %) are normed**. None is measured or imputed.

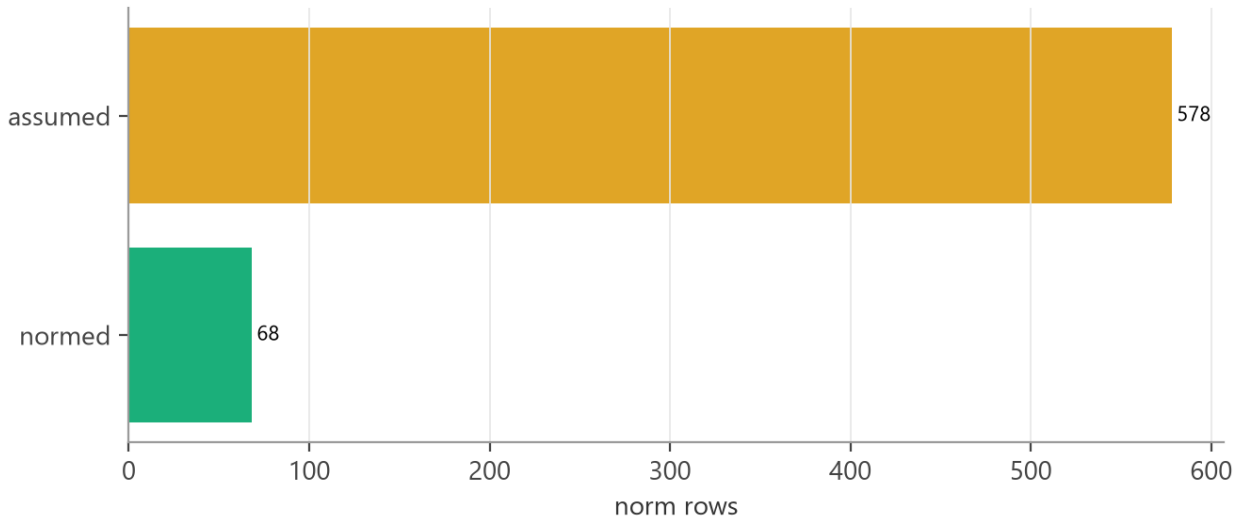

Figure 3.1: The norm set by quality tier: what is evidence and what is still assumed

Source: norm_value table, reference norm version

The 68 normed rows deserve a closer look, because "normed" means "a country standard or policy value, cited", and not all of them are Ugandan policy:

- **16 rows cite external policy or evidence.** Ten are access standards: the all-season road standard for rural settlements (2 km on foot in the wet season, from the Rural Access Index and SDG 9.1.1) and the emergency obstetric care standard (2 hours by motorised transport for women aged 15–49, from the WHO/Lancet standard). Six are the pool shares, taken from the approved budget estimates of FY2024/25 to FY2026/27: the roads share of transport (0.545) and the split of the roads envelope into road-fund maintenance (0.080), government development (0.333), recurrent (0.021), district grants (0.088) and donor (0.478).
- **52 rows cite the design manual itself.** These are the destination eligibility rules (which facility kinds satisfy which destination class, 6 rows), the order of the investment decision rule (9 rows), the works each financing pool may fund (5 rows) and the map from works type and surface to unit-cost code (32 rows). They are structural rules of the method rather than country parameters, and they are tagged normed because the method cites them.

The 578 assumed rows include every design standard (132 rows), every speed (98 rows), every season rule (46 rows), every node connection standard (48 rows), all the vehicle parameters (37 rows), the deterioration parameters (35 rows), every freight generation rate (18 rows), every trip rate and the age-group shares used to compute who must reach a school or a maternity unit. The primary health care standard is tagged "national health policy (to import from HealthModellingx)" but is still assumed, because it has not yet been imported from the health engine. The same is true of the primary school standard.

The conclusion is plain. The engine's method is in place, but nearly every number the method uses is a placeholder awaiting a Ugandan source. Chapter 17 lists the norms table by table, and Chapter 20 the data requests that would replace them.

3.5.4 Norm entry and sign-off

Norms are entered through a workbook with one sheet per table, imported into a draft version with `tipras norms-import`. The design assigns ownership: the road agency's engineers review the design, maintenance and speed norms; the health, education and agriculture ministries sign off the access standards that concern them, because they own them. Where a sector engine already holds a standard (HealthModellingx and EducationModellingx both hold access standards), TIPRAS imports the row instead of retyping it, so that a sector's change of standard changes transport need with no code change.

3.6 Quality and basis: two tags, never merged

Every norm, observation and result row carries two tags. They answer different questions.

Table 3.3: The quality tier: how strong is the evidence?

Quality	Meaning	Uganda examples
`measured`	Derived from a country measurement of the thing itself	None in the reference run: no traffic count, axle-load survey or link condition survey is available
`normed`	A country standard or policy value, cited	The 2 km all-season standard; the budget-evidence pool shares
`imputed`	Estimated from related country data by a documented method	Settlement population raked from the WorldPop grid to the UBOS projection
`assumed`	A default without country evidence	Design standards, speeds, trip rates, the vehicle mix marked "not Ugandan"

Table 3.4: The basis: how did the value arise?

Basis	Meaning	Uganda examples
`observed`	Recorded directly by a survey, count, register or report	A surface tag in OpenStreetMap; a facility coordinate; the poor km in the MoWT statistical abstract
`estimated`	Derived from observations by expansion or conversion	A surface classified from satellite imagery (HeiGIT) or matched from GRIP4; raked population
`modelled`	Computed by a model	Required-state traffic (TI-03); wet-season access (TI-04)
`assumed`	Set by assumption, with no observation behind it	A fallback trip rate; the 3 % a year population growth

Prices carry a third, parallel scale inherited from UCSN, the unit-cost system. The seven UCSN transport coefficients for Uganda are **ILLUSTRATIVE (tier E7)**, the equivalent of assumed. The evidence prices built for this report from contracts, programme reports and regional studies are **EVIDENCE_CANDIDATE (tier E5, low confidence)**. Both are refused by a non-exploratory run.

3.6.1 Rules that bind the tags

- **Valid pairs.** Observed goes with measured or normed; estimated with measured or imputed; modelled with any quality, inheriting the worst of its inputs; assumed only with assumed. Any other pair is a load error.
- **Nothing is ever promoted.** A modelled load is not relabelled observed because it matches a count. A value with basis assumed is never shown without its flag.
- **A result inherits the worst quality of its inputs.** With the tiers ordered measured, normed, imputed, assumed (from strongest to weakest):

$$q(x) = \max_{f \in \text{inputs}(x)} q(f), \quad \text{measured} < \text{normed} < \text{imputed} < \text{assumed} \tag{3.2}$$

- **A national aggregate warns when assumptions dominate.** If more than 30 % of a national total comes from assumed inputs, the figure displays a warning:

$$\omega(X) = \frac{\sum_{x \in X,\ q(x) = \text{assumed}} x}{\sum_{x \in X} x} > 0.30 \tag{3.3}$$

where X is the set of district or link values summed into the national figure.

The worst-quality rule is unforgiving by design. A district's unserved population combines raked population (imputed), an access standard (often assumed) and a wet-season travel time that depends on assumed speeds

and season rules. The result is therefore assumed, however careful the population work. That is the honest answer, and it shows where an improvement would have most effect: replacing the speed and season norms would lift the quality of every access result at once.

One inconsistency is visible in the reference run. The factor trace tags the raking factor as normed with basis estimated, a pair that the valid-pair rule does not list. The factor is derived from the UBOS projection (normed in itself) by a calculation, so imputed would be the consistent tier. It does not change a number, but the trace should be corrected.

3.7 Observations

Observation tables are append-only. Every row carries the date of observation, a source identifier, a basis and a quality. A correction is a new row with a later date and a pointer to the row it supersedes. The design lists ten observation tables, from road attributes and condition surveys to traffic counts, seasonal closures, projects, expenditure, envelopes and population.

For Uganda most of these are empty, and the emptiness is itself recorded:

- **Condition** is available only as administratively reported totals by network for FY2022/23 (national paved, national unpaved, district, community and four urban networks), from the UBOS Statistical Abstract 2024 and the MoWT Statistical Abstract 2023. They are stored with method `administratively_reported` and are never spread over links.
- **Traffic counts** are absent. Every traffic figure in the report is modelled and labelled so.
- **Prices** come from 275 extracted evidence rows, of which 89 are used by the 14 evidence-candidate price codes.
- **Envelopes** come from the approved budget estimates, read line by line (619 rows aggregated by vote, programme, output and funding source).

3.8 Results and the shared GIS layer

Results are keyed by run, spatial key and year, and where relevant by scenario. They are written once per run and never updated. Every module also writes into one shared layer table, which is what the maps in this report draw from. Each row holds a module, a layer identifier, a geometry kind (link, structure, node, settlement, cell, hex or district), a geometry key, a year, a scenario, a value, a unit and a quality.

The reference run wrote district layers for every module, from population density in TI-01 to the allocation per capita of each scenario in TI-08 and the RAI at the horizon in TI-09, and link layers for traffic, axle loading, freight, wet-season closures and works type. The global and local Moran statistics are computed for 49 district layers. To save disk space the store keeps the full layer table only for the latest runs; the `prune` command deletes the results of older runs but keeps their manifest and factor traces, so that a pruned run can be regenerated from the same pins.

3.9 Provenance: the manifest, the trace and `explain()`

3.9.1 The run manifest

Every run writes one manifest row: its pins (Chapter 2), the code version, start and finish times, the gate results and a list of every **fallback**, meaning every interface that used a documented substitute, with its reason. The reference run records eight fallbacks. They range from the freight growth rate (5 % a year, because no sector output trajectories were supplied) and the unknown surface share (68.2 % of network km) to three optimiser specifications whose floors were not fully met, which Chapter 11 reports as findings.

3.9.2 The factor trace and `explain()`

The factor trace records, for a result cell, each factor it depends on: its value and unit, a reference to its source (a norm row, an observation, an exchange file or an upstream result cell), its quality and its basis. `explain(run_id, cell_ref)` walks the trace recursively through upstream cells and returns the full decomposition tree. Each leaf carries its value, source, year, quality and basis.

An example from the reference run shows what this gives a planner. The base-year population of settlement cell `UGA_C100_164`, in Koboko district, is 515.5 persons. `explain()` decomposes it into two factors:

$$P_{i,\,2026} = P_i^{\text{grid}} \cdot \rho_{d,\,u,\,2026} = 411.5 \times 1.2527 = 515.5 \tag{3.4}$$

- the WorldPop grid population of the cell, 411.5 persons (source: the geography snapshot; quality imputed; basis estimated);
- the raking factor for rural Koboko in 2026, 1.2527 (source: the UBOS-based population projection for district UG3078; basis estimated).

For 2030 the same cell has the same grid population and a raking factor of 1.4100, because the projection carries the district total forward at 3 % a year.

The reference run holds **701,468 trace rows**, two for each of the 175,367 settlement cells in the base year and in the horizon year. At present only TI-01 writes traces. The later modules do not yet write factor traces, so `explain()` cannot yet walk from a district's unserved population to the speed norm and season rule behind it. Extending the trace to TI-02 onwards is listed in Chapter 18 as open work.

3.10 Units, currency and time

- **Units.** Lengths are in km, measured in the projected system. A norm expressed per lane-km is multiplied by lanes and km, never by km alone. Equivalent standard axles per year are multiplied by the design life to give design traffic. A unit mismatch is a load-time error.
- **Money.** Money is stored in constant 2026 UGX with the price vintage pinned, and vintages are never mixed in one column. Evidence prices from other years are carried to 2026 with Uganda's GDP deflator; US-dollar prices are converted two ways (at the price-year exchange rate then Uganda's deflator, and with the US deflator then the 2026 exchange rate) and the geometric mean is taken. Conversion to nominal terms happens only in TI-10.
- **Time.** Access is measured in minutes and appraisal in hours, and the conversion happens in one place.

3.11 The missing-data policy

The last rule of the data model is the one most often broken in planning practice: **nothing missing becomes a silent zero**. Table 3.5 shows how the engine behaves when an input is absent, and what that means for Uganda today.

Table 3.5: Missing-data behaviour and its effect on the Uganda run

Situation	Behaviour	Effect in the reference run
Norm missing	`NormNotFound`; the gate aborts with a list	None missing in `draft-f`
Price missing	The project is listed as unpriced and cannot be funded	23 spot-improvement projects unpriced
Illustrative or candidate price in a non-exploratory run	`IllustrativeRefused`; the gate aborts	The run is exploratory, so candidates are allowed and flagged
Link surface unknown	Treated as the lower of gravel or earth for access, flagged; reported as bounds	RAI 28.7 % to 93.5 %; 169,343 km listed as a survey requirement

Situation	Behaviour	Effect in the reference run
Condition survey missing	Not projected; network-level agency statistics used for the condition gap	Condition unknown on all 817,382 links
Traffic count missing	Modelled load used; never an invented count	All traffic modelled, uncalibrated
Node capacity unknown	Listed as capacity unknown; never assumed adequate	276,431 loaded links have unknown lanes
Sector feed missing	Documented fallback, flagged in the manifest and on every page that uses it	Minerals and trade feeds not used; freight growth at the fallback rate

3.12 Implications for planning

- The norm set is the engine's critical path and its weakest part: 578 of 646 rows are assumed. Agreeing and approving a first Ugandan norm set, with design standards from the road design manual, speeds and season rules from the road agencies, and access standards from the health and education sectors, would raise the quality of every result more than any other single step.
- Ownership of norms should follow ownership of policy. The access standards for health and education should be imported from the sector engines and signed off by those ministries, not retyped in the transport model.
- Every figure quoted from this report should carry its quality tier and basis. A modelled traffic load and a counted one are different things, and the engine's tags make the difference visible.
- The worst-quality rule means that a single assumed input marks a result as assumed. This is the correct incentive: it points data collection at the inputs that lower the most results.
- Condition and traffic observations are the empty tables that matter most. The engine is ready to receive link-level surveys and counts, with their dates and methods, and will use them without code changes.
- `explain()` works today for population; extending the factor trace to the needs, gap and cost modules would let any district figure be decomposed in a budget hearing, which is the purpose of the design.

CHAPTER 4

4 TI-01 Spatial base: people, places and the network

This chapter answers the first half of the needs question: *who and what must be connected, and where*. It describes the spatial base the engine builds before it computes anything else: the population on a grid, the network it routes on, the terrain and hazards along that network, and the sector nodes (schools, health facilities, markets, mines, industrial zones, gateways and administrative centres) that people and goods must reach. Every later module works on the objects created here, and no module after TI-01 introduces a new place. The chapter implements design Chapter 3 (`docs/03_TI01_SPATIAL_BASE.md`) and equations (1.1) to (1.6) of Chapter 23.

4.1 What the module does

TI-01 takes five kinds of input and turns them into the objects the chain works on.

- **Population.** The WorldPop 2020 constrained 100 m grid (42.2 million persons), aggregated to 175,367 populated 1 km cells, is raked to the UBOS district projection for each plan year. Each cell is a settlement, the unit at which access is computed.
- **The network.** The SDI noded road network, with OpenStreetMap attributes and surface evidence from satellite classification, is pinned as geography version `UGA-2026.09-c`: 817,354 road links and 28 ferry links.
- **Terrain and hazards.** A slope class for every link from the elevation model, flood exposure from the JRC 100-year flood map, and seasonal crossings and unrecorded river crossings from OpenStreetMap and HydroRIVERS.
- **Sector nodes.** 24,952 points from the health facility register, the school layer, markets, mines, industrial zones, border posts, airports and district headquarters.
- **Agricultural production.** MapSPAM 2010 production by crop on a grid, clustered into production zones that become freight origins.

Its outputs are the raked population of every cell and year, the snapped position and off-network distance of every settlement and node, the agricultural zones with their marketed surplus, and the destination set for each class of service. TI-01 also writes the district layers for population density and the raking factor that this chapter maps.

4.2 Method

4.2.1 Population on the ground

The planning population is gridded, not held at district level, because access is a question of distance. A district total says nothing about who lives 9 km from the nearest all-season road. The grid is raked to the official projection by district and urban or rural class, so that it sums to the official total while keeping the grid's pattern within each district:

$$\rho_{d,\,u,\,t} = \frac{P^{\mathrm{NSO}}_{d,\,u,\,t}}{\sum_{i \,\in\, I_{d,\,u}} P^{\mathrm{grid}}_{i,\,t_0}} \tag{4.1}$$

$$P_{i,\,t} = \rho_{d(i),\,u(i),\,t} \cdot P^{\mathrm{grid}}_{i,\,t_0} \tag{4.2}$$

where

- P with superscript NSO is the official projection for district *d*, class *u* (urban or rural) and year *t*;
- the denominator is the grid population of the cells *i* in that district and class in the grid's base year;
- ρ is the raking factor, logged for every district, class and year;
- P without superscript is the planning population of cell *i* in year *t*.

By construction the raked population sums exactly to the projection for every district, class and year. The engine warns when the raking factor lies outside 0.85 to 1.15, because a factor that large means the grid and the projection disagree about where or how many people live, and access results in that district are fragile.

A destination class may apply to a group rather than to everyone: emergency obstetric care to women aged 15–49, primary school to children aged 6–12. The group population is the cell population times a group share:

$$P^{g}_{i,\,t} = P_{i,\,t} \cdot a_{g,\,d(i),\,u(i),\,t} \tag{4.3}$$

where *a* is the share of group *g* in the district and class. In the reference run the shares are national (Chapter 5).

4.2.2 Snapping to the network

Every settlement and node is snapped to its nearest network link in the projected coordinate system (UTM zone 36N). The snap distance is the length of the perpendicular from the point to the nearest segment, or to the nearest vertex where the perpendicular falls outside the segment:

$$\delta_i = \min_{l \in L} \ \mathrm{dist}_{\mathrm{proj}}(x_i, \mathrm{geom}_l) \tag{4.4}$$

The snap distance becomes the off-network leg of every journey: it is walked at 5 km/h (an assumed norm, with terrain factors of 1.0, 1.15 and 1.4 for flat, rolling and mountainous ground) before the traveller reaches the network. A settlement more than 5 km from any mapped road is listed for review, because it is either a mapping error or one of the strongest access findings in the plan.

4.2.3 Agricultural production zones

Agricultural output arrives as a surface (tonnes per cell by crop), not as points. TI-01 keeps the cells whose production density is at least θ (20 t/km² in the draft norms), joins neighbouring cells into zones by rook contiguity, and computes each zone's marketed surplus:

$$Y^{\mathrm{surplus}}_{c,\,z,\,t} = \mu_c \sum_{i \,\in\, z} Y_{c,\,i,\,t} \tag{4.5}$$

where Y is production of crop *c* in cell *i*, and μ is the marketed share of the crop, a norm (for example 0.45 for maize, 0.20 for cassava, 1.0 for coffee, tea, sugarcane and cotton). A zone becomes a freight origin if its total surplus is at least 5,000 t a year. Its centroid is weighted by production.

4.2.4 Destination sets

For each class of destination the engine builds the set of functional nodes that satisfy it, by year:

$$J_k(t) = \{ n \in N : \mathrm{kind}_n \in \mathrm{Elig}(k),\ \ \mathrm{from}_n \le t < \mathrm{to}_n \} \tag{4.6}$$

where Elig(*k*) is the list of node kinds eligible for class *k*, a norm table, and *from* and *to* are the years a node is functional. A planned hospital is not a destination until its commissioning year. In the reference run every node is functional from 2000 onwards, because none of the registers supplies commissioning dates.

4.3 Results: the population base

4.3.1 How many people, and the raking factor

The base-year planning population is **52.0 million** in 2026, rising to **58.5 million** in 2030. It comes from the UBOS CODPS 2022 district totals (46.2 million for 135 districts) carried forward at an assumed 3 % a year, and split between urban and rural within each district by the grid. About 20 % of the planning population lives in the 2,292 cells the engine classes as urban (density of at least 1,500 persons per km²).

The grid itself sums to 42.2 million, so the national raking factor in 2026 is **1.23**. Almost all of it is a timing effect. The grid describes 2020 and the projection 2026. The ratio of the 2022 UBOS total to the 2020 grid is 1.096, and four years of assumed 3 % growth multiply that by 1.126. As a result **237 of the 254 district-class combinations in 2026 have a raking factor above 1.15**, and none below 0.85; across the five plan years there are 1,235 such warnings. The median factor is 1.22. The factor is highest in Karenga (1.81), Kakumiro (1.59), Nwoya (1.45) and Kaabong (1.41), and lowest in Obongi (0.86), Kotido and Kassanda (1.08 each) and Nakapiripirit and Kitagwenda (1.10 each).

The warning band was designed for a grid and a projection of the same year. Here it mostly flags the age of the grid. The districts at the extremes are still informative: in Karenga and Kakumiro the projection holds far more people than the grid places there, and in Obongi fewer. In those districts every access result is more fragile than elsewhere, because the people the projection adds are placed where the 2020 grid put people, not necessarily where they now live.

4.3.2 Where people live

Figure 4.1 shows population density by district, with the LISA clusters on the right. District density runs from 9.4 persons per km² in Kalangala, an island district that is mostly lake, and 17.8 in Buvuma to 10,493 in Kampala. Outside Kampala the densest districts are Mbale (1,415), Wakiso (1,381), Namisindwa (1,021), Bududa (1,019) and Manafwa (905). The sparsest mainland districts are Lamwo (31.8), Kaabong (32.1), Karenga (32.7) and Obongi (37.1). By region, the median district density is 453 in the Eastern region, 288 in the Western, 239 in the Central and 105 in the Northern region.

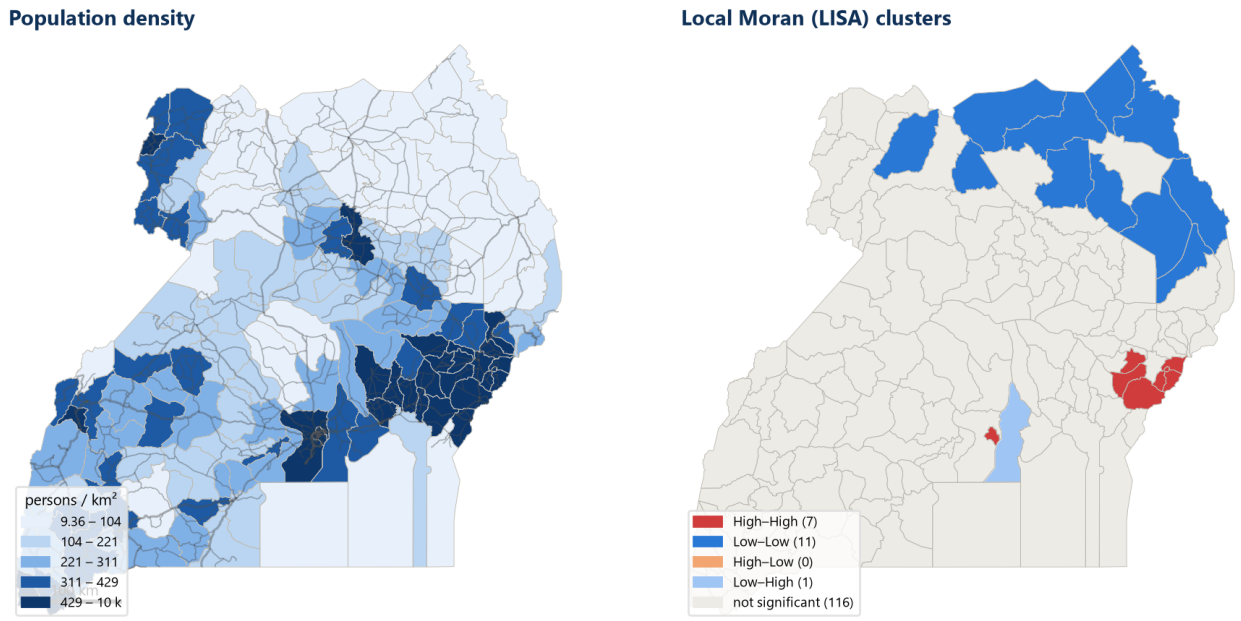

Figure 4.1: Population density by district; right: local Moran (LISA) clusters, 999 permutations

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Density is spatially clustered. The global Moran's I across the 135 districts is 0.100 (p = 0.001). The LISA map picks out two clusters. A **High–High** cluster of seven districts lies on the slopes of Mount Elgon and the plains below it (Bududa, Manafwa, Namisindwa, Budaka, Butaleja and Tororo) together with Kampala. A **Low–Low** cluster of eleven districts covers Karamoja and Acholi (Abim, Adjumani, Agago, Gulu, Kaabong, Karenga, Kitgum, Lamwo, Moroto, Nabilatuk and Napak). Mukono is a Low–High outlier: less dense than Kampala and Wakiso beside it. The pattern matters for the investment question. Where demand is dense and clustered, trunk and urban investment serve many people per kilometre. Where it is sparse and clustered, the same money buys feeder roads that serve few people each but may be their only all-season connection.

4.4 Results: the network

4.4.1 Network size and functional class

Table 4.1 gives the network in the geography snapshot by functional class and surface. The engine routes on **247,972 km** of road. Two-thirds of it (160,842 km) is community roads and a further 46,117 km is tracks. The national classes (trunk, primary and secondary) total 15,537 km and district roads 22,052 km.

Table 4.1: Road network in the geography snapshot: km by functional class and surface

Class	ac	dbst	lvsr	concrete	gravel	earth	unknown	Total km	Unknown share
NAT_TRUNK	641	338	0	0	0	1	2	981	0.2%
NAT_PRIMARY	1,994	920	0	4	27	457	12	3,415	0.4%
NAT_SECONDARY	1,072	1,458	0	0	657	7,624	330	11,140	3.0%
DISTRICT	303	121	0	0	673	10,151	10,804	22,052	49.0%
URBAN_LOCAL	0	0	0	0	0	23	44	67	65.2%
COMMUNITY	210	97	7	9	564	40,498	119,458	160,842	74.3%
TRACK	9	0	1	0	46	9,776	36,284	46,117	78.7%
SERVICE	45	34	1	3	5	858	1,687	2,634	64.1%
UNKNOWN	0	1	0	0	0	0	722	723	99.8%
Total	4,274	2,970	9	16	1,972	69,388	169,343	247,972	68.3%

Snapshot UGA-2026.09-c: SDI noded network with OSM surface tags; HeiGIT satellite and Mapillary surface evidence merged. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_network.csv).

Figure 4.2 draws every class. The national trunk roads (981 km) radiate from Kampala along the Northern Corridor to the Kenyan border, west to Fort Portal and Kasese, and south-west through Masaka and Mbarara. The national primary roads (3,415 km) complete the radial pattern to the north and east, through Gulu, Lira, Soroti and Moroto. The secondary roads (11,140 km) and district roads (22,052 km) form a dense mesh in the south, the west and the east, and a sparser one in the north. The community roads and tracks (grey) fill almost every inhabited area, and are visibly densest around Kampala, in the Elgon area and in the south-west highlands.

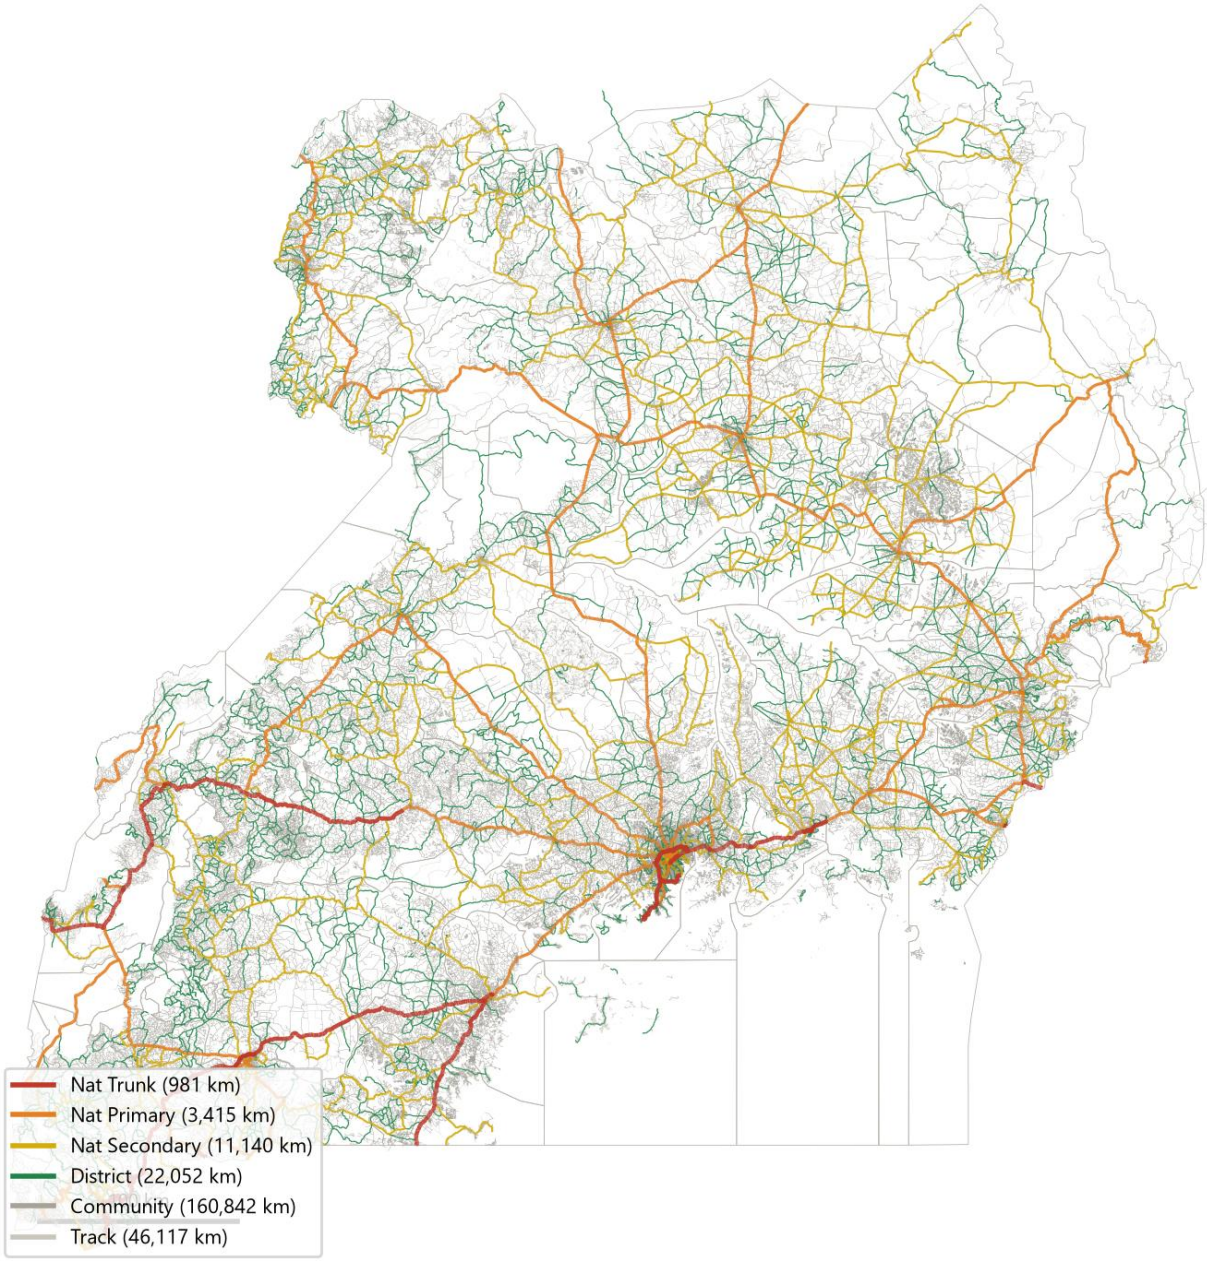

Figure 4.2: The road network in the engine, by functional class (all 248,000 km)

Source: SDI noded network; functional class from OSM highway tags and agency references

Network density varies twenty-fold between districts. The national average is 1.03 km of road per km² of district area (lakes included). The densest districts are Kampala (13.1), Mbale (3.5), Wakiso (3.5), Namisindwa (2.9), Butaleja (2.8) and Iganga (2.7). The sparsest are the island districts of Buvuma (0.06) and Kalangala (0.09), and in Karamoja Napak (0.11), Nabilatuk (0.12), Nakapiripirit (0.15) and Abim (0.20). Of the national and district classes, the Northern region holds 11,742 km and the Western region 11,337 km, against 7,772 km in the Central and 6,737 km in the Eastern region, because the northern districts are large.

The functional classes come from OpenStreetMap road tags and agency road numbers, not from the gazette. The engine's national classes (15,537 km) are therefore shorter than the 21,200 km of national roads the agencies report, and its district class (22,052 km) shorter than the 41,194 km of gazetted district roads. Some gazetted national and district roads are tagged as lower classes in the map. Conversely, the community class (160,842 km) is more than twice the 75,404 km of community access roads in the agency statistics, because the map includes paths and minor roads no agency maintains. Only 11,118 links (3,431 km) carry an agency road number. Matching

agency numbers and chainages to the network is the only reliable way to align the two, and it is listed as a data task in Chapter 20.

4.4.2 Surface: what is known

Figure 4.3 shows the km of each class by surface, and Figure 4.4 the share of each class with a known surface. The engine knows the surface of **78,629 km, 31.7 % of the network**. Of these, 4,274 km are asphalt concrete, 2,970 km double surface treatment, 16 km concrete and 9 km low-volume seal: 7,269 km paved in all. Another 1,972 km are gravel and 69,388 km earth. The remaining **169,343 km (68.3 %)** have no surface record.

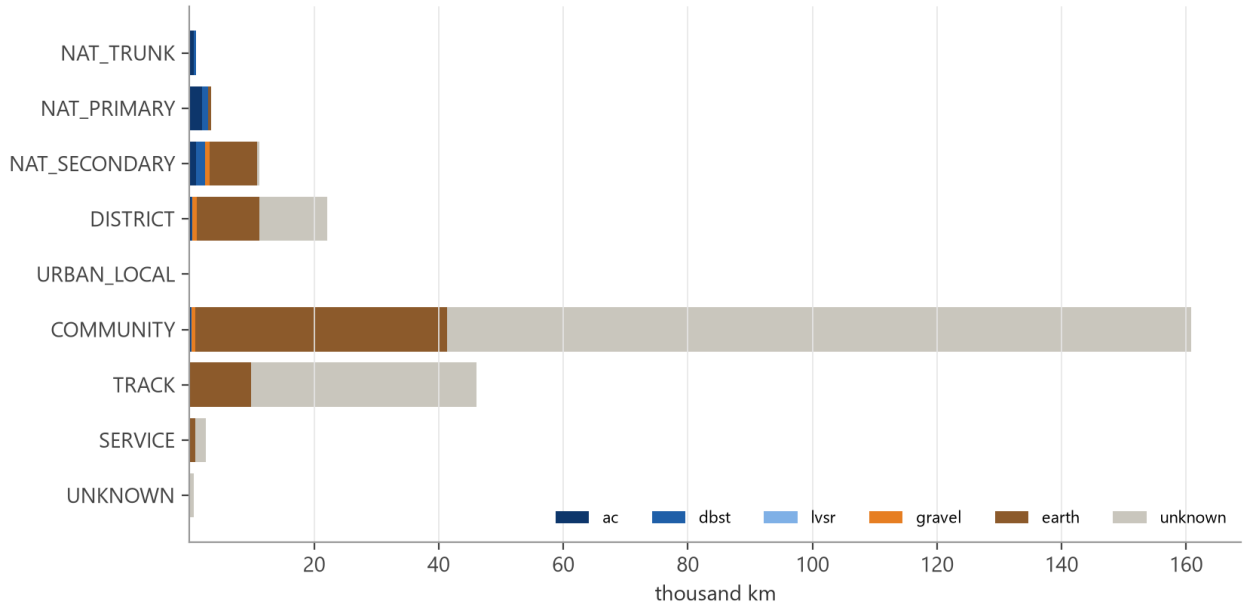

Figure 4.3: Network km by functional class and known surface

Source: Geography snapshot

The unknown surface is concentrated in the lower classes. Surface is known on 100 % of national trunk and primary roads and 97 % of secondary roads, on 51 % of district roads, 26 % of community roads and 21 % of tracks. In other words, the engine knows the national network well and the local networks badly, which is the opposite of where the access problem lies.

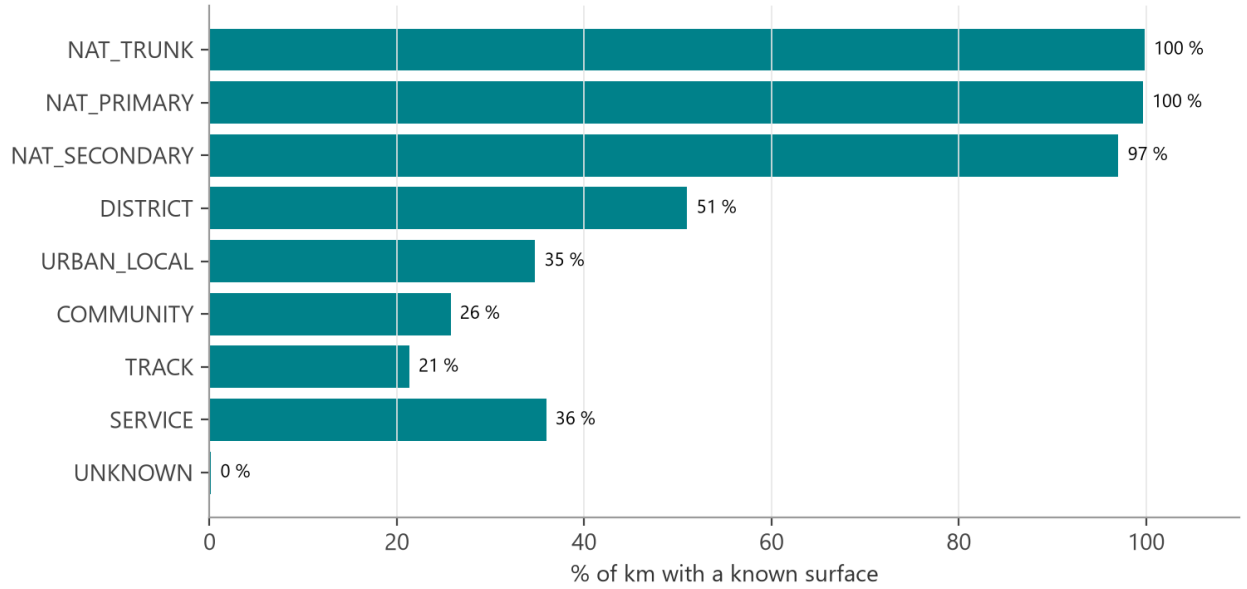

Figure 4.4: How much of each road class has a known surface

Source: Geography snapshot

The surface evidence has four sources, each with its own basis:

- **OpenStreetMap surface and track-type tags**, basis observed: about 74,000 km of links, including the ferry routes.
- **HeiGIT 2024 satellite classification** (a deep-learning paved and unpaved classification of PlanetScope imagery), basis estimated: 1,879 km classed paved and 2,706 km unpaved. Unpaved is recorded as earth unless GRIP4 agrees it is gravel (36 km). The HeiGIT data carry a non-commercial licence.
- **GRIP4** (the global roads inventory, from WFP 1997 and OSM 2014 sources), basis estimated: 210 km where a whole link lies within 40 m of a GRIP segment (150 m for national and district classes).
- **No source**: 169,343 km, recorded as unknown. The engine never infers paved status from class.

Figure 4.5 maps the surface of the national and district roads only (37,588 km): 6,851 km paved (blue), 1,356 km gravel (orange), 18,233 km earth (brown) and 11,148 km unknown (grey). The paved network is the radial trunk and primary system and a handful of secondary links. Earth dominates the district roads of the north and of the central region. Gravel appears in pockets: in West Nile, in Bundibugyo and Kasese, around Soroti and in Kisoro and Kabale. The grey of unknown surface is spread across all regions but is heaviest in the centre and east.

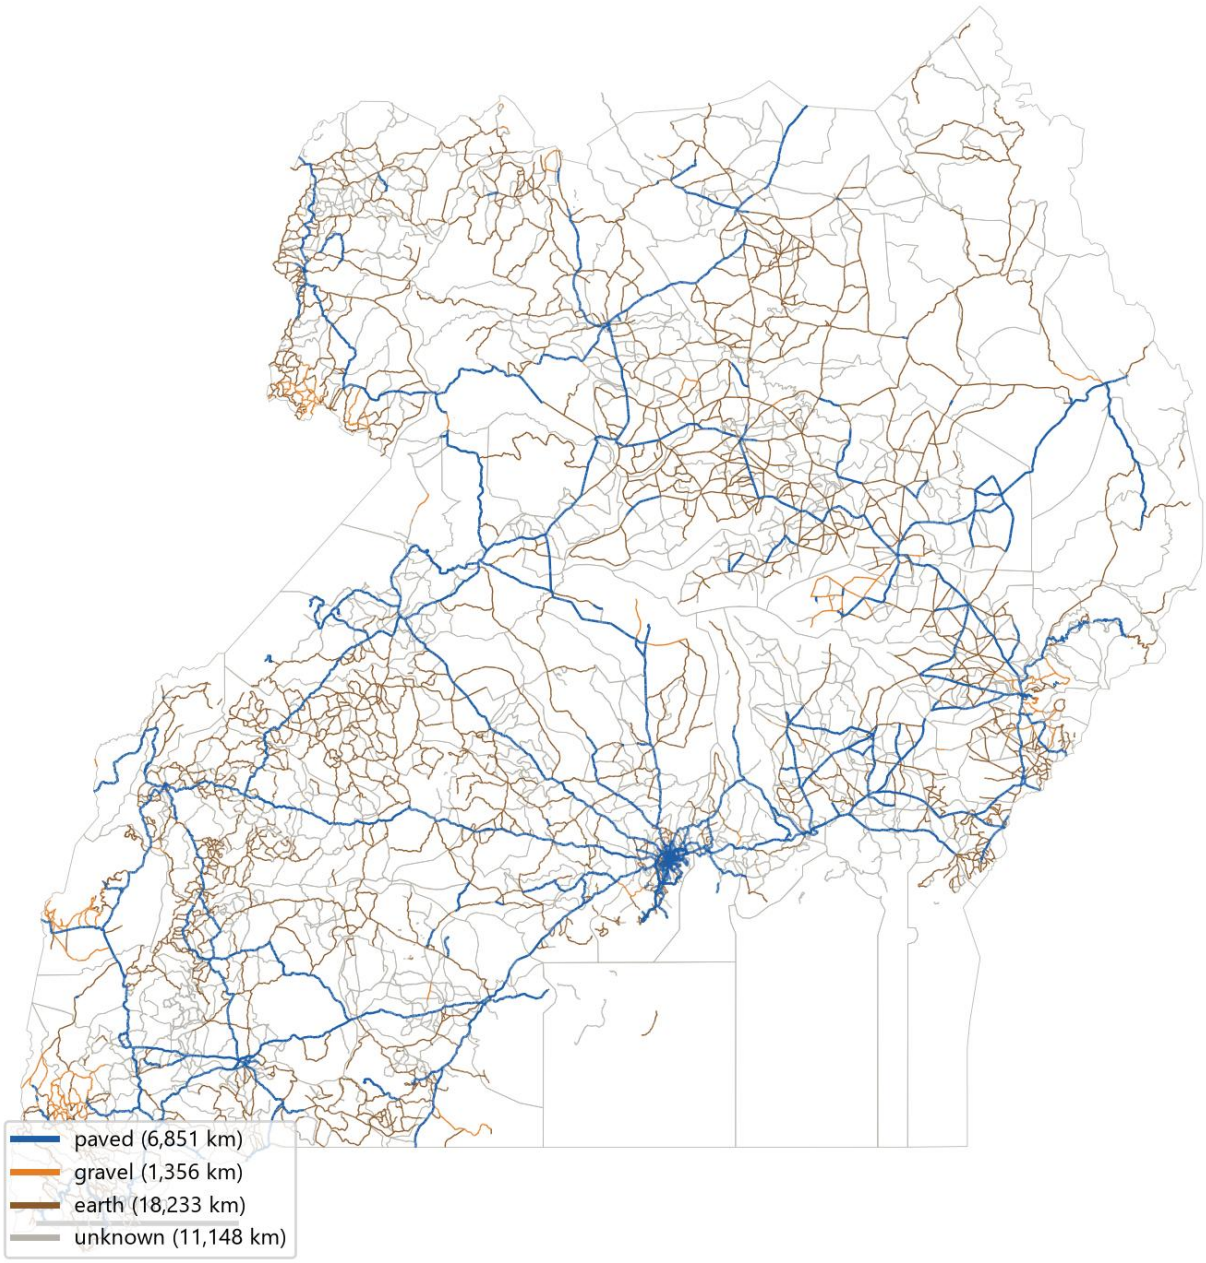

Figure 4.5: Road surface as known to the engine: national and district roads (unknown surface in grey)

Source: OSM surface tags; HeiGIT satellite and Mapillary evidence

By region the known share of all network km is 52 % in the Western region, 28 % in the Northern, 23 % in the Central and 21 % in the Eastern region. By district, the median known share of national and district roads is 72 %, but it is 0 % in Buvuma, 9 % in Buhweju, 16 % in Buyende, 22 % in Kalangala and 25 % in Mitooma. Across all classes, a few districts in the Rwenzori and Kigezi areas are almost fully known (Kanungu, Kisoro and Kasese 99 %, Bundibugyo 96 %, Bunyangabu 95 %), because mappers have tagged their roads, while in Buvuma (2 %), Kayunga, Amuria and Buhweju (5 % each) the surface of almost every community road is unknown.

The share of unknown surface (68.2 % of network km) is far above the 30 % threshold of rule GIS-5. The run is therefore exploratory, and every access figure carries the warning. Chapter 7 shows what this does to the Rural Access Index: the bounds of 28.7 % and 93.5 % differ only in what is assumed about these 169,343 km.

4.4.3 Condition: unknown on every link

No link in the network has a recorded condition. No link-level condition survey exists in the estate: the agency road asset management systems are not shared. The engine therefore does not model condition per link, and it does not project deterioration. Condition enters only as the agency statistics by network (Chapter 7), and is never spread over links.

4.4.4 Terrain

Figure 4.6 shows the terrain class of every link, from the slope along the link in the SDI elevation model: flat below a 3 % grade, rolling from 3 % to 8 %, mountainous above. Most of the network is flat: **189,520 km (76.4 %)**. Rolling links total 45,893 km (18.5 %) and mountainous links 12,559 km (5.1 %).

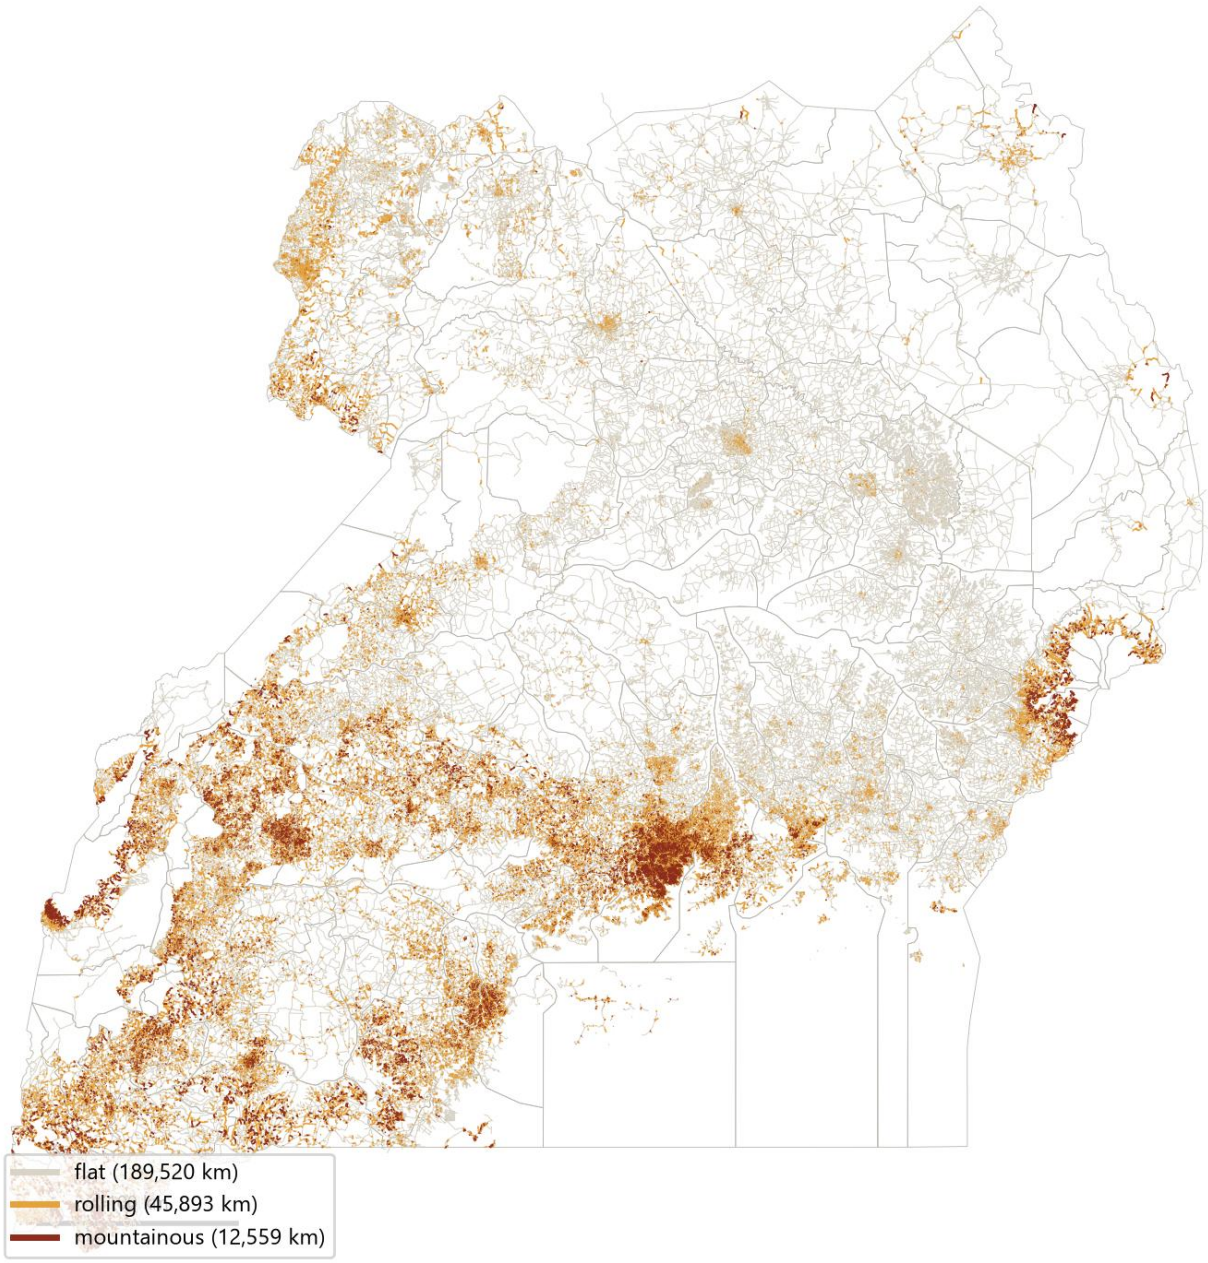

Figure 4.6: Terrain class of every road link (from the DEM slope)

Source: SDI DEM; TIPRAS terrain classification

The mountainous network is concentrated in a few areas, clearly visible on the map: the Elgon massif in the east, the Rwenzori foothills and the south-western highlands of Kigezi, the hills of Buganda north-west of Kampala, and

the Kigezi and Ankole uplands. By district, the mountainous share of network km is highest in Bududa (47 %), Rubanda (34 %), Kabale (28 %), Namisindwa (24 %), Kasese (23 %), Buhweju and Rukiga (20 % each). The national trunk roads are 81 % flat and only 5 of their 981 km mountainous. Terrain enters the engine in three ways: through walking and driving speeds, through season rules for earth roads, and through the unit cost of works. Its weight in the costs is limited, because no source in the price evidence splits paved upgrading by terrain (Chapter 9).

4.4.5 Hazards and crossings

Figure 4.7 shows the two hazards the engine carries on links. The first is **flood exposure**: 3,006 links totalling **1,368 km** intersect the JRC 100-year flood extent. Flood exposure is concentrated in a few places. Ntoroko, on the floodplain of the Semliki and Lake Albert, has 439 km of flood-exposed road, a third of the national total. Adjumani (93 km), Kotido (76 km), Obongi (53 km), Arua (38 km), Serere (35 km), Katakwi (34 km) and Kayunga (33 km) follow, mostly along the Albert Nile, the Kyoga basin and the seasonal rivers of Karamoja. By class, community roads (525 km) and tracks (518 km) carry most of the exposure, district roads 196 km and national secondary and primary roads 105 km. No trunk road is exposed.

The second is **seasonal crossings**: 2,248 links where a ford or low-water crossing is recorded (from 250 ford ways and 12 low-water crossings in OpenStreetMap, on top of 3,943 ford points snapped within 25 m). On these links an earth or gravel road may close in the wet season. The map shows them as red points, densest along the Rwenzori and the western rift (Kasese 296, Bunyangabu 97, Ntoroko 87), in West Nile (Nebbi 163, Zombo 113), in Kigezi (Kanungu 94, Rukungiri 92) and on Elgon (Namisindwa 86). Almost all are on community roads (1,103) and tracks (870); 137 are on district roads and 41 on national roads.

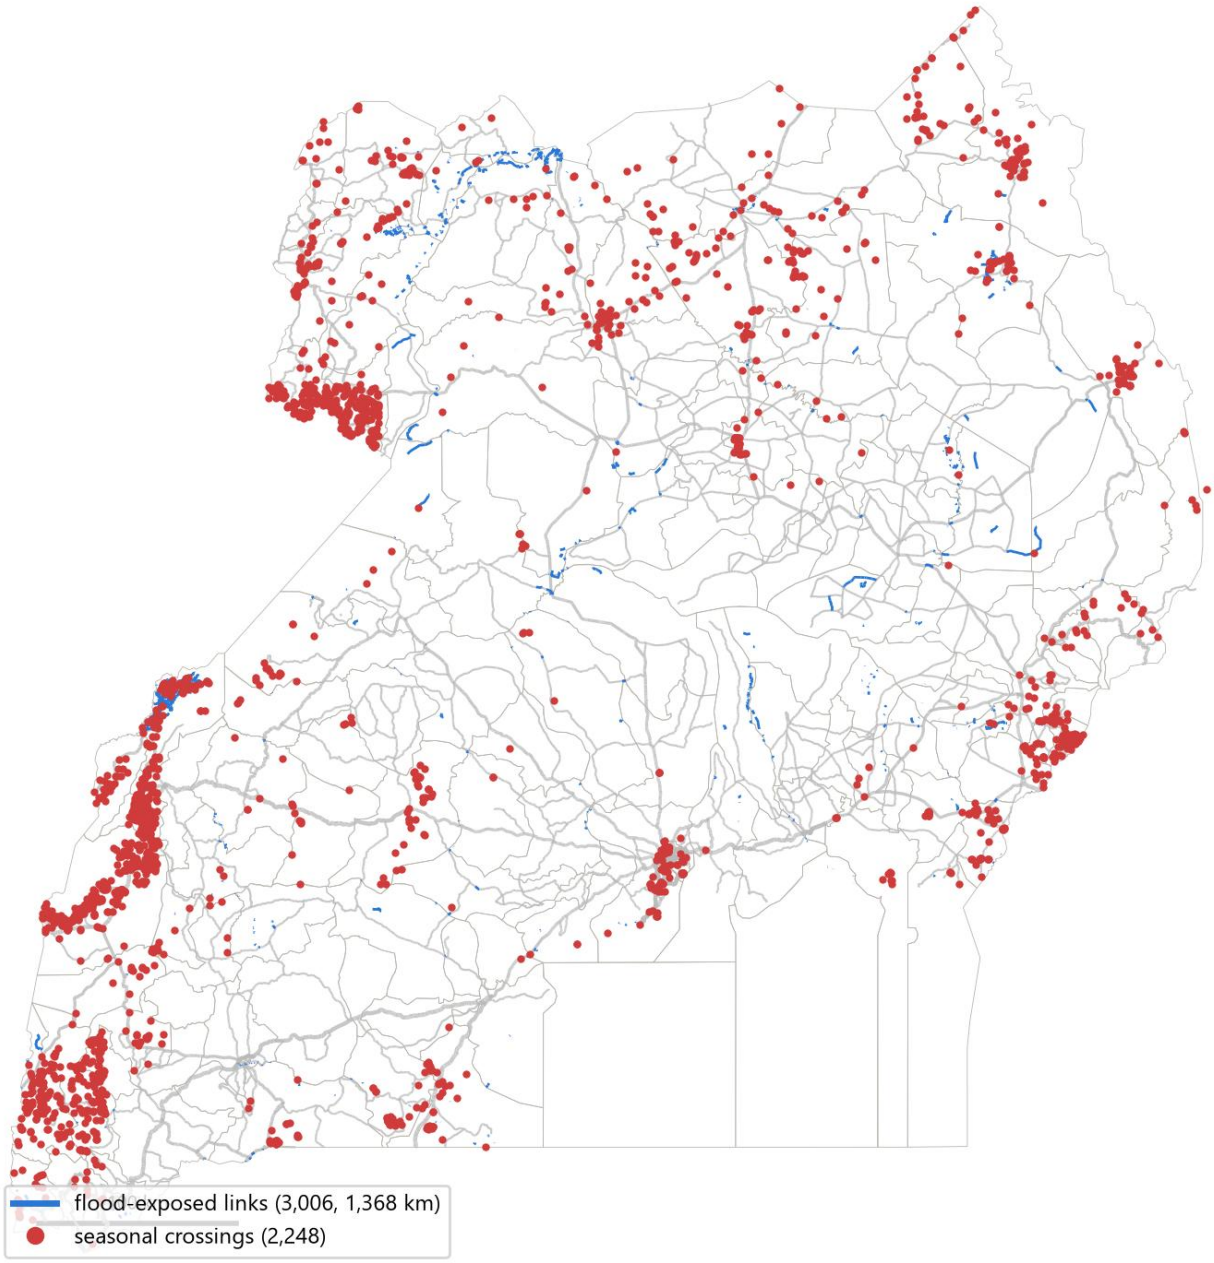

Figure 4.7: Hazard exposure of the network: flood-exposed links (JRC 100-year) and seasonal river crossings

Source: JRC/CEMS flood hazard RP100; HydroRIVERS; OSM fords

Structures are thinly recorded. The snapshot holds 2,188 bridges and 5,039 culverts on 3,676 links, from OpenStreetMap, without span, load rating or condition. A further **1,958 river crossings on 1,849 links** have no recorded bridge, culvert or ford: the road crosses a mapped river and nothing is known about how. The engine does not treat these as closures. It lists them as a structure survey requirement (Chapter 8). The largest numbers are in Zombo (106), Yumbe (104), Gulu (85), Kampala (84), Mukono (73) and Moroto (68).

4.5 Results: the places that must be reached

4.5.1 Sector nodes

Figure 4.8 counts the 24,952 sector nodes by kind, on a logarithmic scale, and Figure 4.9 maps each family. They come from six sources, each with its own coverage:

- **Schools: 20,474** (16,880 primary and 3,594 secondary), from the Ministry of Education school layer held by SDI. These are the only nodes with a magnitude: the enrolment field gives 7.09 million primary and 0.65 million secondary pupils.
- **Health facilities: 3,169**, from the HealthModellingx facility register: 1,492 health centres II, 1,358 health centres III, 202 health centres IV, 102 hospitals and 15 regional referral hospitals.
- **Markets: 814**, of which 722 come from TransportModelling's points of interest and 92 from OpenStreetMap marketplaces.
- **Mines and quarries: 309**, from the same points of interest, without commodity or output.
- **Administrative centres: 135**, one district headquarters per district, from the COD-AB administrative points.
- **Gateways: 28**: 22 border posts from sparse OpenStreetMap border-control tags and 6 airports from OurAirports.
- **Industrial zones: 23**, special economic zones from the points of interest, without output.

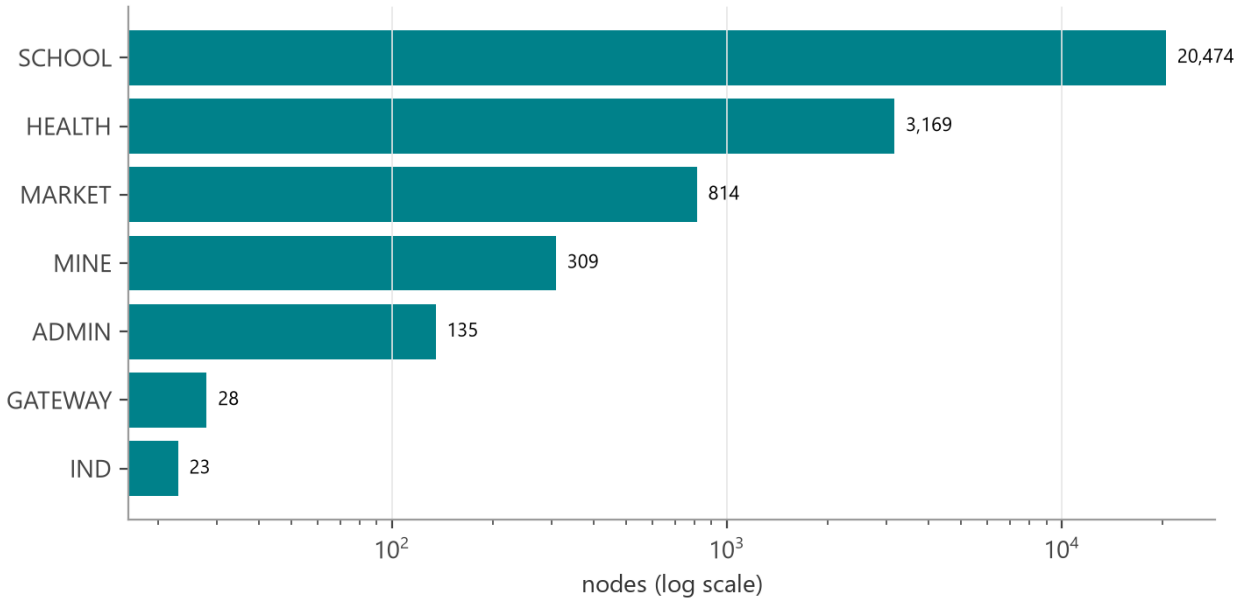

Figure 4.8: Sector nodes by kind

Source: Geography snapshot (nodes.arrow)

The maps in Figure 4.9 show that coverage differs as much as numbers do. Health facilities cover the whole country, densest in the south-west and the east and sparse in Karamoja. Schools form an almost continuous cover in the south and the east, with gaps in the north-east, the parks and around Lake Kyoga. Markets, by contrast, are strongly clustered: a dense cluster in West Nile (Arua and its neighbours), smaller ones around Kampala and in the far south-west, and scattered points elsewhere. The cluster in West Nile reflects where the source data were collected more than where markets are, and it biases every market access result. Mines and quarries cluster around greater Kampala, in the north-west and along the northern border, and in the western and south-western highlands, with scattered points along the eastern border and in Karamoja. The industrial zones and gateways are few and spread along the main corridors and borders.

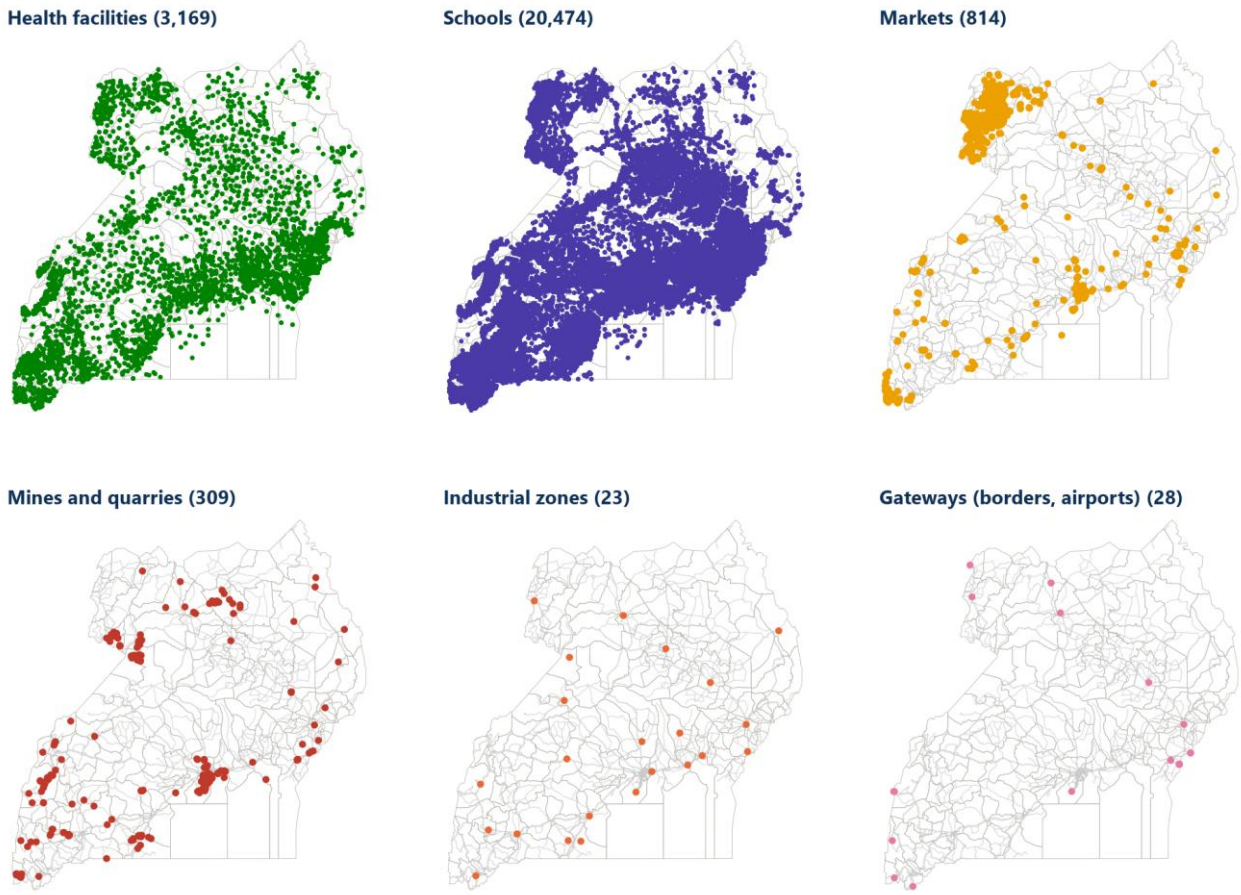

Figure 4.9: Sector nodes the network must serve: health, schools, markets, mines, industry and gateways

Source: Health facility registry; EMIS; markets; mining cadastre; industrial parks; borders and airports (SDI)

Two gaps in the nodes limit later chapters. First, apart from schools, no node carries a magnitude: no catchment population for health facilities, no throughput for markets, no output for mines or industrial zones, no traffic for borders. Mines and industrial zones therefore enter the engine as places that need a connection, not as generators of freight. Second, no node carries a commissioning or closure date, so the destination sets do not change over the plan period.

4.5.2 Destination sets

From the nodes and the eligibility norms, TI-01 builds one destination set per class:

- **All-season road**: any all-season link, a line rather than a point;
- **Primary health**: all 3,169 health facilities;
- **Emergency obstetric and newborn care (EmONC)**: the 319 health centres IV, hospitals and regional referral hospitals. This uses facility level as a proxy; the register does not say which facilities actually provide comprehensive emergency obstetric care;
- **Primary school**: 16,880 schools;
- **Secondary school**: 3,594 schools;
- **Market**: 814 markets;
- **Administrative centre**: 135 district headquarters.

4.5.3 Agricultural production zones

The MapSPAM 2010 grid holds 24.5 million t of production in 45,596 cells. Banana is the largest crop by weight (8.4 million t), then sugarcane (3.4 million t), cassava (2.9 million t), sweet potato (2.7 million t) and maize (2.4

million t). After applying the density threshold, the contiguity rule and the marketed shares, TI-01 finds **10 agricultural zones** with a marketed surplus of at least 5,000 t, together **9.76 million t** a year.

Ten zones is very few for a country with Uganda's agriculture. The reason is the contiguity rule. At a threshold of 20 t/km², most of the cropped south, east and west forms a few very large connected areas, each reduced to one production-weighted centroid. The whole surplus of a region then leaves from a single point, which concentrates freight on a few roads (Chapter 6). The threshold and the minimum zone size are assumed norms. A higher density threshold, or zones bounded by sub-counties, would give a more realistic pattern of freight origins. The production data are also old: MapSPAM describes 2010, not 2026.

4.5.4 Distance to the mapped network

Figure 4.10 shows the distance from each settlement cell to the nearest mapped road link, of any class. The network is dense: the median cell lies **0.24 km** from a road, and the population-weighted mean is 0.22 km. The distribution has a long tail:

- 46,820 cells, holding 4.77 million people (9.2 % of the population), are more than 0.5 km from a road;
- 18,416 cells, with 1.49 million people (2.9 %), are more than 1 km away;
- 5,661 cells, with 430,035 people (0.8 %), are more than 2 km away;
- **1,039 cells, with 78,059 people (0.15 %), are more than 5 km away**, the review threshold; 211 of them, with 19,054 people, are more than 10 km away.

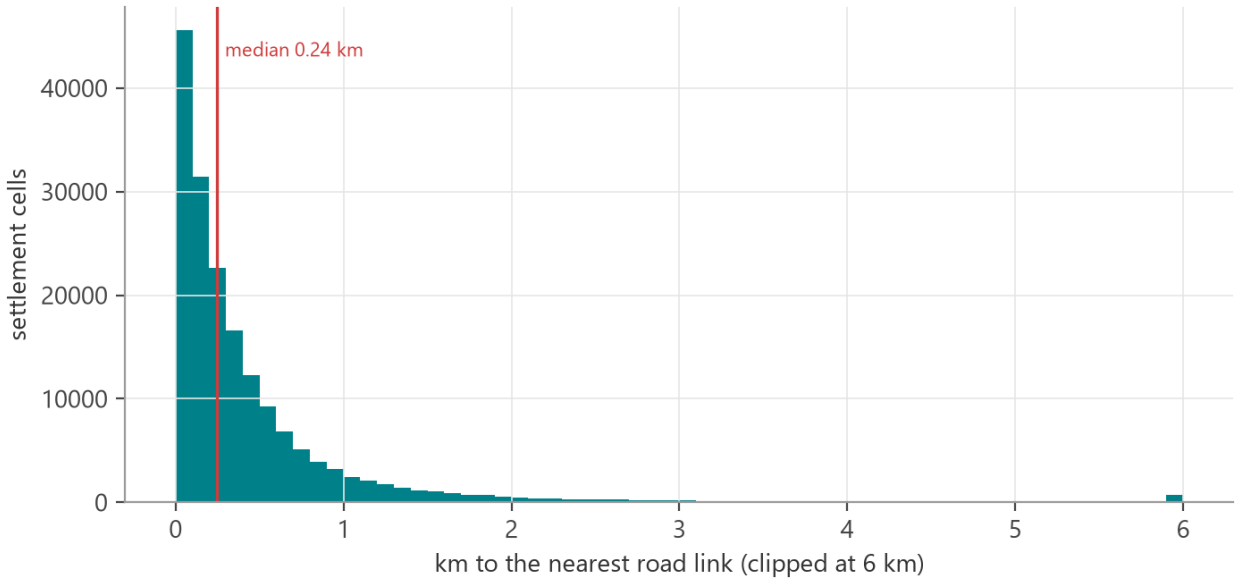

Figure 4.10: Distance from each settlement to the nearest mapped road

Source: TI-01 snapping (KD-tree, segment projection)

The far cells are concentrated in Karamoja. Napak has 325 cells more than 5 km from any mapped road (27,328 people), Kaabong 166, Moroto 116, Nakapiripirit 100 (13,970 people), Kotido 78, Abim 42 and Amudat 20. By district the population-weighted distance to the network is highest in Napak (2.05 km), Nakapiripirit (1.81 km), Kalangala (1.43 km, the islands), Amudat (1.20 km) and Nabilatuk (0.85 km), and lowest in Kampala (0.03 km) and Wakiso (0.05 km).

These distances are to *any* mapped road, including tracks of unknown surface. They are not distances to an all-season road, which is what the Rural Access Index measures (Chapter 7). They show where the map has no road at all near people, which in Karamoja may mean that the map is incomplete as much as that the roads are missing. A person should look at each of the 1,039 far cells before any of them is read as an access finding.

4.6 Data quality and limitations

Table 4.2 summarises the quality of the spatial base.

Table 4.2: Quality of the TI-01 spatial base in the reference run

Element	Basis and quality	Main limitation
Population grid	WorldPop 2020 constrained, estimated	Describes 2020; raked by a national factor of 1.23
Population projection	UBOS CODPS 2022 district totals, carried at an assumed 3 % a year	No 2024 census tables or projections by year in the estate
Network geometry	SDI noded network from OpenStreetMap, observed	Functional class from map tags, not the gazette
Surface	Known on 78,629 km, of which 4,831 km estimated from imagery or GRIP4; unknown on 169,343 km	Unknown on 68.3 % of km and on 74 % of community roads
Condition	None per link	Agency statistics by network only
Terrain	From the SDI elevation model, estimated	Grade thresholds are assumed norms
Flood exposure	JRC 100-year flood map, estimated	Fluvial flood only; no landslide or pluvial layer on links
Structures	OSM bridges and culverts, observed	No span, load or condition; 1,958 unrecorded crossings
Sector nodes	Registers and points of interest, observed locations	Magnitudes only for schools; no commissioning dates; market data clustered in West Nile
Agricultural zones	MapSPAM 2010 with assumed marketed shares	Old data; the contiguity rule yields only 10 zones

Four changes would most improve the spatial base. A road inventory with agency road numbers, chainages, surface and condition for district and community roads would close the largest gap. UBOS projections by district and year, or the 2024 census tables, would replace the assumed growth. Facility catchments and EmONC status from HealthModellingx, and market throughput from the agriculture model, would give the nodes their magnitudes. A current production surface and a sub-county zone rule would give a more realistic set of freight origins.

4.7 Implications for planning

- The spatial base is complete enough to compute access for all 52.0 million planning persons and all 24,952 nodes, but its two largest gaps (surface on 169,343 km and condition on every link) sit on the district and community roads, where access is decided.
- A surface survey of district and community roads, matched to agency road numbers, would narrow the Rural Access Index range more than any other investment in data, and should precede any large programme of upgrading on those networks.
- Population is the base of every requirement. The 2024 census results, or UBOS projections by district and year, should replace the assumed 3 % growth before the plan is used for allocation; the raking warnings in Karenga, Kakumiro, Nwoya and Kaabong show where the current figures are most fragile.
- Flood and seasonal-crossing exposure is concentrated: Ntoroko alone has a third of the flood-exposed km, and Kasese, Nebbi and Zombo the most seasonal crossings. Resilience works can be targeted geographically from this base.
- The 1,958 unrecorded river crossings and the 1,039 settlement cells more than 5 km from any mapped road are review lists, not findings. Each should be checked on the ground or against imagery before it enters a programme.
- Sector nodes need magnitudes and dates from their owning ministries. Without them, health, market, mining and industrial needs are reduced to "connect this point", and the engine cannot weigh one node against another by the people or tonnes it serves.

CHAPTER 5

5 TI-02 Needs: who must be connected, and to what

This chapter answers the core of the first planning question: *what standard of connection each place requires, and how much movement it generates*. It turns the people and places of the spatial base into two kinds of requirement. The first is an **access requirement**: every settlement must reach each class of destination (an all-season road, a health facility, a school, a market, an administrative centre) within a standard time or distance, by an assumed mode, in the wet season. The second is **movement generation**: every production, trade and population node generates tonnes and trips a year. Needs are stated independently of the existing network. The question is what is required, not what is used. The chapter implements design Chapter 4 (`docs/04_TI02_NEEDS.md`) and equations (2.1) to (2.6) of Chapter 23.

5.1 What the module does

TI-02 reads the raked population, the sector nodes, the agricultural zones and the destination sets from TI-01, and the access standards, group shares, node connection standards and generation rates from the norm tier. It produces three tables:

- **the access requirement**: for every settlement cell, destination class and plan year, the number of people who require that class, with the threshold, measure, mode and season of the standard that applies;
- **the node connection requirement**: for every production, gateway and referral node, the type of connection it requires (road class, surface, axle load, maximum connector length);
- **generation**: the tonnes of freight and the passenger trips each node and settlement generates, by commodity or purpose and year.

TI-02 does not say whether anyone meets a standard. That is TI-04's question, answered on the wet-season network in Chapter 7. The gap between the two is TI-05's, in Chapter 8.

Why insist that needs ignore the network? Because a district whose roads carry little traffic may carry little because it has no roads. Planning from observed use alone entrenches the existing network. A need computed from people and services cannot be talked down by pointing at empty roads.

5.2 Method

5.2.1 Access standards

An access standard is a row of data in the norm table `access_standard`, keyed by destination class, settlement type (rural or urban) and population group. It sets a measure (distance in km or time in minutes), a threshold, a mode (walk, motorised or mixed) and a season. Every Uganda standard is tested in the **wet season**, because testing an all-season standard against the dry-season network is the most common way access is overstated. Table 5.1 lists the standards of the reference run.

Table 5.1: Access standards in norm version UGA-2026.09-draft-f

Destination class	Rural standard	Urban standard	Group	Source	Quality
All-season road	2 km, walking	0.5 km, walking	everyone	Rural Access Index, SDG 9.1.1 (urban: illustrative)	rural normed, urban assumed
Primary health facility	60 min, mixed modes	30 min	everyone	national health policy, to be imported from HealthModellingx	assumed

Destination class	Rural standard	Urban standard	Group	Source	Quality
EmONC facility	120 min, motorised	60 min	women aged 15–49	WHO/Lancet 2-hour standard (urban: illustrative)	rural normed, urban assumed
Primary school	3 km, walking	1.5 km	children aged 6–12	education sector norm, to be imported from EducationModellingx	assumed
Secondary school	60 min, mixed	30 min	youth aged 13–18	illustrative	assumed
Market	120 min, motorised	30 min	households	agriculture sector norm, illustrative	assumed
Administrative centre	180 min, motorised	60 min	everyone	illustrative	assumed

Only two of the fourteen standard rows are normed: the rural all-season road standard and the rural emergency obstetric care standard. The design rule is that the health and education standards are **the same rows those sectors' engines use**, imported and not retyped, so that a change of sector standard changes transport need with no code change. That import has not yet been made, and the primary health and primary school standards remain placeholders that carry the right labels but assumed values.

The engine resolves each standard by the most specific row available (district, then region, then urban or rural, then national), and raises an error if none exists:

$$\tau_{k,\,u} = \mathrm{resolve}(\mathrm{access_standard};\ k, u, g) \tag{5.1}$$

where τ is the threshold for destination class *k*, settlement type *u* and group *g*.

5.2.2 The access requirement

For each settlement cell *i*, class *k* and year *t*, the requirement is the number of people in the group the standard applies to:

$$\mathrm{need}_{i,\,k,\,t} = P_{i,\,t} \cdot a_{g(k)} \tag{5.2}$$

$$\mathrm{Req}_{k,\,d,\,t} = \sum_{i \in I_d} \mathrm{need}_{i,\,k,\,t} \tag{5.3}$$

where

- P is the raked population of the cell (Chapter 4);
- *a* is the share of the group *g(k)* that class *k* applies to;
- Req is the requirement of district *d*, the sum over its cells.

Summed over the district, Req is the first answer to Q1: how many people require each kind of access. TI-04 later tests each cell against the standard. For time standards the test is whether the travel time is within the threshold; for the all-season road standard it is whether the distance to the nearest link open in the wet season is within 2 km (0.5 km in urban cells):

$$A_{i,\,k,\,t} = \mathbf{1}\left[tt_{i,\,k,\,t} \le \tau_{k,\,u(i)}\right], \qquad A_{i,\,\mathrm{ASR},\,t} = \mathbf{1}\left[\mathrm{dist}(i, L_t^{\mathrm{wet}}) \le \tau_{\mathrm{ASR},\,u(i)}\right] \tag{5.4}$$

The group shares in the reference run are national and assumed, pending the age structure from UBOS or the household model: 0.235 for women aged 15–49, 0.21 for children aged 6–12 and 0.15 for youth aged 13–18. For the market standard the unit is the household, and the share of 0.222 converts persons into households (one household per 4.5 persons). The market requirement is therefore a count of households, not of persons.

5.2.3 Node connection standards

Production, gateway and referral nodes need a *type* of connection rather than a travel time: a minimum road class and surface on the path to a target, a design axle load, and a maximum length of lower-standard connector. The engine selects, for each node, the row of `node_connection_standard` with the highest magnitude band the node reaches:

$$CS(n) = \text{row } r \;:\; \kappa_r = \kappa_n, \quad m_r = \max\{m_{r'} : \kappa_{r'} = \kappa_n,\ m_{r'} \le \text{mag}_n\} \tag{5.5}$$

where κ is the node kind, *m* the minimum magnitude of a row and mag the node's magnitude. A node that matches no row gets a default all-season connection and is flagged. Table 5.2 lists the standards of the reference run. All are assumed.

Table 5.2: Node connection standards in the reference run

Node kind	Applies when	Required connection	Target
Agricultural zone	surplus ≥ 5,000 t/yr	district class or better, gravel, 8.2 t axle, earth connector ≤ 5 km	nearest market
Agricultural zone	surplus ≥ 50,000 t/yr	national secondary or better, double surface treatment, 10 t axle, connector ≤ 2 km	national road
Mine	output ≥ 100,000 t/yr	national primary, double surface treatment, 11.5 t axle	export gateway
Mine	output ≥ 2 Mt/yr	national trunk, rail or road, 11.5 t axle	export gateway
Industrial zone	any	national primary, double surface treatment, 11.5 t axle	gateway
Gateway (border post, airport)	any	national trunk, double surface treatment, 11.5 t axle	trunk network
Hospital	any	district class or better, gravel, 8.2 t axle	district network
Regional referral hospital	any	district class or better, double surface treatment, 8.2 t axle	district network

5.2.4 Freight generation

Each production node generates outbound freight and each consumption centre attracts inbound freight:

$$O^{c}_{n,t} = \mu_c \, Y_{c,n,t}, \qquad D^{c}_{m,t} = q_c \, P_{m,t} \, (1 - s_{c,m}) \tag{5.6}$$

where

- O is the outbound freight of commodity *c* from production node *n*, the marketed share μ times production Y;
- D is the inbound freight of consumption centre *m*, the per-capita consumption *q* times population, less a self-supply share *s*.

Before distribution the two sides are balanced. The consumption side is scaled to the production side, and the scaling factor is logged:

$$\lambda_c = \frac{\sum_n O^{c}_{n,t}}{\sum_m D^{c}_{m,t}} \tag{5.7}$$

A factor outside 0.8 to 1.25 is a warning that production and consumption data disagree. Where a sector engine supplies node-to-node flows (as MineralModelling is designed to), the engine takes them as given and skips generation for that sector.

5.2.5 Passenger generation and growth

Passenger trips are generated per settlement from a trip rate by settlement type, grown with income where an income path exists:

$$O^{p}_{i,t} = \rho_{p,u(i)}\, P_{i,t} \left(\frac{\bar{y}_t}{\bar{y}_{t_0}} \right)^{\varepsilon_p} \tag{5.8}$$

where ρ is the trip rate for purpose *p* in settlement type *u*, ȳ is GDP per capita from the fiscal and macroeconomic models, and ε is the income elasticity of trips. Movement grows with its drivers: population from TI-01, sector output from the sector models and income from FINEX or CDCGE. A flat traffic growth rate is prohibited as a method and may appear only as a flagged fallback.

5.3 Results: the access requirement

5.3.1 National totals

Table 5.3 and Figure 5.1 give the base-year requirement by class. Three classes apply to everyone, so they require the whole planning population of **52.0 million**: an all-season road within 2 km (0.5 km in towns), a primary health facility within 60 minutes (30 in towns), and the district headquarters within 3 hours (1 hour in towns). The other classes apply to groups:

- **12.2 million women aged 15–49** must be able to reach a facility offering emergency obstetric and newborn care within 2 hours by motorised transport (1 hour in towns);
- **10.9 million children aged 6–12** must be able to walk to a primary school within 3 km (1.5 km in towns), in the wet season;
- **7.8 million young people aged 13–18** must reach a secondary school within 60 minutes (30 in towns);
- **11.6 million households** must reach a market within 2 hours by motorised transport (30 minutes in towns).

Table 5.3: People who must be connected to each destination class, base year

Destination class	People who must reach it (millions)
All-season road within 2 km	52.01
Primary health facility	52.01
EmONC facility (2 h)	12.22
Primary school	10.92
Secondary school	7.80
Market	11.56
Administrative centre	52.01

TI-02 access requirement; population from UBOS projections on WorldPop. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_requirement.csv).

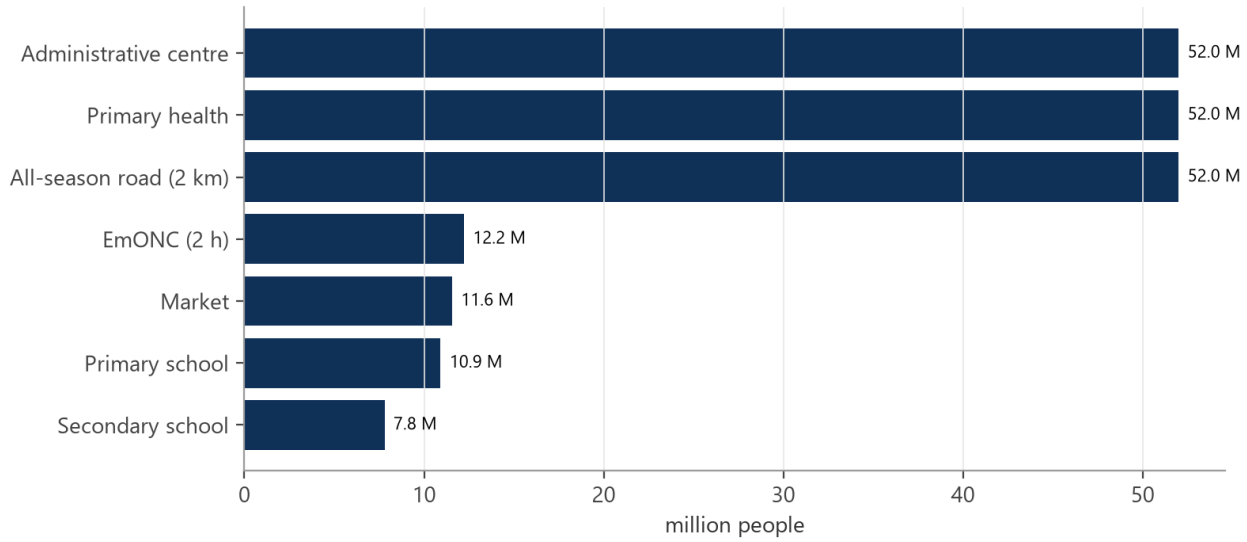

Figure 5.1: People who must reach each destination class (base year)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

These are the needs, and they are large because they are stated for everyone. They say nothing yet about how many people already meet them. The next chapters do: Chapter 7 finds, for example, that with unknown surface counted as earth only 35.9 % of the requiring population meets the all-season road standard in the wet season, while 91.5 % of women aged 15–49 meet the emergency obstetric care standard.

The requirement grows with population. With the projection carried forward at 3 % a year, each class grows by 12.6 % between 2026 and 2030: the planning population rises from 52.0 to 58.5 million.

5.3.2 Where the requirement lies

Figure 5.2 maps the requirement of each class by district, each on its own scale. The six maps look almost identical, and that is a finding about the data, not about Uganda. Because the group shares are national, every group requirement is the district population times a constant. The ranking of districts is the same for every class; only the scale changes. Women of reproductive age, school-age children and households are distributed exactly as the whole population is. When UBOS age structures by district replace the national shares, the maps will differ: districts with young populations, such as many in the north and in Karamoja, will rise in the school maps, and urban districts with more working-age women will rise in the obstetric care map.

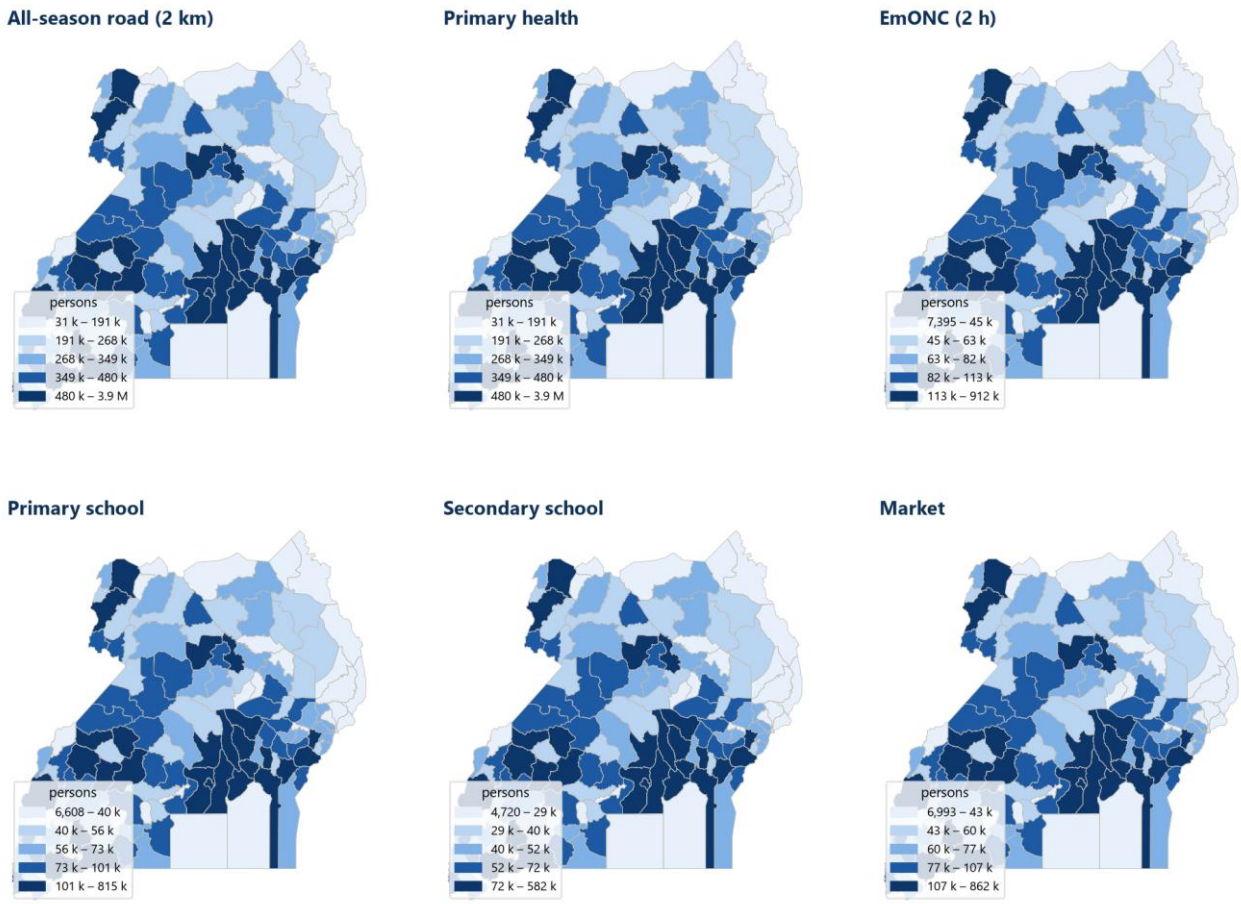

Figure 5.2: People requiring each destination class, by district (each map on its own scale)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The district pattern follows population. By district, the requirement for everyone (all-season road, primary health, administrative centre) ranges from **31,468** persons in Obongi to **3.88 million** in Wakiso. The median district requires 317,004. After Wakiso, the largest requirements are in Kampala (2.05 million), Kasese (975,082), Arua (924,064), Mukono (870,910) and Yumbe (867,130). The smallest, after Obongi, are in Karenga (83,412), Kalangala (84,690), Ntoroko (93,236), Moyo (103,754) and Nabilatuk (115,962). By region, the Central region holds 28.0 % of the requirement (14.5 million), the Eastern 26.0 % (13.5 million), the Western 25.4 % (13.2 million) and the Northern 20.6 % (10.7 million).

5.3.3 The all-season road requirement

Figure 5.3 shows the requirement for an all-season road within 2 km with its LISA clusters. The darkest districts (480,000 persons or more, 27 districts) form four groups: the Kampala–Wakiso–Mukono core with Luwero, Buikwe and Kayunga; a band across Busoga to Mbale and Tororo (Jinja, Mayuge, Kamuli, Iganga, Bugiri, Buyende); Lira and Oyam in Lango, and Arua and Yumbe in West Nile; and a western group from Mubende, Kakumiro, Kagadi, Kyenjojo and Kyegegwa to Kasese, Mbarara, Isingiro and Ntungamo. The lightest districts (under 191,000) are in Karamoja, along the northern border, on the islands and around Lake Albert.

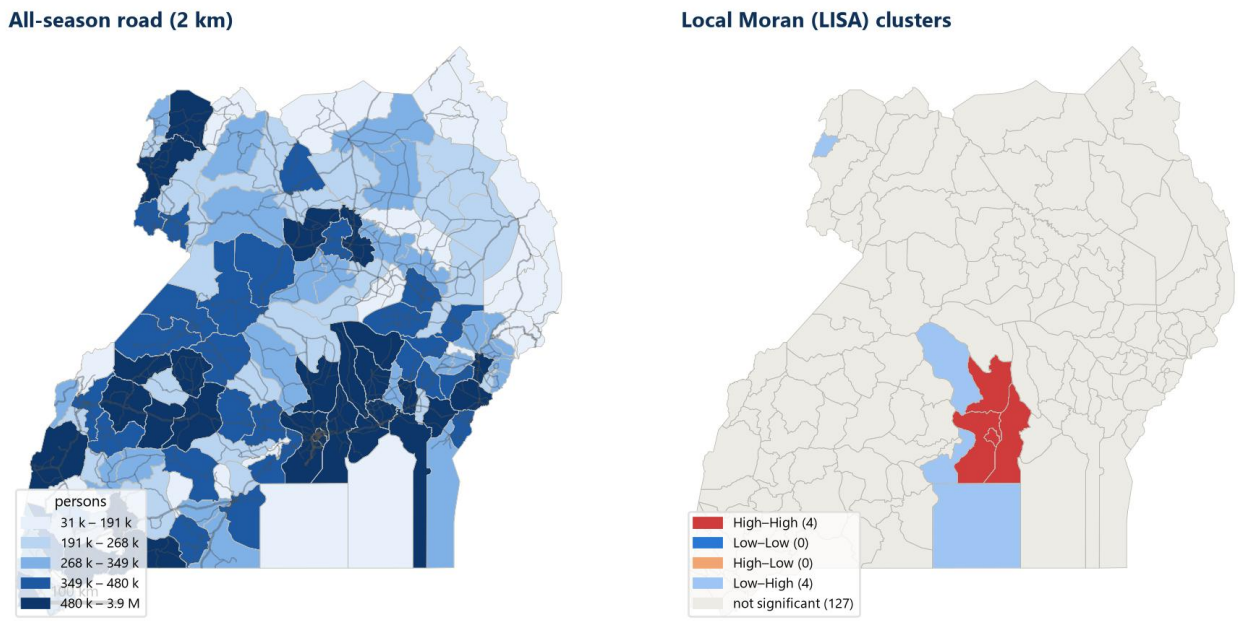

Figure 5.3: People requiring access to: All-season road (2 km), by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The requirement is spatially clustered. Moran's I is 0.275 (p = 0.001), higher than for density, because large districts with large populations sit together around the capital. The LISA panel identifies a **High–High** cluster of four districts (Kampala, Wakiso, Mukono and Luwero) and four **Low–High** districts beside it (Kalangala, Mpigi and Nakaseke, which have smaller populations than their neighbours, and Maracha in West Nile). No Low–Low cluster is significant: the sparsely populated districts of the north-east are small enough in number and different enough from one another not to form one.

Because the all-season road requirement applies to everyone, its district values are the district populations. The question it frames is not how many people require an all-season road (everyone does), but how many live more than 2 km from a road open in the wet season. That is the Rural Access Index, which Chapter 7 computes and Chapter 8 turns into the G1 gap.

5.3.4 Emergency obstetric care

Figure 5.4 shows the requirement for emergency obstetric and newborn care: women aged 15–49 who must reach a health centre IV, hospital or regional referral hospital within 2 hours by motorised transport (1 hour in towns). The district range is from 7,395 women in Obongi to 911,512 in Wakiso; the median district holds 74,496. The pattern and the clusters are those of the all-season road map, scaled by 0.235. With 319 eligible facilities, each serves on average about 38,000 women of reproductive age, but the average hides the difference between Wakiso, where many facilities are close together, and Karamoja, where few facilities serve large areas.

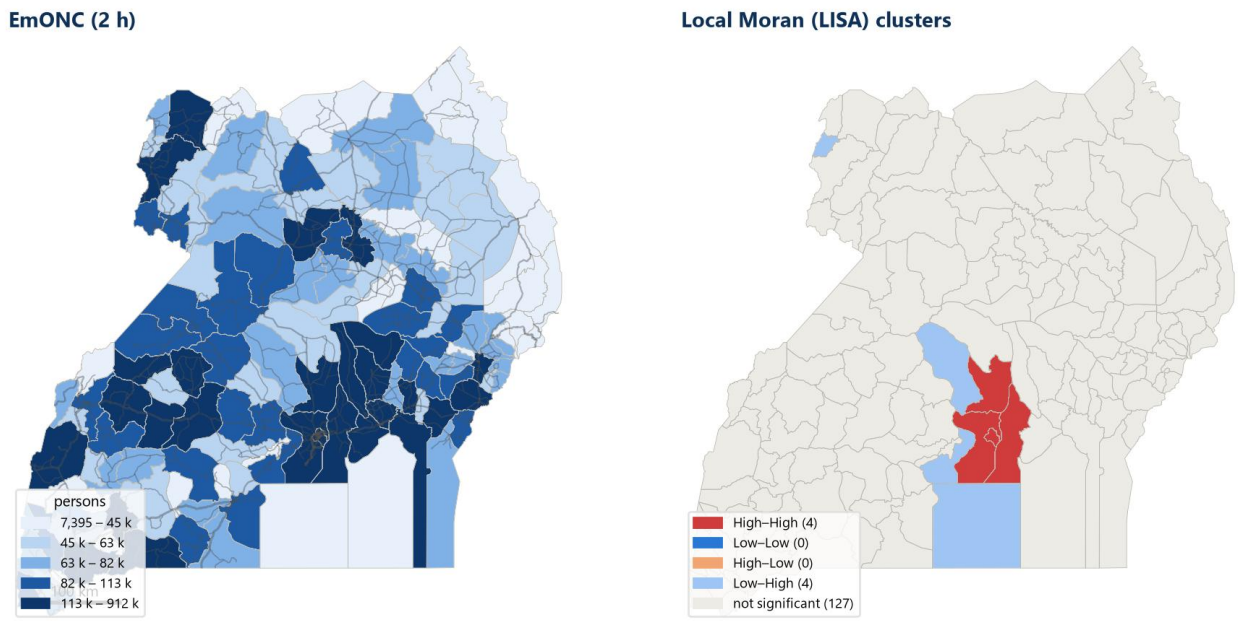

Figure 5.4: People requiring access to: EmONC (2 h), by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

This is the class where the destination set matters most. The engine treats facility level (health centre IV and above) as a proxy for emergency obstetric capability, because the facility register does not record which facilities actually provide it. If some health centres IV do not provide caesarean sections or blood transfusion, the true destination set is smaller and the access gap larger. The fix belongs in HealthModellingx's facility register, not in the transport engine.

5.3.5 Secondary school

Figure 5.5 shows the requirement for secondary school: young people aged 13–18 who must reach one of 3,594 secondary schools within 60 minutes by mixed modes (30 minutes in towns). The district range is 4,720 in Obongi to 581,816 in Wakiso, with a median of 47,551. The pattern again mirrors population. The school layer gives an enrolment for every secondary school: 647,524 pupils in all. That is less than a tenth of the 7.8 million young people the standard covers, which says something about enrolment, but it is a statement for the education engine to make. TIPRAS uses the school only as a destination.

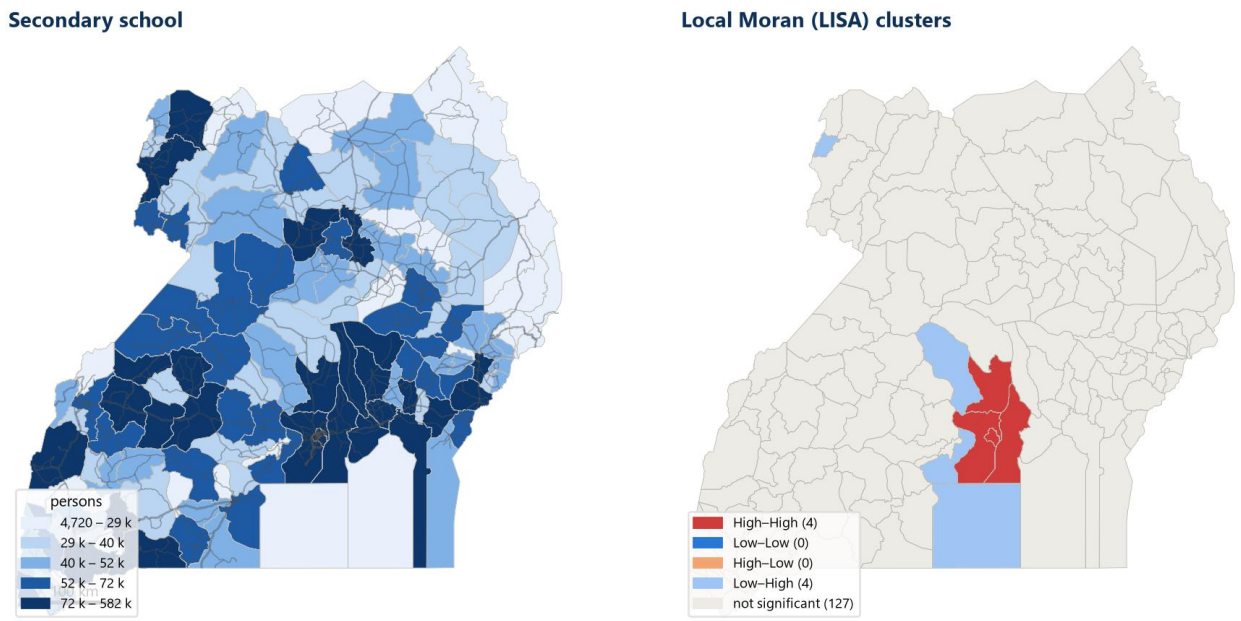

Figure 5.5: People requiring access to: Secondary school, by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

5.3.6 Markets

Figure 5.6 shows the market requirement: 11.6 million households that must reach one of 814 markets within 2 hours by motorised transport (30 minutes in towns). The district range is 6,993 households in Obongi to 861,950 in Wakiso, with a median of 70,445. The requirement is spread as population is; the markets are not. Chapter 4 showed that the market nodes cluster in West Nile, around Kampala and in the far south-west. Where the market layer is thin, the travel time to the nearest market will be long, and part of the resulting gap will reflect the data rather than the absence of markets.

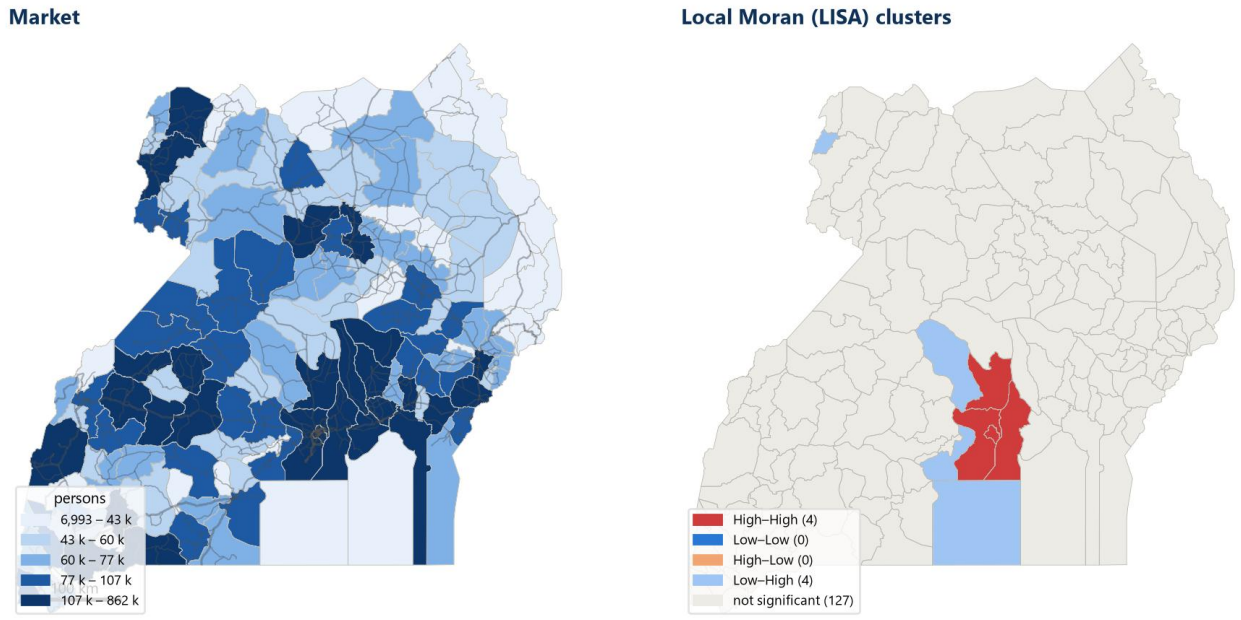

Figure 5.6: People requiring access to: Market, by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The market requirement is the farm-to-market standard of the design. It is distinct from the freight chain below: the access requirement asks whether a household can reach a market, while the freight chain asks which roads carry the marketed surplus of the production zones.

5.4 Results: node connection standards

Of the 24,952 sector nodes, **178** carry a connection standard in the reference run:

- **117 health nodes**: the 102 hospitals require at least a district-class gravel connection, and the 15 regional referral hospitals a district-class road with double surface treatment;
- **28 gateways**: the 22 border posts and 6 airports require a national trunk connection with double surface treatment and an 11.5 t axle;
- **23 industrial zones**, each requiring a national primary connection with double surface treatment and an 11.5 t axle to a gateway;
- **10 agricultural zones**, each with a surplus above 5,000 t and requiring at least a gravel district-class connection to its nearest market; zones above 50,000 t require a sealed national secondary connection.

None of the 309 mines carries a standard. The mine rows apply from 100,000 t a year of output, and no mine has an output in the node register: the minerals exchange in the estate is synthetic demonstration data and is not used. Mines are therefore present as destinations a connection could serve but not as generators of a requirement. Chapter 8 tests each of the 178 nodes against its standard (the G6 gap) and finds, for example, that 103 of the 117 health nodes and all 10 agricultural zones lack the connection their standard requires.

5.5 Results: movement generation

5.5.1 Freight

TI-02 generated **9.76 million t** of agricultural freight a year in the base year: the marketed surplus of the 10 agricultural zones, from MapSPAM 2010 production and the assumed marketed shares. It generated **18.2 million t** of consumption demand, from 52.0 million persons at an assumed 0.35 t per person per year. No mineral or industrial freight was generated, because no mine or industrial zone carries an output. No gateway flows were generated, because the trade exchange is also synthetic and is not used.

The balancing factor between agricultural production and consumption is therefore 9.76 / 18.2 = **0.54**, outside the plausibility band of 0.8 to 1.25. The engine scales consumption down to match production, as the method requires, but the gap is a warning. The production side counts only the marketed surplus of ten very large zones from 2010 data, while the consumption side counts everyone's food consumption in 2026. Production from cells outside the zones, imports, and non-agricultural freight are all missing. The freight that TI-03 loads onto the network is therefore a partial picture of national freight, dominated by the agricultural surplus of a few large zones.

5.5.2 Passenger trips

Passenger trips come from assumed motorised trip rates: 0.15 trips per person per day in rural cells and 0.6 in urban cells, with 10 % of trips leaving the district. With 80 % of the planning population rural, this gives about 0.24 motorised trips per person per day, or about 12.5 million trips a day and 4.6 billion a year. Rural and urban cells generate about the same number of trips (some 6.2 million a day each), because the urban rate is four times the rural one and the urban population a quarter of the rural.

The design's plausibility band of 0.5 to 3.0 trips per person per day applies to all trips, including walking. The engine generates only motorised trips, so its rate of 0.24 is not comparable with the band. A household travel survey, or the travel module of the Uganda National Panel Survey, would replace both rates.

5.5.3 Growth

No sector engine supplied an output trajectory and no income path was available. The engine therefore grows passenger trips with population only (12.6 % from 2026 to 2030), applies no income elasticity, and grows freight at the flagged fallback rate of 5 % a year (21.6 % over four years). The run manifest records the fallback, and every traffic result in Chapter 6 inherits it.

5.6 Data quality and limitations

The needs are as good as their two ingredients, population and norms, and the norms are mostly assumed.

- **Standards.** Two of fourteen access standard rows are normed; the rest are illustrative. The primary health and primary school standards are labelled as sector standards but have not been imported from the health and education engines. Until they are, the transport engine's health and education needs may disagree with those sectors' own.
- **Groups.** The group shares are national and assumed. Every group requirement is therefore a scaled copy of the population map, and districts with unusual age structures are misstated.
- **Destinations.** Emergency obstetric capability is proxied by facility level; market coverage is uneven; no node has a commissioning date.
- **Generation.** Agricultural freight comes from 2010 production and ten zones; mineral, industrial and gateway freight is absent; trip rates are assumed; growth is a fallback.
- **Quality inherited.** By the worst-quality rule of Chapter 3, every requirement combines imputed population with an assumed share or standard, and is therefore assumed except for the rural all-season road requirement of everyone, whose standard is normed and whose share is 1.

None of this changes the size of the everyone-requirements: 52.0 million people need an all-season road, a health facility and their district headquarters within reach, whatever the norms say about the threshold. What the norms change is who meets the standard, and that is the subject of Chapter 7.

5.7 Implications for planning

- The access requirement is universal for three classes: all 52.0 million planning persons need an all-season road within reach, a primary health facility and their district headquarters. The planning question is not whether a district needs access but how far its people are from meeting the standard.
- The health and education ministries should own their access standards. Importing the primary health and primary school standards from HealthModellingx and EducationModellingx, and signing them off, would make transport needs move with sector policy instead of with the transport engine's placeholders.
- District age structures from UBOS are needed before the group requirements (12.2 million women of reproductive age, 10.9 million primary-age children, 7.8 million secondary-age youth) can differ between districts in any way other than by population.
- Emergency obstetric care deserves a destination set based on capability rather than facility level. The transport gap for the 12.2 million women it covers depends directly on which facilities qualify.
- Freight needs are under-stated. Mines, industrial zones and gateways need outputs and flows from their sector engines before the engine can say what standard their roads require. Until then, the freight-driven standards of Chapter 6 rest on agricultural surplus alone.
- Growth should come from drivers, not from a rate. Sector output trajectories and an income path from the fiscal models would replace the 5 % fallback that every traffic figure now carries.

CHAPTER 6

6 TI-03 Demand and the required standard of every link

This chapter completes the needs side of the chain. It asks *which links carry the movement generated in TI-02, and what standard each link therefore requires*. The module distributes freight and passenger movement between origins and destinations, loads it onto the network, converts the loads into vehicles and axle loads, and then sets a required standard for every link from four tests: traffic, functional role, the connection needs of sector nodes, and the access paths of settlements. The required standard is transport's equivalent of a bill of materials, and the test that sets it, the binding test, tells a planner *why* the standard is required. The chapter implements design Chapter 5 (`docs/05_TI03_DEMAND.md`) and equations (3.1) to (3.10) of Chapter 23.

6.1 What the module does

TI-03 takes the generation tables of TI-02 (9.76 million t of agricultural surplus from ten zones, 18.2 million t of consumption demand and about 12.5 million motorised trips a day), the node connection standards and the access standards, and the planning graph built by TI-04. It produces, for every link:

- the **required-state load**: annual average daily traffic (AADT), passenger car units, freight tonnes and equivalent standard axles (ESA) per year, in the base year and the horizon year;
- the **required standard**: surface, width, lanes, all-season status, design axle load, design speed and structural capacity;
- the **binding test**: which of the four tests set that standard.

It also writes the district total of road km whose required surface is above its current surface, the input to the standard gap G4 in Chapter 8.

6.1.1 Why demand is loaded on a required-state network

The design makes one choice that shapes every result in this chapter. Future demand is loaded on a **required-state** network, not on today's network. Loading demand on today's network produces a self-fulfilling plan: traffic avoids the impassable road, so the impassable road shows no demand, so it is never upgraded. In the required-state graph every link is brought up to at least the minimum standard of its functional role, in good condition, with the speed that standard allows. A district road of unknown or earth surface is treated as gravel, a community road as engineered earth, a national primary road as double surface treatment. The loads that result answer the question "what would this link carry if it did its job". They are labelled modelled, state required, everywhere they appear.

6.2 Method

6.2.1 Distribution: who travels to where

Where a counterpart engine supplies an origin–destination matrix for a commodity or purpose, the engine takes it as given. No such matrix was supplied for Uganda, so every flow is distributed by a **doubly constrained gravity model**, solved by Furness balancing:

$$T_{od} = A_o B_d O_o D_d \exp(-\beta c_{od}) \tag{6.1}$$

$$A_o = \left[\sum_d B_d D_d \exp(-\beta c_{od})\right]^{-1}, \quad B_d = \left[\sum_o A_o O_o \exp(-\beta c_{od})\right]^{-1} \tag{6.2}$$

where

- T is the flow from origin *o* to destination *d* (tonnes or trips a year);
- O and D are the totals generated at the origin and attracted to the destination;
- *c* is the cost between them on the required-state graph;
- β is the distance-decay parameter, a norm;
- A and B are the balancing factors, updated in turn until every row and column sum matches its total to within 0.1 %.

The design defines *c* as a generalised cost (time valued at the value of time, plus vehicle operating cost, plus gateway penalties). The reference run uses travel time in hours on the required-state graph as the cost, with β = 0.2 per hour, a value taken from TransportModelling and marked uncalibrated. Where observed trip lengths exist, β would be calibrated by Hyman's method; none exist for Uganda.

6.2.2 The two-stage agricultural chain

Agricultural surplus does not go straight to the capital. It goes first to a market and is then re-shipped to consumers. The design splits each zone's surplus across the markets within its standard by a logit share:

$$\text{share}_{z,\,m} = \frac{\exp(-\theta\, c_{zm})}{\sum_{m' \in M_z} \exp(-\theta\, c_{zm'})} \tag{6.3}$$

where θ is a norm and M is the set of markets within reach of zone *z*. The reference run uses the limiting case in which θ is very large, so each zone sends all its surplus to its nearest market by travel time. In stage 2 the market throughput becomes an origin of the gravity model above, with the 135 district centres as destinations. This two-stage chain is what makes feeder roads visible to the plan: the first leg runs from the production zone along district and community roads to a market.

6.2.3 Passengers

Passenger trips are loaded in two parts. The 90 % of motorised trips that stay within their district are loaded from each settlement cell to its district centre. The 10 % that leave the district are distributed among the 135 district centres by the same gravity model.

6.2.4 Assignment and the conversion to vehicles and axles

Each flow is assigned to its shortest path on the required-state graph, all or nothing. The design adds a capacity-restrained incremental assignment (four increments of 40, 30, 20 and 10 %) on links where volume exceeds 60 % of capacity. **That step is not yet built.** Loads on the busiest corridors are therefore those of an unconstrained shortest-path loading, and are overstated where real traffic would spread over alternative routes.

Loads are converted to vehicles and axle loads with the vehicle norms:

$$\text{AADT}_l = \tfrac{1}{365}\left[x_l^{\,f} \sum_h \frac{\eta_h (1+e)}{\text{pay}_h\, lf} + \frac{x_l^{\,p}}{\text{occ}}\right] \tag{6.4}$$

$$\text{ESA}_l = x_l^{\,f} \sum_h \frac{\eta_h\, EF_h}{\text{pay}_h\, lf} + \frac{x_l^{\,p}}{\text{occ}} \sum_{h \in H^p} w_h EF_h \tag{6.5}$$

where

- x with superscript *f* is the freight on link *l* in tonnes a year, and with superscript *p* the passenger trips;
- η is the share of freight carried by truck class *h* (heavy 0.50, medium 0.35, light 0.15), pay its payload (20, 8 and 3 t), *lf* the load factor (0.7) and *e* the empty-return factor (0.6);
- occ is the average vehicle occupancy (2.2) and *w* the share of each passenger vehicle class in the passenger mix;
- EF is the equivalent standard axle factor of a vehicle (3.5 for a heavy truck, 1.6 medium, 0.15 light, 1.2 for a bus, 0.02 for a minibus, 0.0001 for a car, zero for a motorcycle).

With these norms one tonne of freight a year generates 0.271 truck movements and 0.236 ESA (empty trucks are assumed to cause no pavement damage), and one passenger vehicle 0.053 ESA. The design computes the ESA factor axle by axle with the fourth-power law and a separate overloading term; the reference run uses per-vehicle factors from the norm table, with no overloading, because no axle-load survey is available. All vehicle norms are TransportModelling defaults marked "not Ugandan".

Design traffic over a 15-year design life, with growth *g* and a directional split of one half, is:

$$N_{\mathrm{des},\,l} = \mathrm{ESA}_{l,\,t_H} \cdot \frac{(1+g)^n - 1}{g} \cdot f_{\mathrm{dir}} \tag{6.6}$$

With the fallback growth rate of 5 % a year and *n* = 15, the multiplier is 21.6 × 0.5 = 10.8 times the horizon-year ESA.

6.2.5 The required standard

The required standard is the most demanding of four tests, each a lookup in the norm table `design_standard`. A standard is a record (surface, width, lanes, all-season status, design axle, design speed, structural capacity), and the engine takes the field-by-field maximum, the **join** of the four records:

$$S_l^{*} = S_l^{\mathrm{traffic}} \sqcup S_l^{\mathrm{function}} \sqcup S_l^{\mathrm{node}} \sqcup S_l^{\mathrm{access}} \tag{6.7}$$

where

- the **traffic** standard comes from the horizon-year AADT band and design ESA (Table 6.1);
- the **function** standard is the minimum for the link's role: engineered earth for community roads, gravel for district and national secondary roads, double surface treatment for national primary and trunk roads, low-volume seal for urban local roads;
- the **node** standard is the join of the connection standards of every sector node whose path to its target uses the link;
- the **access** standard is all-season if the link lies on a path by which a settlement meets an access standard.

If the design traffic exceeds the structural capacity of the resulting standard, the surface is promoted to the lowest rank whose capacity suffices:

$$N_{\mathrm{des},\,l} > \mathrm{cap}(S_l^{*}) \Rightarrow S_l^{*} \leftarrow S_l^{*} \sqcup \min\left\{ S^{b} : \mathrm{cap}(S^{b}) \geq N_{\mathrm{des},\,l} \right\} \tag{6.8}$$

The **binding test** is the test whose surface rank is highest, with ties broken in the order node, access, traffic, function. The lowest traffic band only restates the functional minimum, so the traffic test binds only from the second band upwards (50 vehicles a day). Otherwise every near-empty rural road would be reported as traffic-driven. Ferries carry no road standard; their binding test is recorded as "mode".

Table 6.1: Traffic bands of the design standard (assumed norms, to be replaced by the national road design manual)

Horizon-year AADT	Required surface	Width	Lanes	Design axle	Structural capacity (ESA)
below 50	engineered earth	5.0 m	1	8.2 t	50,000

Horizon-year AADT	Required surface	Width	Lanes	Design axle	Structural capacity (ESA)
50–200	gravel	6.0 m	2	8.2 t	100,000
200–400	gravel	6.5 m	2	8.2 t	300,000
400–1,000	double surface treatment	7.0 m	2	10 t	1,000,000
1,000–8,000	asphalt concrete	7.0 m	2	11.5 t	10,000,000
above 8,000	asphalt concrete, dual carriageway	14.0 m	4	11.5 t	30,000,000

6.3 Results: the distribution of movement

6.3.1 Agricultural freight

In stage 1 the 9.76 million t of marketed surplus of the ten agricultural zones reaches **nine markets**: each zone sends its whole surplus to its nearest market, and two zones share a market. In stage 2 the gravity model distributes the markets' throughput to the 135 district centres. The Furness iteration stopped at its limit of 500 iterations with **9.69 million t** distributed, 0.7 % short of the 9.76 million t to be distributed. The design's validation rule requires balance to 0.1 %. The shortfall is small, but it is a validation failure, and it suggests that some district centres are poorly reachable from the markets on the required-state graph within the decay set by β.

The freight loaded on the network adds up to **2.72 billion tonne-km** a year. It travels on 22,371 links totalling 9,709 km: 4,054 km of national secondary roads, 2,990 km of national primary roads, 1,501 km of district roads, 704 km of trunk roads and 441 km of community roads.

Figure 6.1 shows the result. Freight follows the radial pattern of the national network into Kampala, with the heaviest flows (over a million tonnes a year) on the trunk roads from the west and south-west, the Northern Corridor east of Kampala, the Gulu road and the routes into West Nile. The darkest segment, north of Kampala in Nakaseke, carries **15.2 million t** a year on a district road, more than the whole surplus, because it carries both the stage-1 flows of a large zone into its market and the stage-2 flows out of that market. Freight is concentrated: the 10 % of freight-carrying km with the most tonne-km carry 41.7 % of all tonne-km. The median freight link carries 80,826 t a year, and one in ten carries more than 929,000 t.

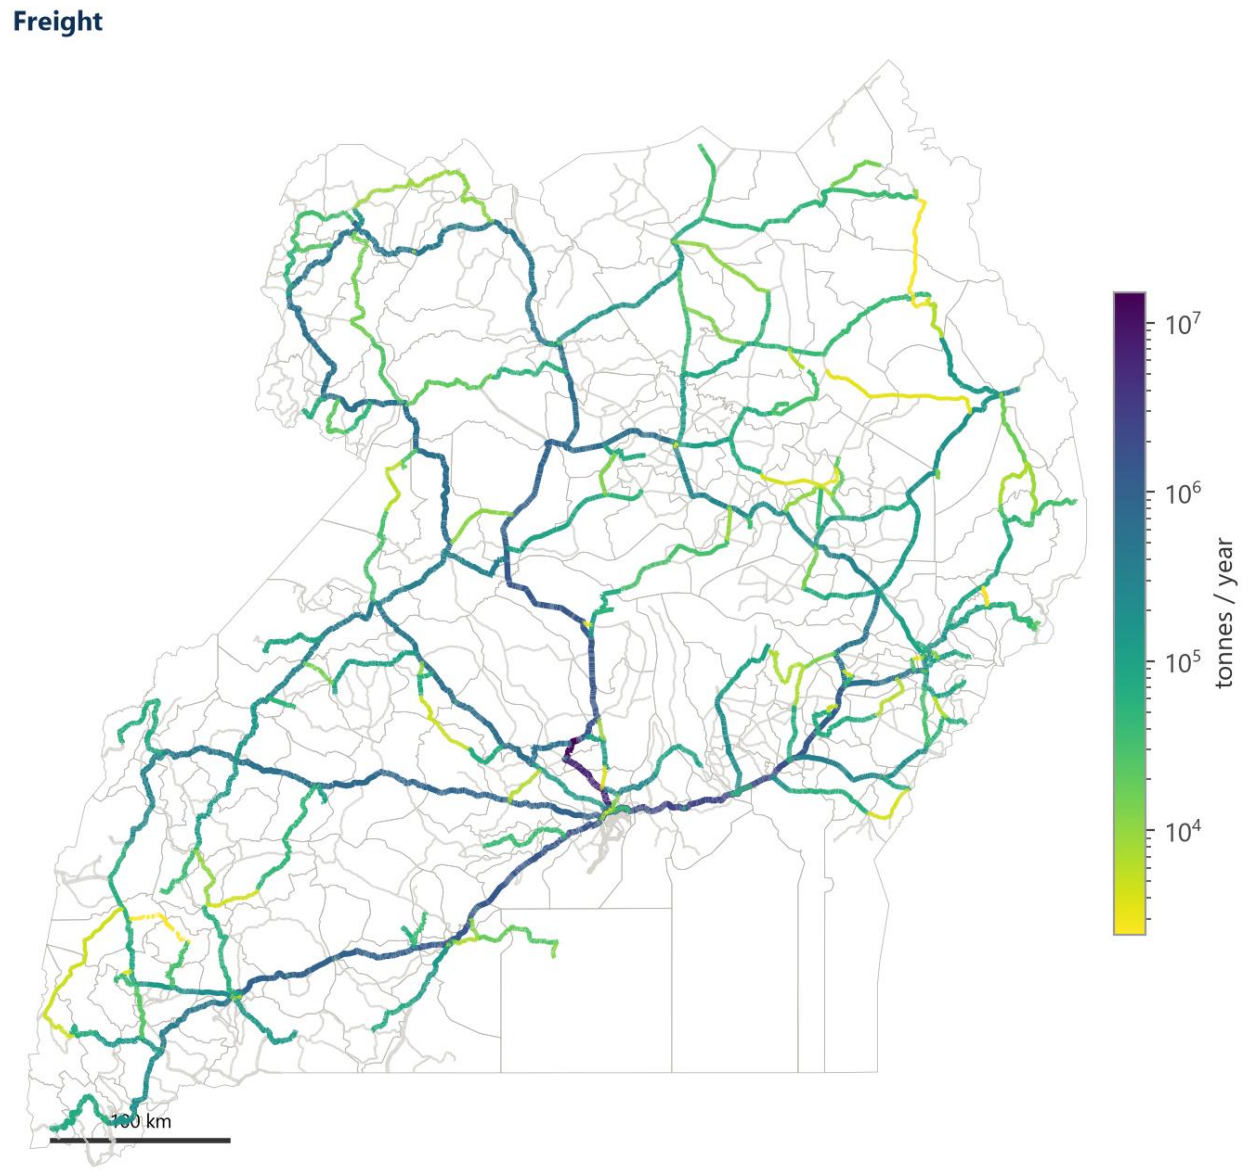

Figure 6.1: Freight flows on the network: agricultural surplus to markets

Source: TI-03 two-stage agriculture chain on MapSPAM 2010 surplus

The pattern is as much a product of the method as of Uganda's agriculture. With only ten zones (Chapter 4), each a large region collapsed to a single centroid, and each sending all its surplus to one market, the first leg of the chain concentrates on a handful of roads, and most feeder roads receive no agricultural freight at all. A finer zoning and a logit split across markets would spread the first leg over many more district and community roads.

6.3.2 Required-state traffic

The engine loaded traffic on **288,725 of the 817,382 links** (35 %), totalling about 139,000 km. Tracks carry no load, because their functional standard (none) gives them no speed in the required state. Of the other classes, more than 84 % of the links of every national and district class are loaded, but only 38.8 % of community road links: most community roads lie off the shortest paths between settlement cells and district centres.

Figure 6.2 shows the distribution of base-year AADT over the loaded links on a logarithmic scale, with the band thresholds marked. It has two humps. The main hump, centred near 30 vehicles a day, is the mass of local roads that carry the trips of a few settlement cells. The second, flatter hump, from about 1,000 to 100,000 vehicles a day, is the national and district network into which those trips converge on their way to the district centres. By band:

- 135,447 links (62,685 km) carry fewer than 50 vehicles a day;
- 64,499 links (37,949 km) carry 50–200;
- 17,495 links (10,741 km) carry 200–400;
- 20,371 links (9,625 km) carry 400–1,000;
- 32,370 links (12,264 km) carry 1,000–8,000;
- 18,543 links (5,735 km) carry more than 8,000.

The median loaded link carries 58 vehicles a day; the 75th percentile 383; the 90th percentile 3,369; the 99th percentile 65,258.

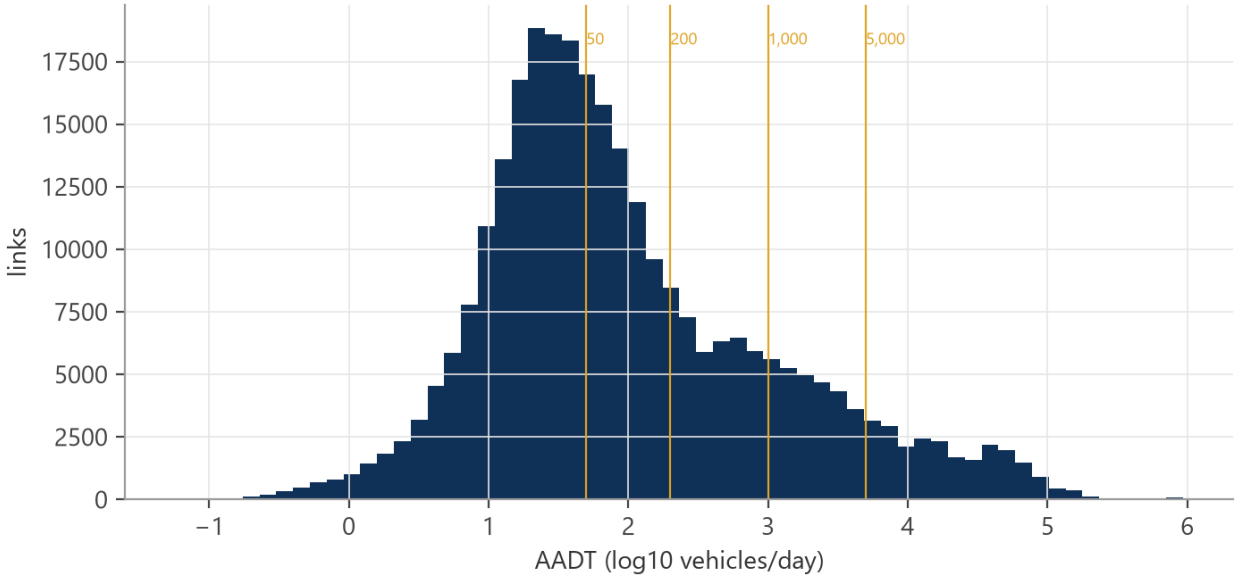

Figure 6.2: Distribution of required-state traffic over loaded links

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Figure 6.3 maps the 20,000 busiest links. Traffic converges on Kampala along every radial road, and is heaviest on the Northern Corridor towards Jinja and the Kenyan border, on the Masaka road to the south-west and on the roads into Kampala from the north and west. Regional networks around Mbale and Tororo, Gulu and Lira, Arua, Fort Portal and Kasese, and Mbarara and Kabale stand out as secondary concentrations. By class, the median loaded trunk link carries 41,902 vehicles a day, primary 28,595, secondary 2,685, district 380 and community 34.

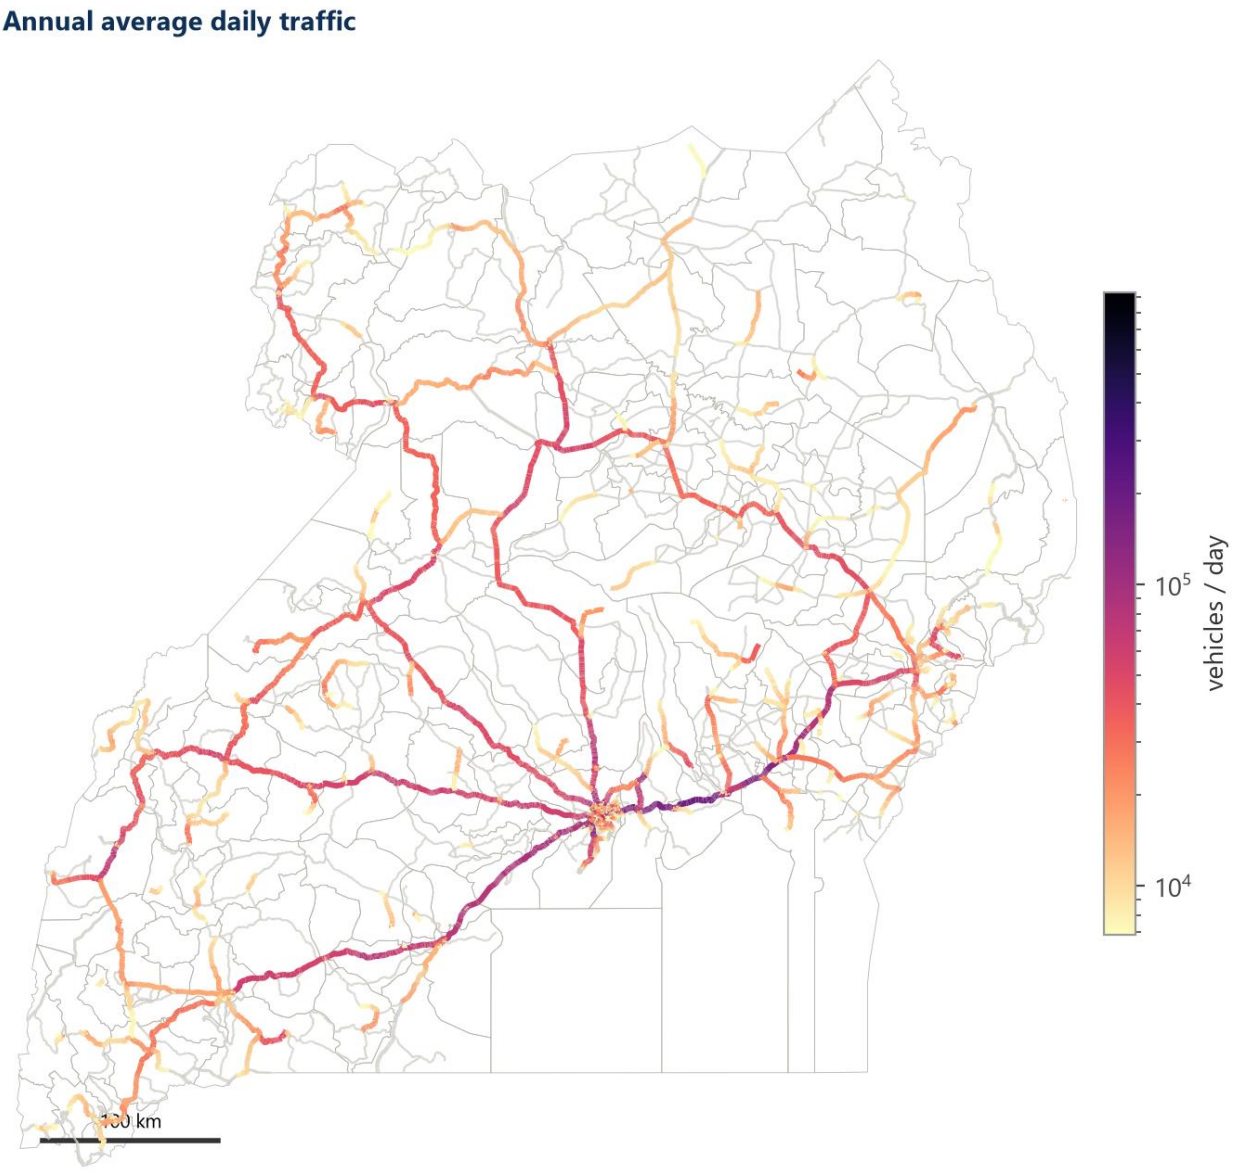

Figure 6.3: Required-state traffic (AADT) on the 20,000 busiest links, from the all-or-nothing assignment

Source: TI-03 Furness gravity + AON assignment; TransportModelling vehicle mix (uncalibrated)

The top of this distribution is not credible as a traffic forecast, and should not be read as one. The highest value, **935,037 vehicles a day**, is on short links in Wakiso where all the intra-district trips of 3.9 million planning persons converge on the single point chosen as the district centre. It arises from three simplifications that compound one another:

- every intra-district motorised trip is loaded to the district centre, as if every trip in a district had the same destination;
- the loading is all or nothing, with no capacity restraint to spread traffic over parallel roads;
- the trip rates (0.15 a day rural, 0.6 urban) and the occupancy (2.2) are assumed.

The pattern of the map is informative: it shows which roads lie on the shortest paths that connect people to their district centres and markets. The levels on the busiest corridors are not. No traffic count is available to calibrate them (the design's calibration statistic, the share of count stations with a GEH below 10, cannot be computed), and none should be inferred from them.

6.3.3 Pavement loading

Figure 6.4 shows ESA per year on the most loaded links. Its pattern follows freight more than traffic, because trucks cause almost all pavement damage. The heaviest loading, above a million ESA a year, is on the trunk roads into Kampala, the Northern Corridor and the Masaka road, and on the Nakaseke freight link north of Kampala. Over the loaded links, the median is 1,121 ESA a year and the 99th percentile 1.42 million. **5,425 links (1,399 km) carry more than a million ESA a year**, and 7,842 km more than 100,000. The maximum is 18.2 million ESA a year.

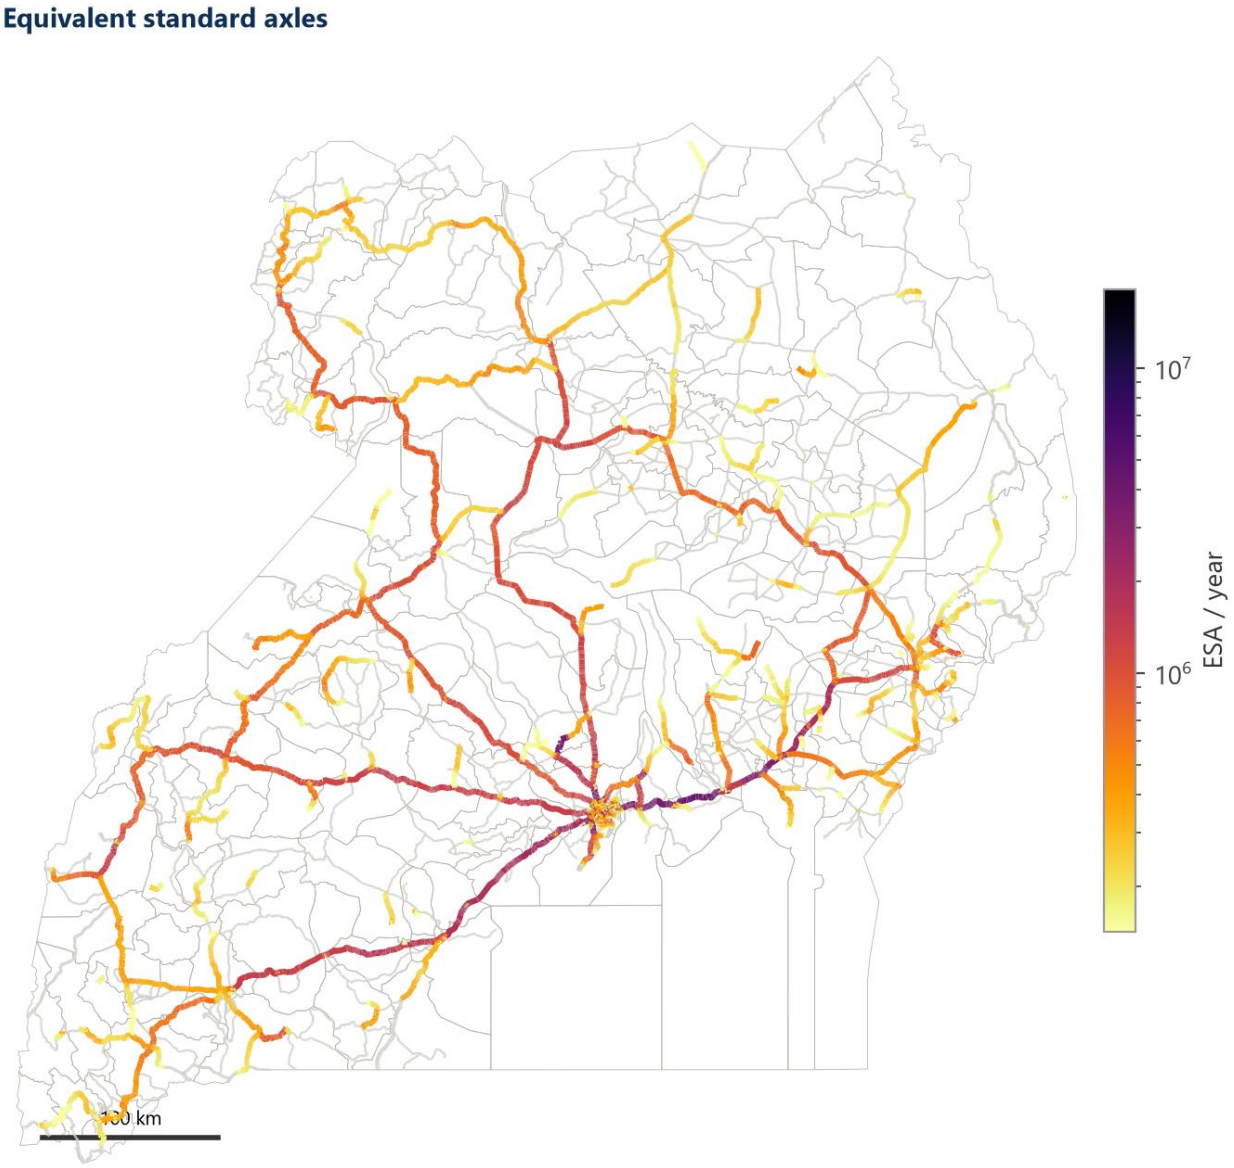

Figure 6.4: Pavement loading: equivalent standard axles per year on the most loaded links

Source: TI-03 ESA from the vehicle mix

At the 5 % fallback growth and a 15-year life, a link with a million ESA a year in 2030 needs a design capacity of about 10.8 million ESA, the capacity of the asphalt concrete standard. The structural promotion rule therefore matters on the heavy corridors: it can lift a link whose traffic alone requires double surface treatment to asphalt concrete.

6.4 Results: the required standard

6.4.1 Which test binds

Table 6.2 and Figure 6.5 show which test sets the required standard. The functional role binds on **656,479 links (80.3 %)**. Traffic binds on **160,189 links (19.6 %)**. The connection standard of a sector node binds on **686 links (0.1 %)**, and the 28 ferry links carry no road standard ("mode").

Table 6.2: Which test sets the required standard on each link (TI-03)

Binding test	Links	Share
function	656,479	80.3%
traffic	160,189	19.6%
node	686	0.1%
mode	28	0.0%

The lowest traffic band never binds by construction. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_binding.csv).

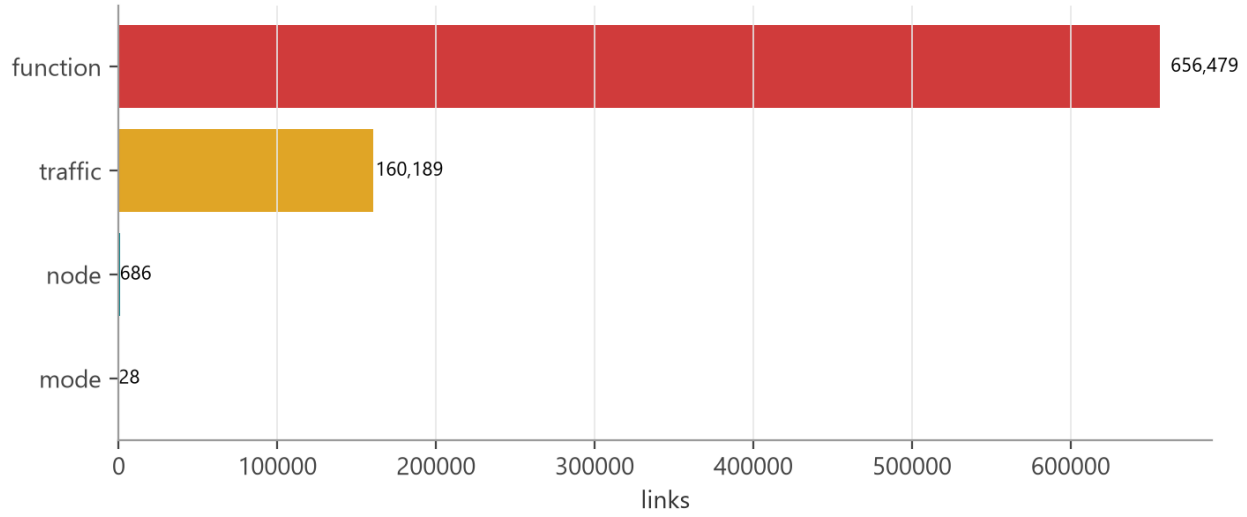

Figure 6.5: Which test sets the required standard on each link

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Three observations follow.

- **Function dominates because most links are quiet.** For most of the network the required standard is simply the minimum for its class: engineered earth for a community road, gravel for a district road. That is the expected result in a network where most links carry fewer than 50 vehicles a day.
- **Traffic binds on one link in five**, the links whose horizon-year AADT reaches at least 50 vehicles a day and whose traffic band requires more than their class. These include many district and community roads that lie on the paths into district centres. The level of traffic on them is uncalibrated, so the traffic-bound share should be read as an upper bound on where traffic drives the standard.
- **Nodes and access rarely bind.** The 686 node-bound links lie on the paths of the 178 nodes with connection standards, mostly the hospitals, gateways and industrial zones whose standards require double surface treatment. Access never binds in the reference run. The access test requires only all-season status, and every functional class that settlements use already requires an all-season standard, so access ties with function and the tie goes to function. In this run the access test confirms the functional standards rather than adding to them.

6.4.2 How far the network is below its required standard

The required surface is above the current surface on **242,068 km**, 97.6 % of the road network. The number is large, and it is important to understand why.

- **The unknown and the earth.** Every road link requires at least engineered earth, the lowest band of the traffic test, which applies even to near-empty links and to tracks. Earth surfaces, and unknown surfaces treated as earth, rank below engineered earth. The 238,731 km of road with an earth or unknown surface therefore all count as below standard. They account for 98.6 % of the total.
- **The rest.** The remaining 3,337 km or so are links with a known gravel or better surface whose required standard is higher: gravel roads whose traffic requires a seal, sealed roads whose traffic or node standard requires asphalt concrete.

The standard gap is therefore dominated by what is not known. On 169,343 km the engine does not know whether the road is earth, engineered earth or gravel, and its answer, in TI-07, is to list those km as a survey requirement rather than as works. The 69,388 km recorded as earth are a firmer figure: they are below the engineered-earth or gravel standard of their class by observation.

Figure 6.6 maps the km below required standard by district, with the LISA clusters. The largest totals are in Wakiso (9,055 km), Mukono (5,463 km), Kyenjojo (4,787 km), Mubende (4,694 km), Nakaseke (4,503 km) and Yumbe (4,368 km): large districts with dense networks of earth and unknown roads. The smallest are in Nabilatuk (207 km), Nakapiripirit (338 km), Kapchorwa (356 km), Bududa (382 km), Kaberamaido (435 km) and Butebo (438 km): small districts or districts with sparse networks. By region the totals are 69,932 km in the Central, 65,590 km in the Western, 64,140 km in the Northern and 42,406 km in the Eastern region. The median district has 1,449 km below standard.

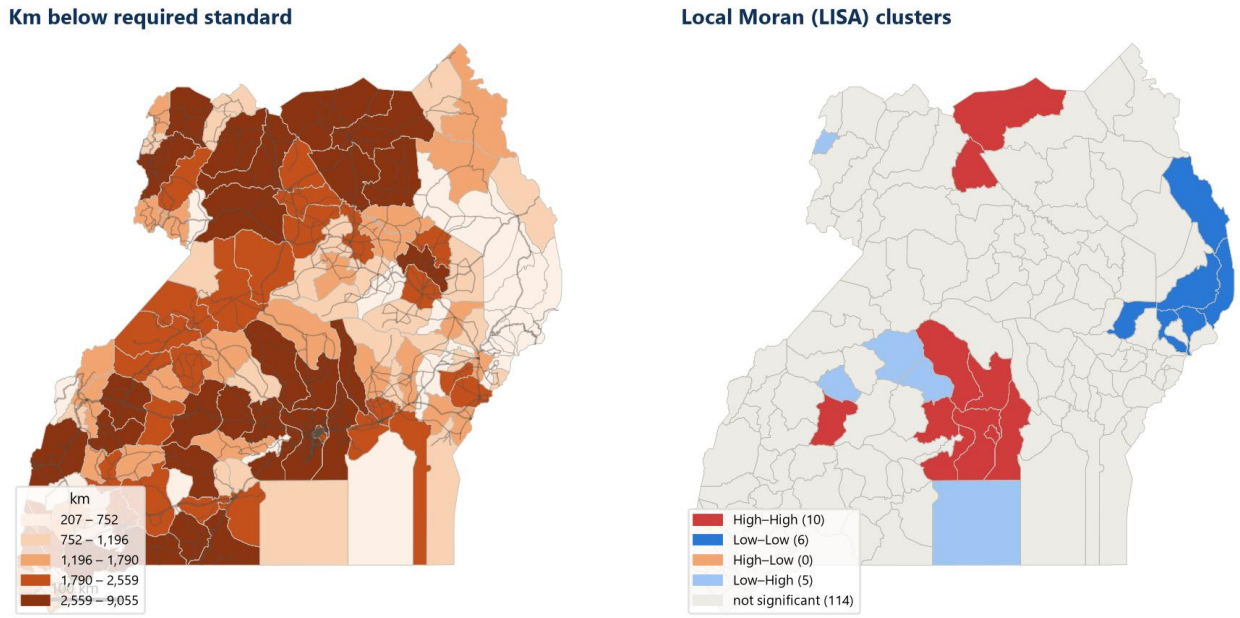

Figure 6.6: Km whose required standard is above the current surface, by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Two normalisations show the same result from different angles. As a share of the district's road km, the median district has 98.1 % of its network below the required standard. The share is lowest in Kampala (85.9 %), Kapchorwa (89.0 %) and Mbale (90.2 %), where more roads are paved or surface is better known, and reaches 100 % in Buvuma, Kalangala, Buhweju, Kapelebyong and Rukiga. Per 1,000 planning persons, the median district has 4.3 km below standard. The ratio is highest in sparsely populated districts with long networks: Obongi (27.8 km per 1,000 persons), Lamwo (20.0), Kapelebyong (19.1), Nakaseke (15.4) and Pader (14.4); and lowest in Kampala (1.1) and Bududa (1.2).

The district totals are spatially clustered (Moran's I 0.295, p = 0.001). A **High–High** cluster of ten districts runs from Kampala, Wakiso, Mukono and Mpigi through Luwero, Nakaseke, Mityana and Kyegegwa, with a second group in Gulu and Lamwo in the north. A **Low–Low** cluster of six districts lies along the eastern border in Karamoja and Elgon (Amudat, Moroto, Nakapiripirit, Bulambuli, Kween and Kumi). Five **Low–High** districts (Kalangala, Kibaale, Kiboga, Kyankwanzi and Maracha) have less km below standard than their neighbours.

6.5 Data quality and limitations

TI-03 is the module of the chain with the most assumed inputs, and its results should be read with that in mind.

Table 6.3: Quality of the TI-03 results in the reference run

Element	Status	Effect
Origin–destination matrices	None supplied; gravity on travel time with uncalibrated β = 0.2 per hour	Distribution pattern indicative only
Freight generation	Agricultural surplus of 10 zones only (MapSPAM 2010); no mineral, industrial or gateway freight	Freight under-stated and concentrated
Stage 1 of the agricultural chain	Nearest market only (θ very large)	First leg concentrated on few roads
Furness balance	Stopped at 500 iterations, 9.69 of 9.76 Mt distributed	0.7 % imbalance; validation failure
Passenger loading	All intra-district trips loaded to one district centre; assumed trip rates and occupancy	Extreme loads near district centres
Assignment	All or nothing; capacity-restrained increments not built	Busiest corridors overstated
Vehicle and axle norms	TransportModelling defaults, "not Ugandan"; per-vehicle ESA, no overloading	Pavement loading indicative only
Growth to the horizon	Freight at the 5 % fallback; passengers with population	Recorded as a run fallback
Calibration	No traffic counts; GEH not computable	No check on levels
Node throughput requirement	Not yet computed	Gateway and hub capacity gaps not assessed
Design standards	Assumed bands; to be replaced by the national road design manual	Standards indicative

All TI-03 layers carry quality assumed and basis modelled. They are fit for showing which roads lie on the paths that connect people and producers to services and markets, and for showing why a standard is required. They are not fit for the design of individual roads, which needs counts, axle-load surveys and the design manual.

Five improvements would change the results most. Traffic counts on the national network and a sample of district roads would calibrate β and the trip rates. Destinations for intra-district trips (markets, schools, health facilities and trading centres) in place of a single district centre would remove the artificial convergence. The incremental capacity-restrained assignment of the design would spread corridor loads. Mineral, industrial and trade flows from the sector engines would add the freight that is missing. And an axle-load survey from the weighbridges would replace the per-vehicle ESA factors.

6.6 Implications for planning

- For most of the network the required standard is set by function, not by traffic: 80.3 % of links need only the minimum standard of their class. Planning for these roads is a question of bringing the network to an all-season minimum, and the cost of doing so is the subject of Chapters 9 and 10.

- The standard gap of 242,068 km is overwhelmingly a gap of knowledge. 169,343 km have no recorded surface. A surface survey would convert most of this figure into either confirmed works or confirmed adequacy, and should come before any programme sized on it.
- Traffic-driven standards are concentrated on the radial national roads and on the district roads that feed district centres. Their levels are uncalibrated and the busiest corridors are overstated; the corridor standards (asphalt, dual carriageways) should not be taken from this run without counts.
- Pavement loading is concentrated on a small part of the network: 1,399 km carry more than a million ESA a year. Axle-load control and the structural design of these corridors matter far more for maintenance costs than their length suggests.
- Freight standards rest on agricultural surplus alone. Until the minerals, industry and trade engines supply flows, the engine cannot say what the roads to mines, industrial parks and border posts must carry, and the node-bound standards (686 links) are the only freight-related signal.
- The required-state approach is a deliberate choice: it shows what roads would carry if they did their job, not what they carry today. That is the right basis for a needs-based plan, but it means that observed counts, when they arrive, should be reported beside the modelled loads, never in place of them.

CHAPTER 7

7 The network today: supply and access

This chapter answers the second planning question: *what does Uganda have today?* It describes the road network as the engine sees it in the base year 2026, which links stay open in the wet season, what is known about their condition, and how long people take to reach the services and places that the needs module (Chapter 5) says they must reach. It implements the TI-04 module, design chapter 6 (`docs/06_TI04_SUPPLY.md`), with the notation of `docs/23_MATHEMATICAL_MODELS.md`. All results come from the reference run `014c9cd7` (geography `UGA-2026.09-c`, norms `UGA-2026.09-draft-f`, prices `UCSN+EV-UGA-2026.09b`). The run is **exploratory**: the norms are drafts, and 68.3 % of network kilometres have no known surface, which raises the GIS-5 flag on every access figure in this chapter.

7.1 What the module does

TI-04 turns the inventory into a routable picture of the network *as it actually is*, and then measures access on that network.

- **Inputs.** The noded road network from the spatial base (Chapter 4): 817,354 road links and 28 ferry links, 247,972 km in all, each with its functional class, surface (where known), terrain, structure flags and flood exposure. The season rules, speed profiles and walking-leg norms from `tpl_norm`. The 24,952 sector nodes (schools, health facilities, markets, district headquarters, gateways, industrial parks, mines). The 175,367 settlement cells with their 2026 population. The agency condition statistics published by the Ministry of Works and Transport (MoWT) and the Uganda Bureau of Statistics (UBOS).
- **What it computes.**
  - the *seasonal availability* of every link: open or closed in the wet season;
  - the *speed and travel time* on every link, by mode and season;
  - the *capacity* of every link whose lane count is known;
  - the *do-nothing condition* of every link whose condition is known (none, in Uganda today);
  - the *current accessibility* of every settlement to every destination class: the travel time (or distance) to the nearest functional destination, on the wet-season network;
  - the *Rural Access Index* (RAI), the SDG 9.1.1 indicator, by district and nationally, as a bound.
- **Outputs.** Link supply and access tables in the store, the RAI by district (`tables/ti04_rai_by_district.csv`), the RAI bounds (`tables/ti04_rai_bounds.csv`), and GIS layers of wet-season closures and access by settlement. TI-05 (Chapter 8) compares these with the requirement to find the gaps.

Two principles govern the module. First, **no attribute is inferred**. Surface comes only from an observation or a model of that specific road, never from its class; condition comes only from a survey. Second, **the planning season is the wet season**. A road that is open in the dry season but cut by a drift in the rains does not give all-season access, and the access figures say so.

7.2 Method

7.2.1 Seasonal availability of a link

A link *l* is available in season *s* only if three tests pass: its surface is acceptable in that season, its condition is acceptable, and every structure on it is passable. Each factor is 0 or 1 and comes from the season-rule table `tpl_norm.season_rule`:

$$a_{l,\,s} = \mathrm{surf}_{ok}(l,\,s) \cdot \mathrm{cond}_{ok}(l) \cdot \prod_{b \,\in\, B(l)} \mathrm{pass}(b,\,s) \tag{7.1}$$

Here *B(l)* is the set of structures on the link. In the Uganda implementation:

- **Surface test.** Paved (asphalt concrete, double bituminous surface treatment, low-volume seal, concrete), gravel and engineered-earth links are open in the wet season unless they are in bad condition. Unengineered earth and unknown-surface links are open unless they are flood-exposed (JRC 100-year flood depth above zero at the link midpoint, or tagged `flood_prone` in OpenStreetMap).
- **Condition test.** A gravel or engineered-earth link in *bad* condition closes. Because no link has a known condition, this test never closes a link in the reference run.
- **Structure test.** A link with a seasonal crossing (an OSM ford point within 25 m, a way tagged `ford`, or a low-water crossing) closes in the wet season unless a funded spot improvement has fixed the crossing. Ferries are assumed to run all year (an assumed norm).

7.2.2 Link travel time

The travel time on a link in minutes is its length divided by the speed for its mode, surface, condition, terrain and season, looked up in `tpl_norm.speed_profile`:

$$t_{l,\,s} = \frac{60\,\mathrm{len}_l}{V(\mathrm{mode}_l,\,\mathrm{surf}_l,\,\mathrm{cond}_l,\,\mathrm{terrain}_l,\,s)} \tag{7.2}$$

A missing lookup raises `NormNotFound`; the engine never guesses a speed. Where OSM carries a posted speed limit, it caps the motorised speed but never raises it. A closed link has an infinite time. The speed profiles are TransportModelling generic defaults, marked "NOT Ugandan" in the norm table; ferries are given an assumed 8 km/h including waiting time.

7.2.3 Current accessibility

For each destination class *k*, the engine runs one multi-source shortest-path search (Dijkstra's algorithm) on the reversed wet-season graph, starting from every functional destination at once. It joins a super-source to the snapped vertex of every destination *j* in the set *J(k)*, with an arc cost equal to the destination's own walking leg *w(j)*:

$$D_k(v) = \min_{j \in J_k} \left(d_{G^{\,wet}}(v,\,\mathrm{snap}(j)) + w_j \right) \tag{7.3}$$

D(k,v) is then the shortest time from network vertex *v* to the nearest destination of class *k*. A settlement *i* adds its own walking leg from the cell to the network:

$$tt_{i,\,k} = w_i + D_k(\mathrm{snap}(i)), \qquad w_i = \frac{60\,\delta_i\,\varphi_i}{v_{walk}} \tag{7.4}$$

where *δ(i)* is the straight-line snap distance in km, *φ(i)* the terrain factor of the snapped link (1.0 flat, 1.15 rolling, 1.4 mountainous, all assumed until SDI supplies slope) and *v(walk)* the walking speed of 5 km/h. The settlement meets the standard of class *k* if its time (or distance, for classes measured in km) is within the threshold *τ*, which depends on whether the cell is rural or urban:

$$A_{i,\,k} = \mathbf{1}\left(tt_{i,\,k} \le \tau_{k,\,u(i)}\right) \tag{7.5}$$

The thresholds come from `tpl_norm.access_standard`, summarised in Table 7.1. Only the RAI distance and the two-hour EmONC time are *normed* (taken from an international standard); the others are *assumed* until the Health, Education and Agriculture engines supply national values.

Table 7.1: Access standards used in the reference run (rural / urban)

Destination class	Measure	Rural	Urban	Mode	Quality
All-season road (RAI)	distance	2 km	0.5 km	walk	normed (SDG 9.1.1)
Primary health facility	time	60 min	30 min	mixed	assumed
EmONC facility	time	120 min	60 min	motorised	normed (WHO / Lancet)
Primary school	network distance	3 km	1.5 km	walk	assumed
Secondary school	time	60 min	30 min	mixed	assumed
Market	time	120 min	30 min	motorised	assumed
Administrative (district) centre	time	180 min	60 min	motorised	assumed

7.2.4 The Rural Access Index

The RAI is the share of rural people living within 2 km of an all-season road. For the all-season class the engine measures the straight-line distance from each rural cell to the nearest link that is wet-season available *and* has an all-season surface (paved, gravel or engineered earth, not in bad condition). For district *d*, with *P(i)* the 2026 population of rural cell *i*:

$$\mathrm{RAI}_d = \frac{\sum_{i \,\in\, I_d^R} P_i \,\mathbf{1}(\delta_i^{as} \le 2)}{\sum_{i \,\in\, I_d^R} P_i} \tag{7.6}$$

Because surface is unknown on most of the network, the engine computes the RAI twice. The **plan figure** counts unknown surface as earth, which is never all-season. The **upper bound** counts it as all-season gravel. The truth lies between the two, and the published SDG benchmark is printed beside them.

7.2.5 Capacity and condition

Road capacity in passenger-car units per day is computed only where the lane count is known, from the Highway Capacity Manual-derived norm in `tpl_norm.road_capacity`. Lanes are known on 23,521 links (4,464 km). Every other link has *unknown* capacity; it is listed as such in Chapter 8 and never assumed adequate.

Condition would be projected forward on the do-nothing path with the deterioration model of `docs/06` (§4.3). In Uganda today **no link carries a measured condition**: the engine reports "condition known on 0 of 817,382 links", and unknown condition is not projected. Network-level condition is taken instead from agency statistics, by network class, and is never spread onto links.

7.3 Results: what the network is made of

The network that TI-04 routes on is the one described in Chapter 4: 247,972 km, of which 160,842 km are community roads, 46,117 km tracks, 22,052 km district roads and 15,536 km national roads (981 km trunk, 3,415 km primary and 11,140 km secondary). Three features of this inventory shape every result in the chapter.

- **Surface is known on 31.7 % of km.** The known part is 7,269 km paved (4,274 km asphalt concrete, 2,970 km surface dressing, 25 km low-volume seal and concrete), 1,972 km gravel and 69,388 km earth. The remaining 169,343 km have no surface record. Of the known surface, 73,748 km comes from OSM `surface` tags (*observed*), 327 km from OSM track grades, and 4,831 km from the HeiGIT 2024 deep-

learning classification and the GRIP4 inventory (*estimated*). Unknown surface is concentrated on the lower classes: 74.3 % of community-road km, 78.7 % of track km and 49.0 % of district-road km, against 3.0 % or less on national roads.

- **Few dimensional attributes exist.** Width is recorded on 847 km and lane count on 4,464 km. Posted speeds exist on 2,182 km and OSM smoothness on 989 km; smoothness is stored as a proxy for validation only and never sets condition.
- **Structures are recorded only partly.** OSM marks 2,188 bridges and 5,039 culverts (snapped to 3,676 links). 2,248 links carry a seasonal crossing (3,943 ford points, 250 ford ways and 12 low-water crossings). A further 1,958 river crossings on 1,849 links have no recorded bridge, culvert or ford; they are a survey list (Chapter 8), not a closure. 3,006 links (1,368 km) are flood-exposed.

7.4 Results: links closed in the wet season

Figure 7.1 maps the links that the season rules close in the wet season. There are **5,013 closed links, totalling 3,299 km**, or 1.3 % of the network by length. The closures have two causes, which rarely coincide:

- **Seasonal crossings** close 2,248 links (2,139 km of link length). A ford or drift cuts the whole link, so a long link with one drift counts in full.
- **Flood exposure** closes 2,874 earth and unknown-surface links (1,276 km). Only 109 links are closed by both.

By class, community roads account for 2,383 closed links (1,525 km) and tracks for 1,860 (1,187 km). District roads have 400 closed links (305 km). The national network is not immune: 114 national secondary links (183 km), 15 national primary links (24 km) and 2 trunk links (23 km) close, almost all because a ford is mapped on them. By surface, 2,772 closed links are earth (1,977 km) and 2,201 are of unknown surface (1,275 km); 28 gravel links and 12 paved links close because of a mapped crossing.

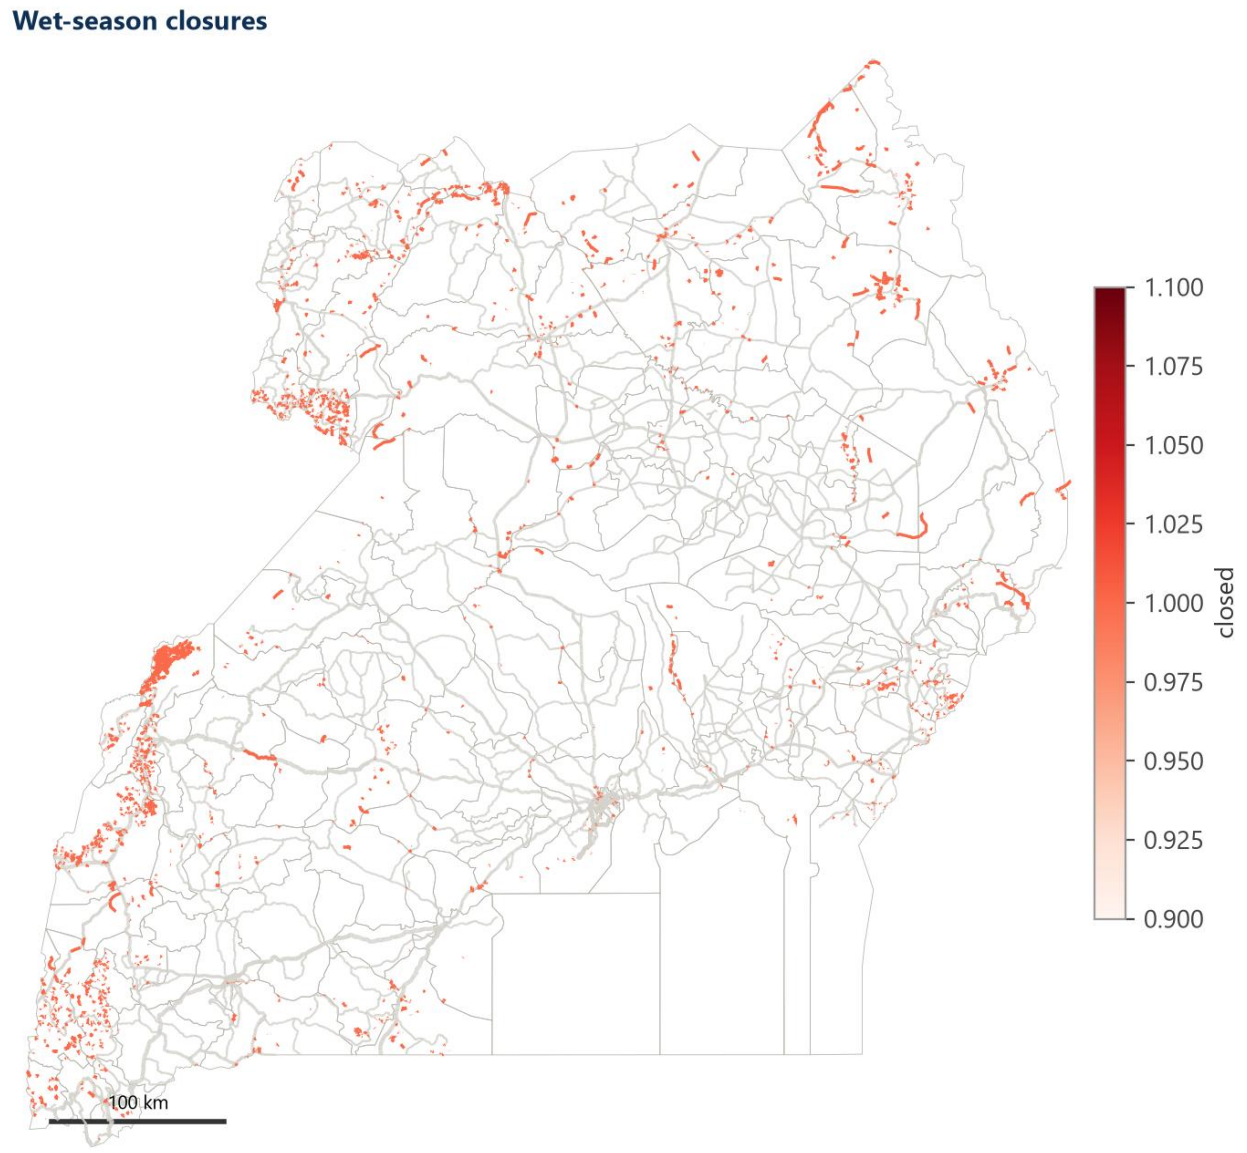

Figure 7.1: Links closed in the wet season (season rules on surface, terrain and river crossings)

Source: TI-04 season rules (norms)

The closures are strongly clustered. The map shows three concentrations.

- **The Semliki floodplain and the Rwenzori foothills.** Ntoroko has 1,058 closed links and 463 km closed, by far the largest in the country; Kasese has 299 links and 225 km, and Bunyangabu 97 links. Here flood exposure on the low-lying plain dominates.
- **The West Nile and Albert Nile belt.** Nebbi (166 km), Arua (230 links, 86 km), Zombo (87 km), Adjumani (175 links, 88 km), Obongi (171 links) and Amuru (138 links) all show dense clusters along the river and its tributaries.
- **Karamoja.** Kotido (145 km), Karenga (120 km) and Kaabong (60 km) have long closures where seasonal river crossings cut the few roads that exist.

By region, the Northern districts have 1,417 km closed and the Western districts 1,367 km, against 305 km in the East and 209 km in the Central region. Kanungu (115 km) and Rukungiri (93 km) in the south-west are the other notable cases. These kilometres are the G4 "not all-season" gap of Chapter 8; the closures are what turns a road on the map into a road that does not count for access in the rains.

7.5 Results: the Rural Access Index

7.5.1 The national figure is a bound

Figure 7.2 and Table 7.2 give the national Rural Access Index. Counting unknown surface as earth, **28.7 % of rural people live within 2 km of an all-season road**. Counting it as gravel, the figure is **93.5 %**. The published SDG 9.1.1 value for Uganda is 53.1 % (2015, World Bank).

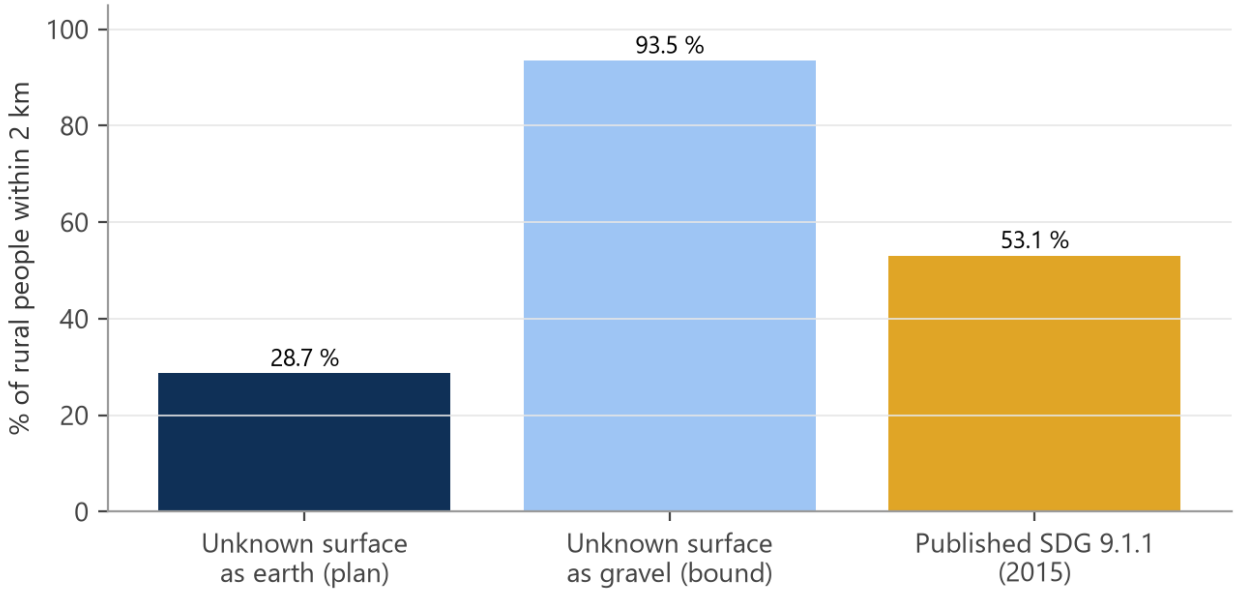

Figure 7.2: Rural access today: plan figure, upper bound from unknown surface, and the published benchmark

Source: TI-04; UNSD SDG 9.1.1

Table 7.2: Rural Access Index today: plan figure, upper bound and published benchmark

Measure	Value
RAI_unknown_surface_as_earth (plan figure)	28.7%
RAI_unknown_surface_as_gravel (upper bound)	93.5%
published SDG 9.1.1 benchmark	53.1 % (2015, World Bank)

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_rai_bounds.csv).

The width of this bound, 64.8 percentage points, is a statement about data, not about roads. It is the same 169,343 km of unknown-surface links that moves the figure: if they are earth, most rural people are far from an all-season road; if they are gravel, almost everyone is near one. The published benchmark falls between the two bounds, as it should. The plan figure is the conservative one, and it is the figure the engine plans against, because the norms forbid treating an unrecorded road as all-season. The 2015 benchmark used a different network and population base, so the gap between 28.7 % and 53.1 % does not show that access has fallen; it shows how much of the network is unrecorded. Adding the HeiGIT and GRIP4 surface sources in run 3 raised the plan figure from 20.2 % to 28.6 % (`IMPLEMENTATION_STATUS.md`), which shows how directly it responds to surface evidence; Chapter 20 covers the data acquisition that would close the bound.

Across all people, urban and rural, 35.9 % of the population meets the all-season-road standard (urban cells must be within 0.5 km). The national RAI, which covers only the 41.6 million rural people, is the lower 28.7 %.

7.5.2 Districts

Figure 7.3 maps the plan figure by district, and Figure 7.4 shows its distribution. District values range from 0 to 83.0 %, with a median of 26.8 %. Fifty districts are below 20 %, 56 lie between 20 % and 40 %, and 29 are at 40 %

or more; only 14 reach 50 %. The distribution is wide and skewed: a long tail of a few well-served districts sits above 75 %.

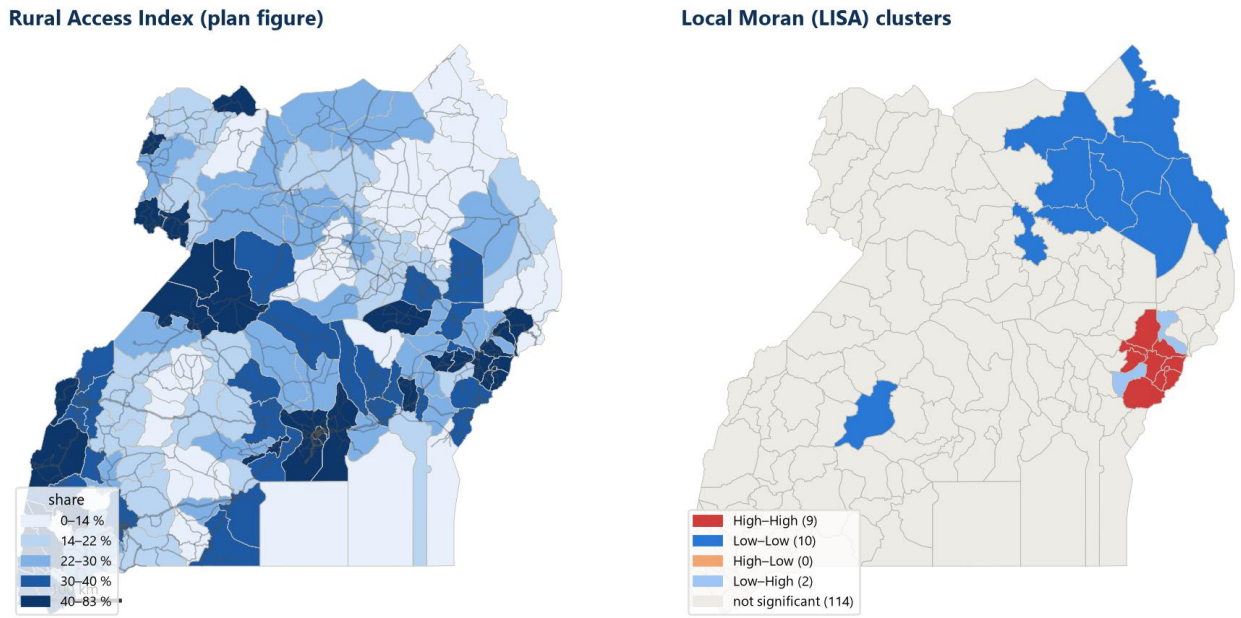

Figure 7.3: Rural Access Index today by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

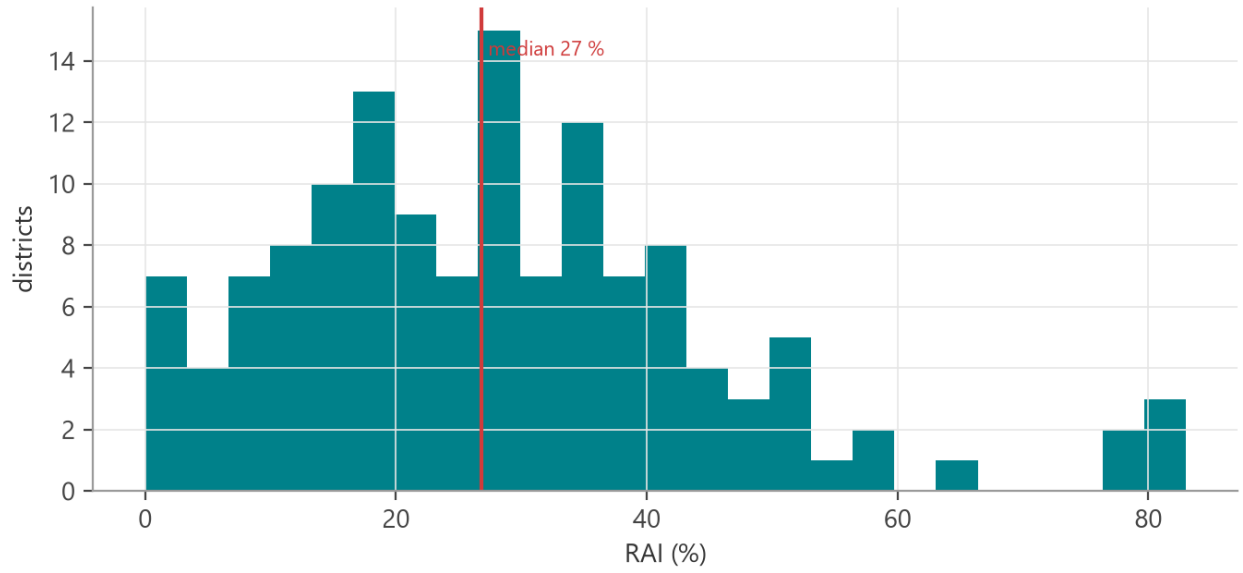

Figure 7.4: Distribution of the Rural Access Index across districts

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The highest values are in the Mount Elgon area and the capital: Manafwa 83.0 %, Kampala 82.6 %, Mbale 81.5 %, Bududa 79.5 % and Kapchorwa 76.4 %, followed by Sironko (66.3 %), Zombo (59.4 %), Iganga (58.8 %), Serere (55.1 %) and Kasese (52.7 %). The plan figure can only count roads whose surface is recorded, so in part these high values may reflect denser surface tagging in OSM rather than better roads; the upper bound is the check.

The right panel of Figure 7.3 shows the local Moran (LISA) clusters. Global Moran's *I* for the district RAI is 0.363 (p = 0.001), the second-strongest spatial clustering of any layer in the run (Chapter 21).

- **High–high (9 districts):** Budaka, Bududa, Bukedea, Butebo, Manafwa, Mbale, Namisindwa, Sironko and Tororo, a contiguous block in the east.

- **Low–low (10 districts):** Abim, Agago, Kaabong, Kitgum, Kotido, Lira, Moroto, Napak and Otuke in the north and north-east, and Mubende in the centre-west.
- **Low–high (2):** Bulambuli and Butaleja, low-access districts next to the eastern high cluster.

At regional level, population-weighted, the rural RAI is 34.7 % in the Eastern region, 28.8 % in the Central, 28.3 % in the Western and 21.5 % in the Northern region.

7.5.3 The districts with the lowest access

Figure 7.5 lists the 25 districts with the lowest RAI, and Table 7.3 gives the 15 lowest.

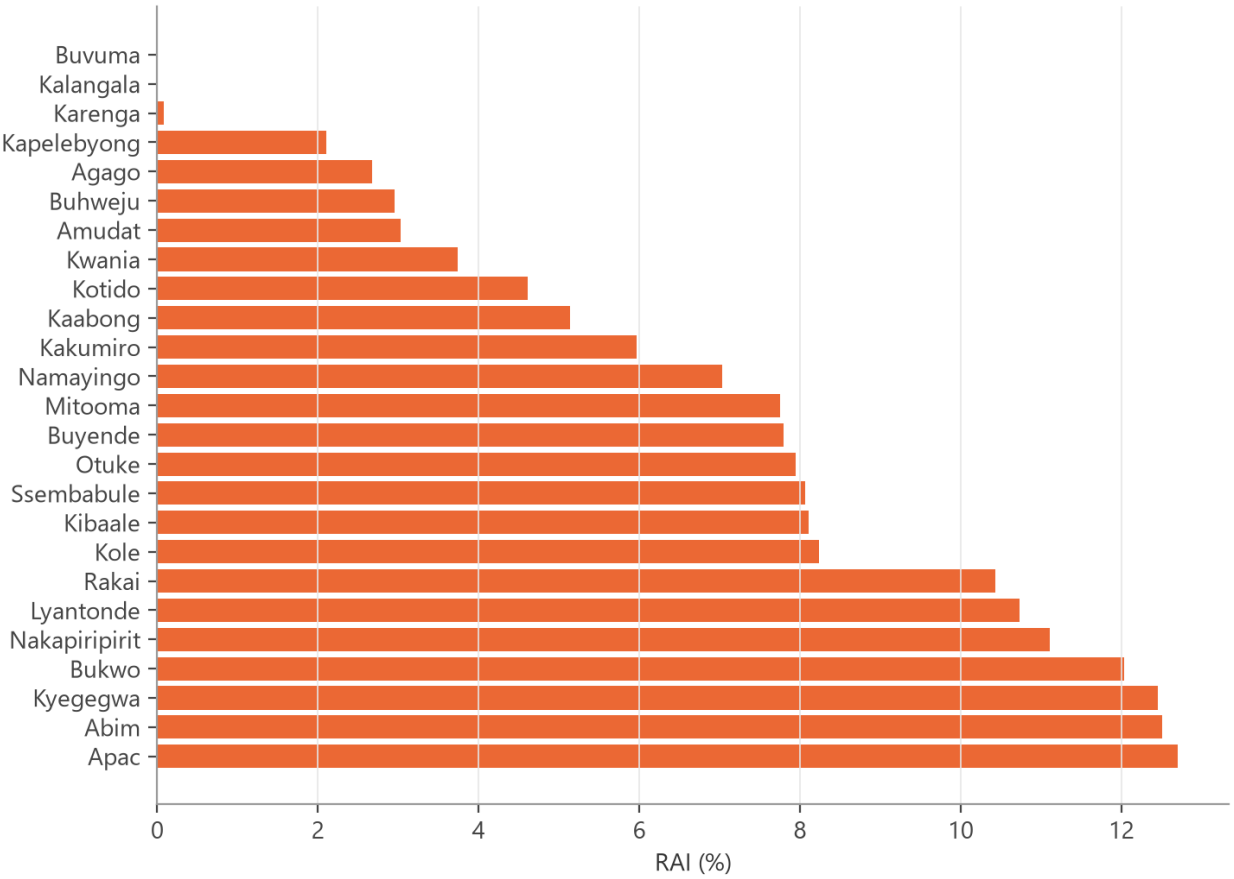

Figure 7.5: The 25 districts with the lowest rural access

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 7.3: The 15 districts with the lowest Rural Access Index today

District	RAI today
Buvuma	0.0%
Kalangala	0.0%
Karenga	0.1%
Kapelebyong	2.1%
Agago	2.7%
Buhweju	2.9%
Amudat	3.0%
Kwania	3.7%
Kotido	4.6%
Kaabong	5.1%
Kakumiro	6.0%
Namayingo	7.0%
Mitooma	7.7%

District	RAI today
Buyende	7.8%
Otuke	7.9%

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_rai_lowest.csv).

Three groups stand out.

- **Island districts.** Buvuma and Kalangala have an RAI of 0.0 %. Their roads have no recorded all-season surface, and the islands are not joined to the mainland network except by ferry. A ferry link counts for travel time but not as an all-season road.
- **Karamoja and the north-east.** Karenga (0.1 %), Amudat (3.0 %), Kotido (4.6 %), Kaabong (5.1 %), Kapelebyong (2.1 %), Agago (2.7 %), Otuke (7.9 %) and Nakapiripirit (11.1 %) form the bulk of the list. These are the districts of the low–low LISA cluster.
- **Districts with long unrecorded networks in the west and centre.** Buhweju (2.9 %), Kwania (3.7 %), Kakumiro (6.0 %), Mitooma (7.7 %), Buyende (7.8 %), Ssembabule (8.1 %), Kibaale (8.1 %) and Kole (8.2 %). Here the network is extensive but mostly unrecorded or earth, so the plan figure is low and the upper bound is high. These districts would move most if a surface survey found gravel.

For planners, the second and third groups call for different first steps. In Karamoja the roads are few and the gap is physical. In the third group the first step is to find out what the roads are.

7.6 Results: travel time to services

7.6.1 Five destination classes in the wet season

Figure 7.6 to Figure 7.10 map the wet-season travel time from every settlement to the nearest destination of each class, as hexagon medians capped at 240 minutes. Figure 7.11 maps the distance to an all-season road. White areas carry no value: either no settlement is there (game reserves, lakes, forests) or the settlement has no route to that class on the wet-season network.

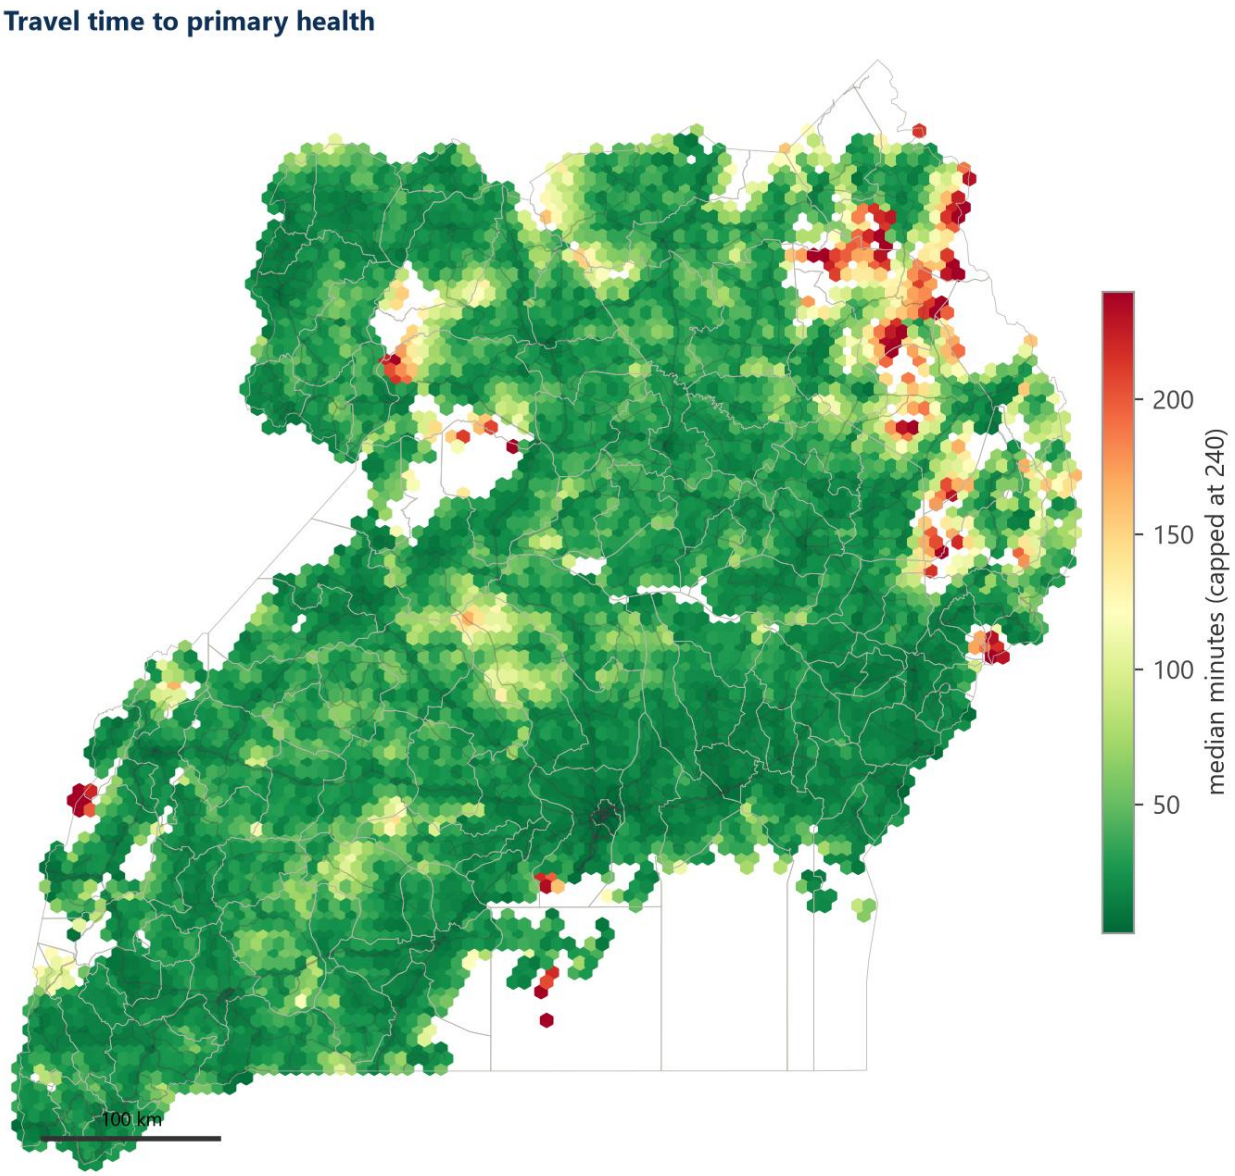

Figure 7.6: Travel time to primary health, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

Primary health. Most of the country is within an hour of a health facility (Figure 7.6). The population-weighted median time is 15.4 minutes, and 93.8 % of the requiring population meets the standard of 60 minutes rural, 30 minutes urban. The red patches are few and local: north-east Karamoja (Kotido, Kaabong, Moroto, Napak), small pockets on the Uganda–South Sudan and Uganda–Kenya borders, the islands of Lake Victoria and parts of Kasese and Bundibugyo. The Northern region has the slowest access: a median of 20.1 minutes and 88.7 % within standard, against 9.5 minutes and 95.2 % in the Central region.

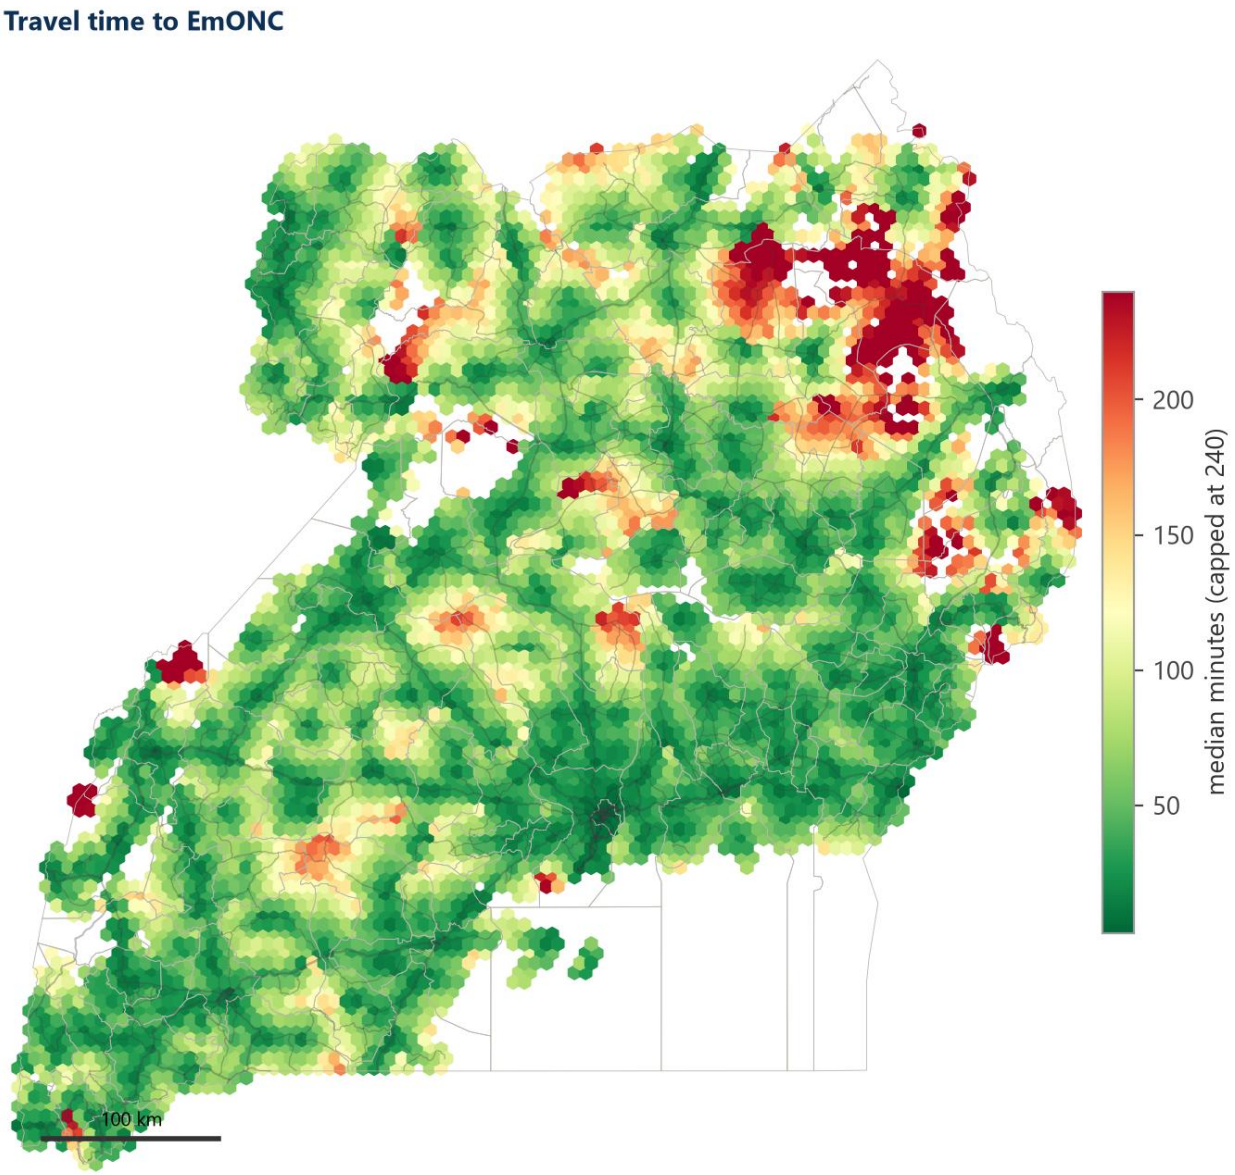

Figure 7.7: Travel time to EmONC, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

Emergency obstetric and newborn care (EmONC). The two-hour EmONC standard is met for 91.5 % of women of reproductive age (Figure 7.7). The median time is 33.7 minutes, but the map shows a large contiguous red area across Karamoja and Agago, where times exceed 200 minutes, and smaller hot spots in Kwania and Apac, Mubende, Ssembabule and the West Nile border. In the Northern region the median is 53.7 minutes and only 81.5 % are within two hours, against 94 % in each of the other regions.

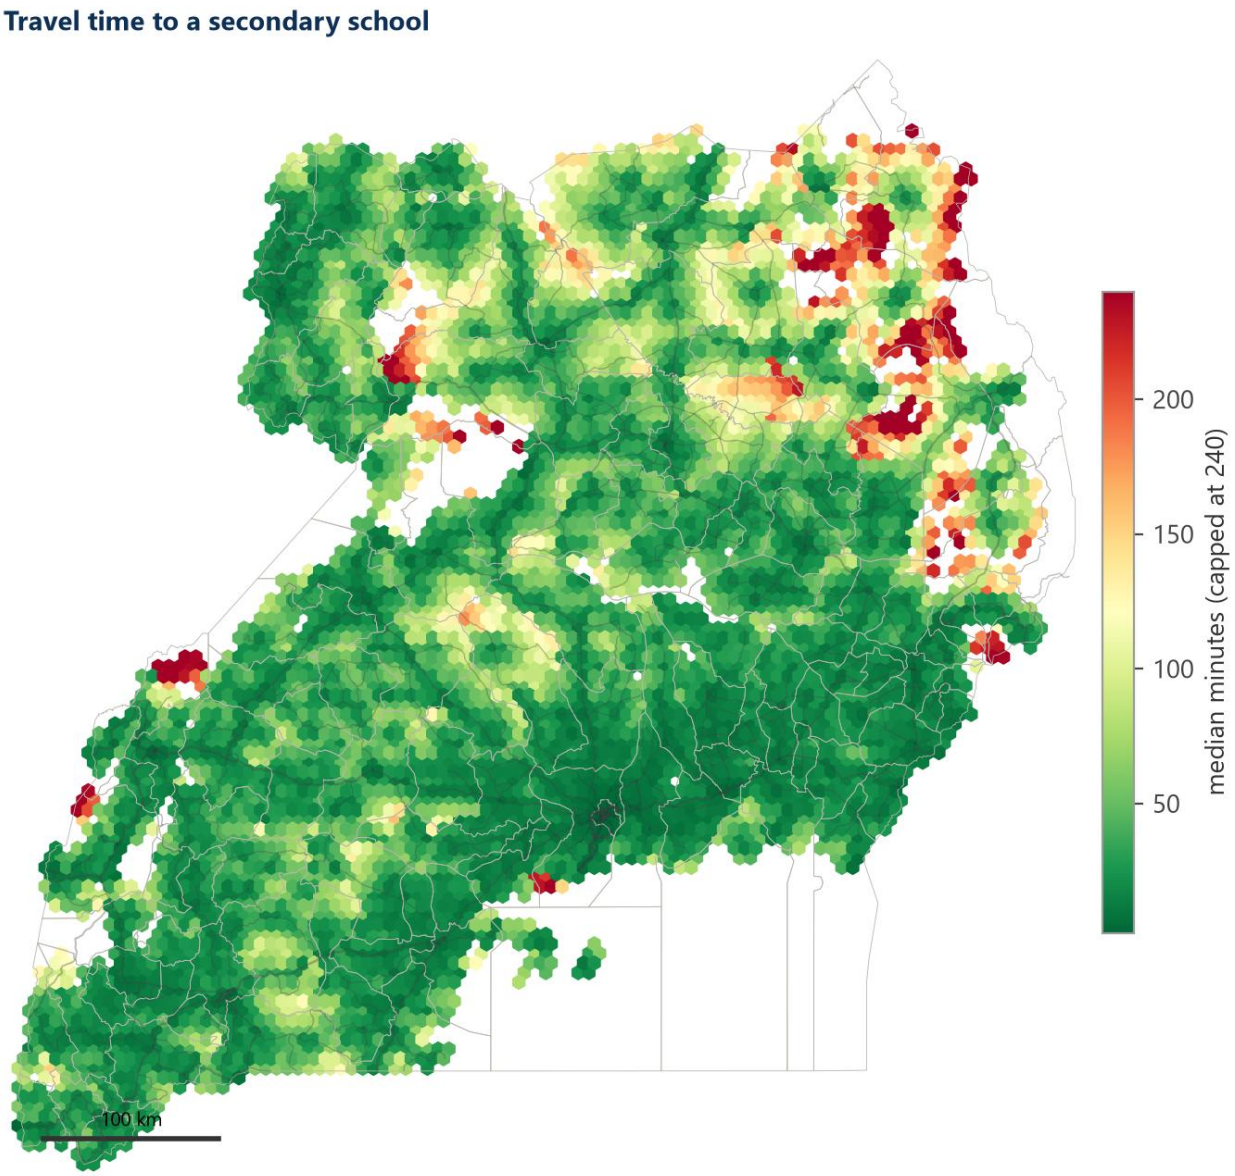

Figure 7.8: Travel time to a secondary school, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

Secondary school. The pattern resembles primary health but is slower (Figure 7.8). The median is 16.5 minutes, and 89.0 % of young people aged 13–18 meet the assumed 60-minute standard. Northern Uganda stands apart again: a median of 31.5 minutes and 72.5 % within standard. Alebtong, Otuke, Amudat, Agago, Kotido and Napak have the largest numbers beyond the standard (Chapter 8).

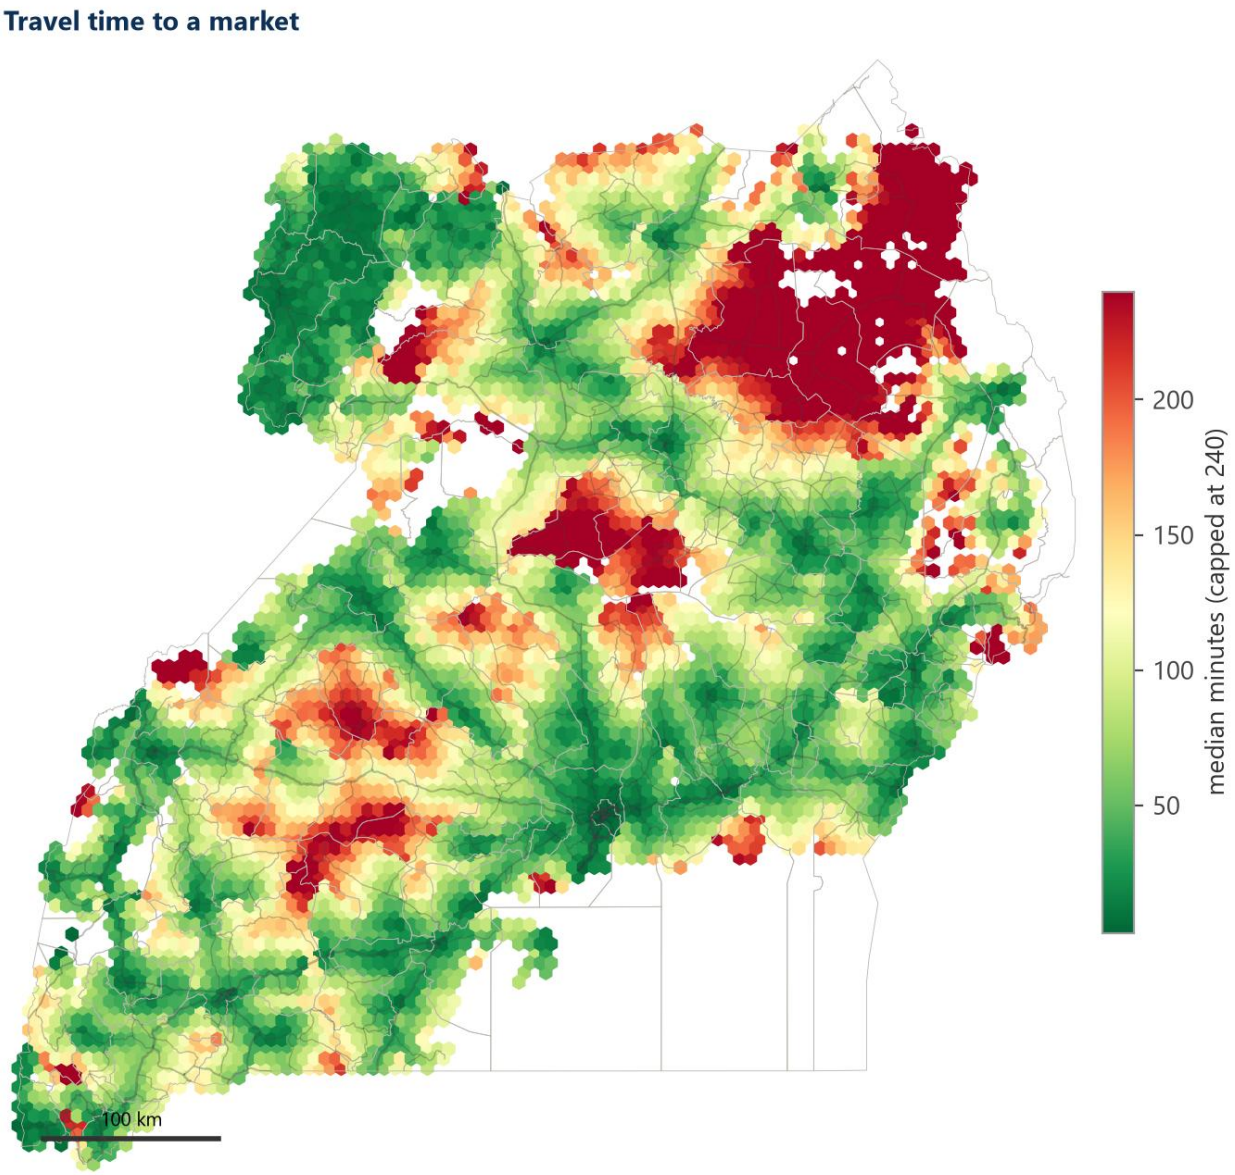

Figure 7.9: Travel time to a market, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

Markets. Market access is the weakest of the service classes (Figure 7.9). The median is 47.1 minutes, and 79.8 % of households meet the assumed two-hour standard. The map shows three large red zones: the whole of Karamoja east of Kitgum and Agago; a central block around Nakasongola, Kyankwanzi and Kiboga; and a western block across Kibaale, Kakumiro, Kagadi, Mubende and Kyegegwa. These zones partly reflect the market register: the snapshot holds 814 crop markets, so a district with few registered markets appears remote even where informal markets exist. The Western region has the highest median (67.7 minutes), ahead of the Northern (60.2 minutes).

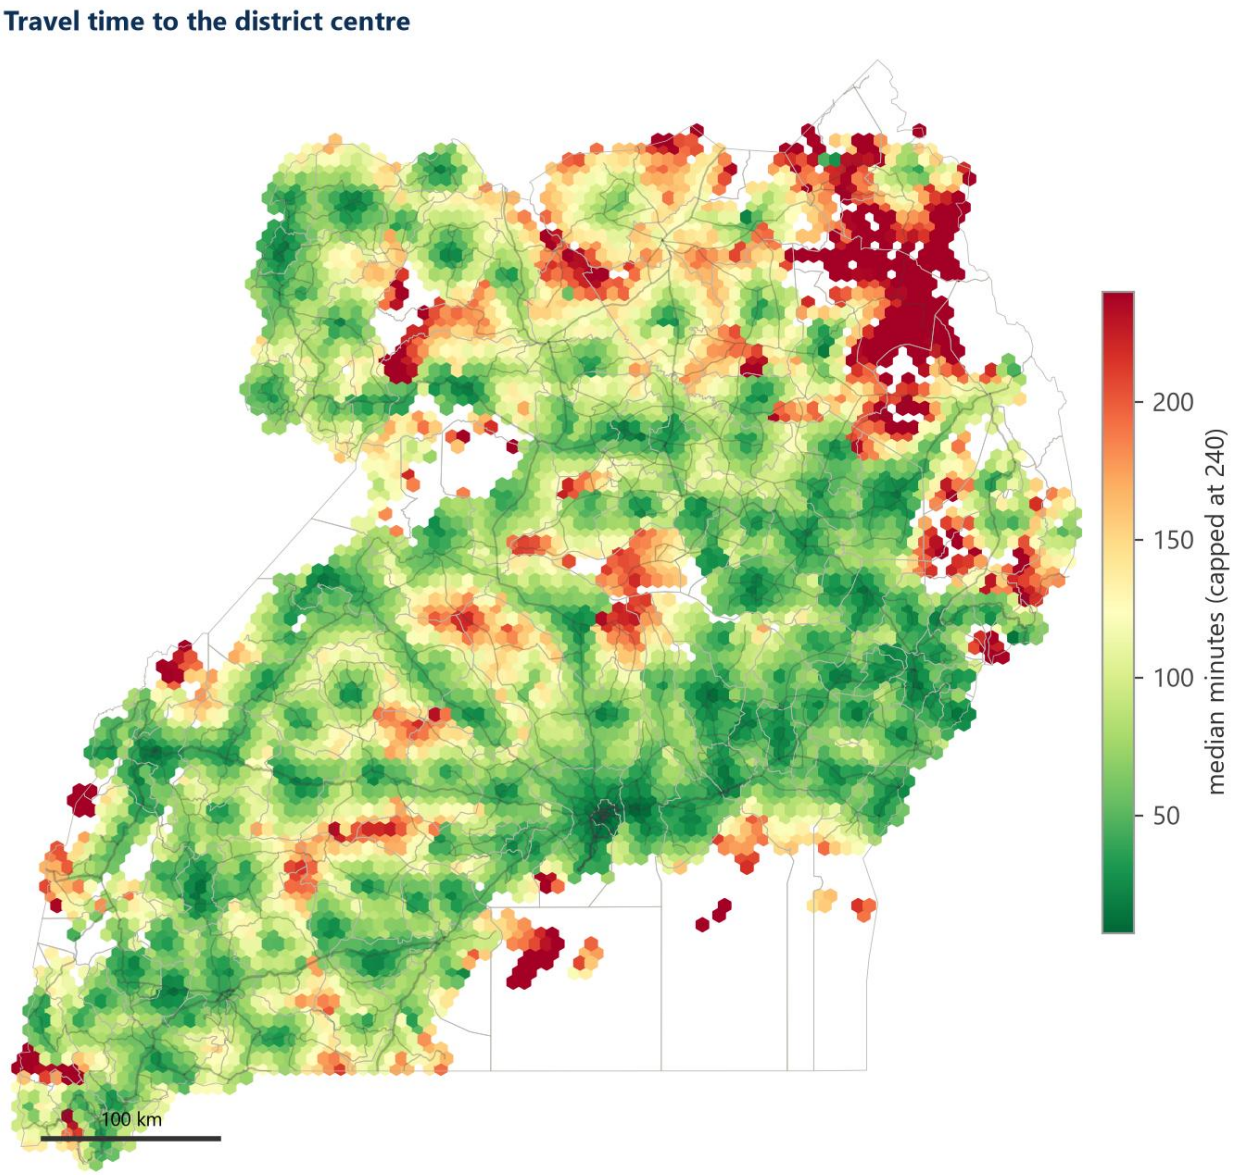

Figure 7.10: Travel time to the district centre, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

District headquarters. The district centre is the only class with one destination per district (135 in all), so times are longer everywhere (Figure 7.10): the median is 56.1 minutes and 91.9 % meet the three-hour rural standard. Karamoja is again the largest red area, with further patches in Kasese and Bundibugyo, Buvuma and the islands, northern Kitgum and Lamwo, and Gulu district.

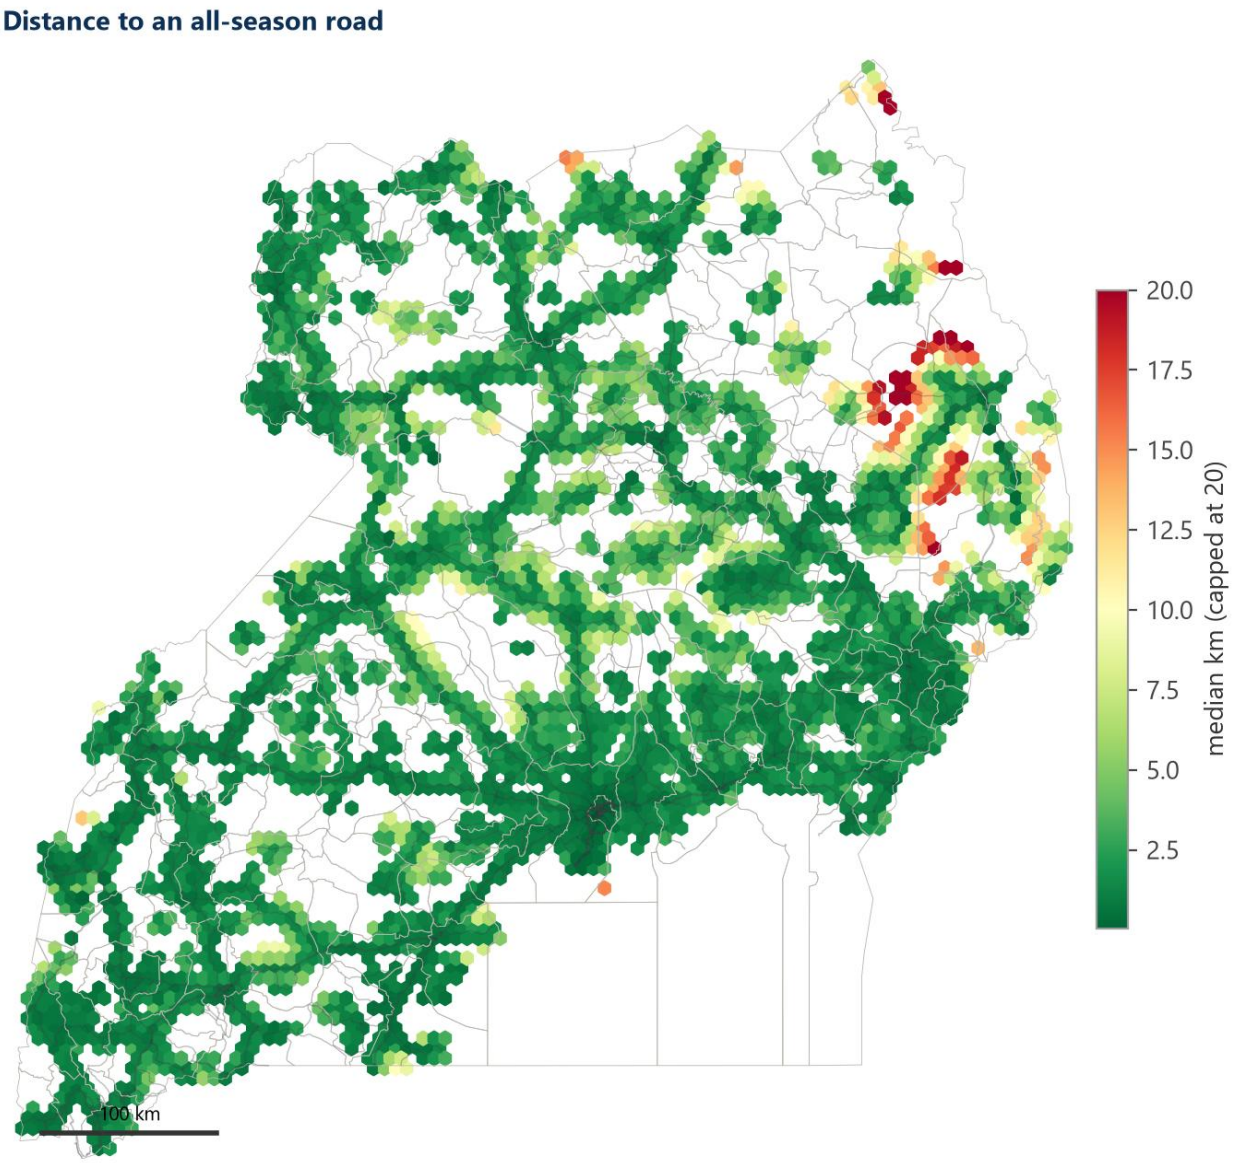

Figure 7.11: Distance to an all-season road, wet season, from every settlement (hexagon median)

Source: TI-04 multi-source Dijkstra on the seasonal network; walk legs per norms

All-season road. Figure 7.11 shows distance to the nearest all-season link. The green corridors follow the national roads and the surfaced and gravel roads that OSM records; the large white areas between them are settlements for which no all-season link was found in the nearest-neighbour search, which on this network means far beyond 2 km. They hold 28.0 million people, 54 % of the population. This map is the spatial form of the RAI bound: the white is mostly where surface is unknown.

The primary-school class is measured as walking distance on the network (3 km rural). Its median is 1.72 km; 73.0 % of children aged 6–12 meet the standard, the lowest share of any class after the all-season road.

7.6.2 Access curves

Figure 7.12 summarises the five time-based classes as cumulative curves: the share of people (with a wet-season route) reached within a given time.

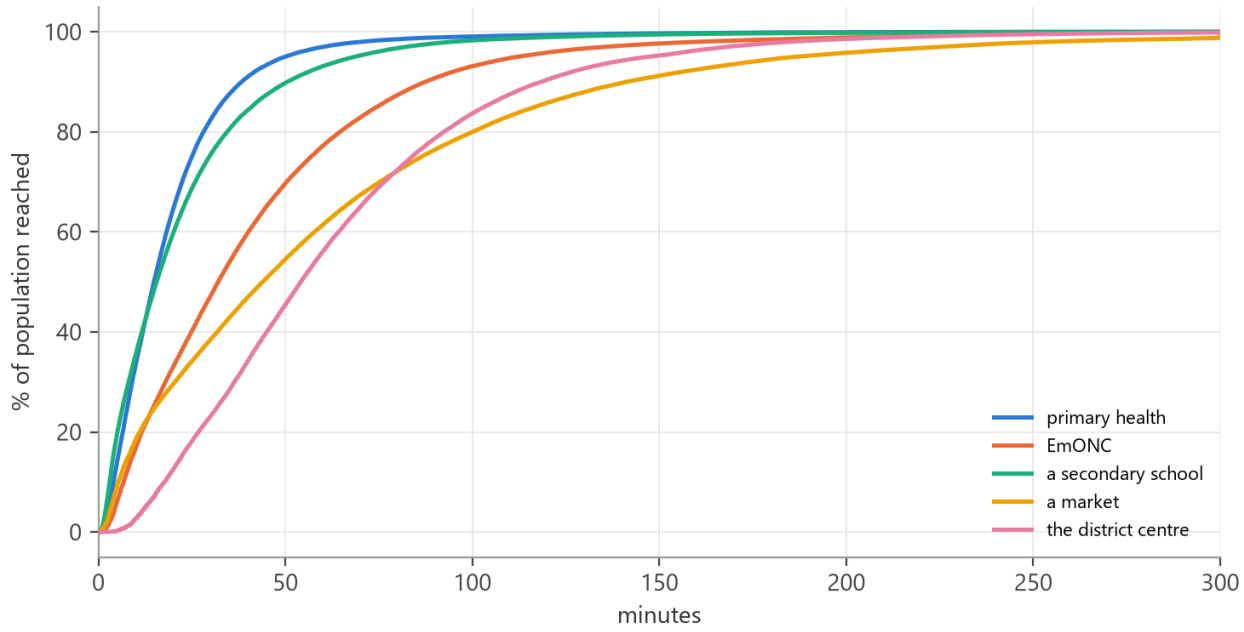

Figure 7.12: Share of people reached within a given travel time, by destination class (wet season)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The curves separate into three groups.

- **Primary health and secondary school** rise steeply: about 82 % of people with a route reach a health facility within 30 minutes and 97 % within an hour; for secondary school the shares are about 75 % and 93 %.
- **EmONC** rises more slowly because there are fewer facilities: about 47 % within 30 minutes, 77 % within an hour and 96 % within two hours.
- **Markets and the district centre** rise slowest. About 61 % reach a market within an hour and 85 % within two hours; for the district centre the shares are 56 % within an hour and 90 % within two hours. The market curve has the longest tail: about 6 % of people with a route need more than three hours.

Beneath the curves are the people with no wet-season route at all: about 3.1 % of the population for primary health and 3.9–4.1 % for the other time-based classes (1.6 to 2.1 million people). They are cut off by a closed crossing, a flooded road or a network component without that class of destination. Chapter 8 counts them as unserved and explains why.

7.6.3 Inequality of access

Access is markedly worse for poorer people. Using the International Wealth Index of each settlement cell, the population-weighted median time to primary health is 30.1 minutes in the poorest fifth of cells and 7.9 minutes in the richest fifth. For EmONC the medians are 67.1 and 13.8 minutes; for markets 102.4 and 16.2 minutes; for secondary school 42.9 and 6.1 minutes. The poorest quintile is also the furthest from an all-season road: its median rural cell has no all-season link nearby, while the richest quintile's median distance is 0.63 km. Chapter 12 reports the RAI by wealth quintile for each scenario (0.100 in the poorest quintile against 0.521 in the richest on the do-nothing path).

The population-weighted Gini coefficient of travel time to primary health, computed over settlements with a route, is 0.44 (0.52 for markets, 0.46 for EmONC). These numbers describe dispersion of travel time across people, not inequality of income, and they exclude the people with no route.

7.7 Results: condition from agency statistics

No link-level condition survey exists in the estate. The agency road-condition systems (UNRA and DUCAR road asset management systems) hold IRI and visual surveys, but they are not public. TI-04 therefore carries condition as *unknown* on every link, and the condition picture comes from published statistics by network class.

Figure 7.13 and Table 7.4 give the FY2022/23 figures, from the UBOS Statistical Abstract 2024 (Figures 10.1 and 10.2, national roads) and the MoWT Statistical Abstract 2023 (Table 3, district, community and urban roads, from the Uganda Road Fund DUCAR status report). Their quality tier is *observed*, administratively reported.

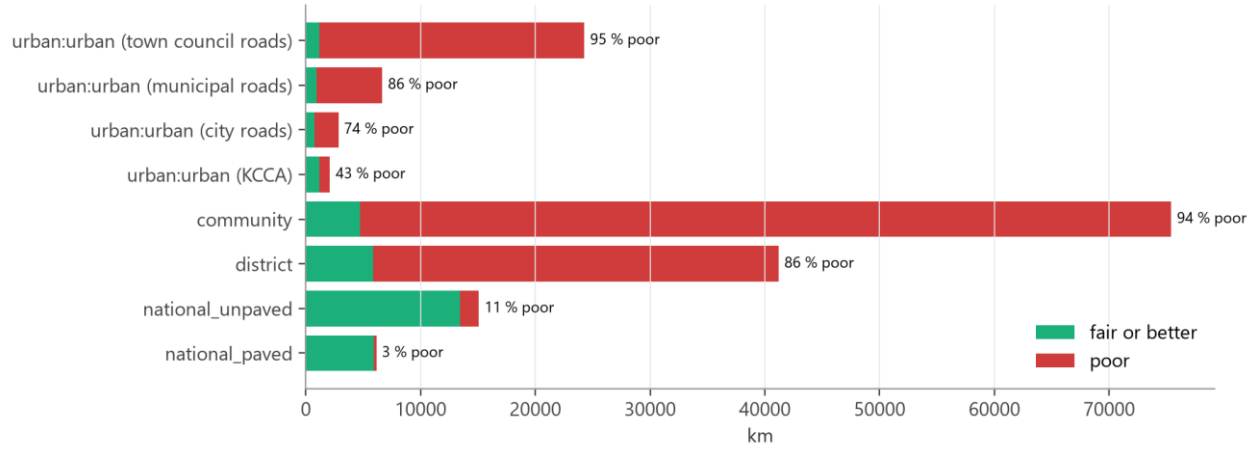

Figure 7.13: Road condition by network: agency-reported km in poor condition

Source: UBOS Statistical Abstract 2024; MoWT Statistical Abstract 2023 (URF DUCAR status)

Table 7.4: Road condition by network, agency statistics

Network	Km	Km poor	Poor share	Year	Basis
national_paved	6,133	204	3.3%	FY2022/23	observed
national_unpaved	15,067	1,657	11.0%	FY2022/23	observed
district	41,194	35,320	85.7%	FY2022/23	observed
community	75,404	70,656	93.7%	FY2022/23	observed
urban:urban (KCCA)	2,103	899	42.7%	FY2022/23	observed
urban:urban (city roads)	2,830	2,100	74.2%	FY2022/23	observed
urban:urban (municipal roads)	6,656	5,740	86.2%	FY2022/23	observed
urban:urban (town council roads)	24,269	23,115	95.2%	FY2022/23	observed

Administratively reported (UNRA/URF via UBOS and MoWT abstracts). No link-level condition survey exists. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_condition.csv).

The figures describe two different networks.

- **The national network is mostly fair or better.** Of 6,133 km of paved national roads, 204 km (3.3 %) are poor. Of 15,067 km of unpaved national roads, 1,657 km (11.0 %) are poor.
- **The local network is mostly poor.** Of 41,194 km of district roads, 35,320 km (85.7 %) are poor. Of 75,404 km of community access roads, 70,656 km (93.7 %) are poor. Urban roads range from 42.7 % poor in Kampala (KCCA, 899 of 2,103 km) to 95.2 % on town-council roads (23,115 of 24,269 km).

In total, 139,691 km of public roads are reported poor; 107,837 km of them are on the non-urban networks. Three cautions apply when reading these numbers beside the network map.

- The agency kilometres are not the snapshot's kilometres. The agency counts 41,194 km of district roads; the snapshot classifies 22,052 km as DISTRICT, and much of the rest falls under COMMUNITY or TRACK. The two classifications cannot be reconciled link by link without the agency's road register.

- The statistics say how many kilometres are poor, not which ones. The engine therefore never paints a link poor. A map of condition would be an invention, and the design forbids it.
- Because condition is unknown per link, condition does not slow or close any link in the access calculation. The access results above are what the network gives if condition were no obstacle, which on the local networks it clearly is.

Chapter 8 uses these statistics as the G2 condition gap, and Chapter 9 costs the treatment of the poor kilometres.

7.8 Data quality and limitations

The access results are the best the evidence allows, and they are exploratory. What would change them, roughly in order of effect:

- **Surface on 169,343 km.** This is the single largest uncertainty. It moves the national RAI anywhere between 28.7 % and 93.5 %, and it decides whether the 5,013 closed links are the whole story. A surface survey of district and community roads, or the agencies' road inventories, would collapse the bound.
- **Condition.** It is unknown on every link. If the agency statistics are right, most district and community roads are poor, and many would be slower or impassable in the rains. The service-access figures are therefore optimistic on the local networks.
- **Speeds.** The speed profiles are generic TransportModelling defaults, not Ugandan. Travel times on earth roads, where most of the remote population lives, depend directly on them.
- **Crossings.** 1,958 river crossings have no recorded structure. They are treated as passable. If many are fords, more links would close.
- **Destination registers.** Market access depends on the 814 registered crop markets. Health and school access depend on the facility registers of the snapshot. A missing facility shows as remoteness.
- **Season rules and the ferry norm** are assumed. Rules on flood exposure use the JRC 100-year depth at the link midpoint, a coarse test.
- **Population** is WorldPop 2020 raked to the UBOS district totals and carried forward at an assumed 3 % a year.
- **Validation.** The design asks for the health access to be checked against SDI's `health_accessibility` layer and for the RAI to reproduce the published figure within 5 points. Neither check passes yet: the RAI plan figure is 24.4 points below the benchmark, and the reason (surface data) is reported. The health cross-check is open.

7.9 Implications for planning

- **The national RAI should be quoted as a range, 28.7 % to 93.5 %, until surface is surveyed.** A single figure would hide that two-thirds of the network is unrecorded.
- **A surface and condition survey of district and community roads is the first investment the plan needs.** It changes the access baseline more than any feasible amount of road works in one budget cycle.
- **Wet-season closures are concentrated and fixable.** 3,299 km close, clustered in the Semliki plain, West Nile and Karamoja; spot improvements at the 2,248 mapped crossings are the cheapest way to turn them into all-season links (Chapters 8 and 10).
- **The north and north-east are behind on every class.** Karamoja and the Acholi and Lango sub-regions form the low–low cluster of the RAI and the largest red areas for EmONC, secondary school, market and district-centre access.
- **Remote is also poor.** Median travel times in the poorest fifth of settlement cells are three to six times those in the richest fifth, so access investments that target remoteness also target poverty.

- **Service access is often a sector question.** Health and school access is good for most people, and where it is not, the cause is often the absence of a facility (Chapter 8). The road programme should be read beside the Health and Education engines' facility plans.
- **Condition data must come from the agencies.** The statistics show 85.7 % of district and 93.7 % of community roads in poor condition, but not where. Access to the UNRA and DUCAR road asset management data is a precondition for a condition programme that can be mapped and sequenced.

CHAPTER 8

8 Gaps: seven dimensions, never combined

This chapter answers the third planning question: *where are the gaps?* It puts the requirement of Chapters 5 and 6 against the network of Chapter 7 and reports, district by district, how far the network falls short. It implements the TI-05 module, design chapter 7 (`docs/07_TI05_GAPS.md`), with the equations of `docs/23_MATHEMATICAL_MODELS.md`. The result is a **gap register** with seven dimensions, each in its own unit, each with its own maps and rankings. They are never added, weighted or averaged into one score. All figures come from the reference run `014c9cd7`, which is exploratory: the norms are drafts, surface is unknown on 68.3 % of network km, and condition is unknown on every link.

8.1 What the module does

TI-05 compares three things for every settlement, link and sector node: what is required (TI-02 and TI-03), what exists (TI-04) and what the difference means. It writes:

- a **settlement-level access gap** (G1) for each of the seven destination classes, with the number of people unserved, how far beyond the standard they are, and *why* they are unserved;
- **link-level gaps** for condition (G2), capacity (G3), standard (G4), connectivity (G5) and resilience (G7), each with its own measure and unit;
- a **node-level gap** for sector connections (G6): which hospitals, industrial parks, border posts and farm zones lack the road their sector needs;
- a **structure-survey list**: river crossings with no recorded structure, and gap links whose surface is unknown;
- the **district gap matrix** (`tables/ti05_gap_matrix.csv`), one row per district and one block of columns per dimension. The table has no "overall" column, and a schema test fails the run if a column is named or computed as a combination of dimensions.

The gaps feed TI-06 (costing, Chapter 9) and TI-07 (the investment programme, Chapter 10), where one intervention may close several gaps and is costed once.

8.2 Why the seven gaps are never added

The seven dimensions answer different questions, in different units, and are closed by different instruments (Table 8.1).

Table 8.1: The seven gap dimensions, their units and the instruments that close them

Gap	Question	Unit	Closed by
G1 Access	Who cannot reach each destination class within its standard?	persons, by class	upgrading, spot improvements, new links; or a new facility (sector)
G2 Condition	Which roads are below their target condition?	km poor	maintenance and rehabilitation
G3 Capacity	Where does traffic exceed capacity?	links over the volume/capacity threshold	capacity expansion
G4 Standard	Where is the existing standard below the required one?	km by field (surface, all-season)	upgrading, spot improvements

Gap	Question	Unit	Closed by
G5 Connectivity	Where is the network fragile or missing?	cut edges; persons behind them; unrecorded crossings	new links, crossings, structure surveys
G6 Sector connection	Which sector nodes lack their required connection?	nodes and output at stake, by sector	targeted upgrades of connectors
G7 Resilience	Which critical links are exposed to hazard?	persons behind exposed critical links	climate adaptation, alternative routes

A person, a kilometre and a link over capacity have no common unit. Converting them to one would need weights, and the weights would decide the answer. A district can have a small access gap and a large capacity gap: Wakiso has the most people beyond 2 km of an all-season road (1.71 million) *and* the most links over capacity (1,014), while Ntoroko has one of the smallest numbers of people beyond the RAI standard (59,000) but the largest wet-season closure (463 km) and the largest flood-resilience gap. A composite score would place both somewhere in the middle and hide what each needs. The engine therefore ranks districts on one column at a time, which the planner chooses, and every map shows one dimension. Overlaps are real: a washed-out drift is an access gap, a standard gap and possibly a resilience gap at once. They are resolved in TI-07, which costs the one intervention that closes all three, not here by weighting.

The sign convention is the same in every dimension: positive is a shortfall, negative a surplus.

8.3 Method

8.3.1 G1 access and its depth

For settlement *i* and destination class *k*, the unserved population is the requiring population of the class's target group (all people, women aged 15–49, children 6–12, youth 13–18 or households) times the indicator of failing the standard, from Equation Equation (7.5):

$$U_{i,k} = P_i^{g(k)}\,(1 - A_{i,k}), \qquad \mathrm{cov}_{k,d} = 1 - \frac{\sum_{i \in I_d} U_{i,k}}{\sum_{i \in I_d} P_i^{g(k)}} \tag{8.1}$$

The depth of the gap is how far beyond the standard the settlement lies, in minutes or km:

$$\mathrm{depth}_{i,k} = \max(0,\ tt_{i,k} - \tau_{k,u(i)}) \tag{8.2}$$

A settlement five minutes over and one three hours over are both unserved, but they need different responses, so the engine keeps both numbers.

8.3.2 G1 cause decomposition

Every unserved person-class pair is given exactly one cause, by recomputing access on two altered networks:

1. **Network deteriorated** if the standard is met once every link is restored to good condition.
2. **Network sub-standard** otherwise, if the standard is met on the *required-state* network: every link raised to its required surface and every seasonal crossing fixed.
3. **No destination** otherwise: even the required network does not bring the destination within the standard.

The causes add up to the unserved total for every settlement and class:

$$\sum_c U_{i,k}^c = U_{i,k} \tag{8.3}$$

The design has a fourth cause, *network missing*, for people who would be served only through a link that does not yet exist. It needs candidate alignments for new roads from SDI or TransportModelling, and none were supplied. In the reference run, **"no destination" therefore also contains the people a missing link would serve**; the run log records this as a fallback. And because no link has a known condition, restoring condition changes nothing, so no person is assigned to *network deteriorated*.

8.3.3 G2 to G7

- **G2 condition** is the km below the target condition (fair or better). At link level it would be the sum of lengths of poor or bad links, on the base-year and do-nothing paths. With no link condition, it is reported **by network class** from the agency statistics of Chapter 7, never spread onto links.
- **G3 capacity** is the ratio of required-state design traffic, in passenger-car units, to daily capacity, on links whose capacity is known:

$$vc_l = \frac{PCU_l}{Q_l^{day}}, \qquad \mathrm{G3}: \; vc_l > \bar{v} = 0.8 \tag{8.4}$$

- **G4 standard** compares the required standard *S** of TI-03 with the existing standard *S*, field by field. For surface, with fields ranked from earth to asphalt (unknown surface ranks as earth):

$$\Delta_l = \mathrm{rank}(S_l^*) - \mathrm{rank}(S_l), \qquad km^{G4} = \sum_l \mathrm{len}_l \, \mathbf{1}(\Delta_l > 0) \tag{8.5}$$

A second G4 measure counts the km of links that must be all-season but are closed in the wet season.

- **G5 connectivity** finds *cut edges*: links whose removal splits the wet-season network in two (Tarjan's bridge-finding algorithm). For each cut edge *l*, the population isolated is that of the smaller side, where *C(l)* is the side not containing the larger part of the component *C*:

$$\mathrm{iso}_l = \min(P(C_l), \; P(C) - P(C_l)) \tag{8.6}$$

Cut edges nest along pendant chains, so their isolated populations are **never summed**. The matrix reports, per district, the number of cut edges isolating at least 2,000 people (`nc_min_pop`) and the largest population behind a single cut.

- **G6 sector connection** tests whether the shortest wet-season path from each sector node to its target (a road of the required class, or a market) meets the node's connection standard in surface and axle load. A link of unknown surface cannot be shown to meet a gravel or paved standard.
- **G7 resilience** flags cut edges that are flood-exposed and isolate at least 2,000 people:

$$RG_l = \mathbf{1}(\mathrm{flood}_l) \cdot \mathbf{1}(\mathrm{iso}_l \geq 2000) \tag{8.7}$$

Exposure and criticality are reported side by side, never multiplied into a risk score.

The bands on the maps are the run's own quintiles, per dimension, always shown with the underlying number.

8.4 Results: the gap register

Table 8.2 is the national gap register: one row per measure, in its native unit, with the district where it is largest.

Table 8.2: The gap register: every dimension in its own unit (never combined)

Dimension	Measure	Unit	National	Largest district	Its value	Districts > 0
G1 Access	rural people beyond 2 km of an all-season road	persons	33,320,032	Wakiso	1,709,126	135
G1 Access	people unserved: primary health	persons	3,208,741	Kotido	227,078	135

Dimension	Measure	Unit	National	Largest district	Its value	Districts > 0
G1 Access	people unserved: EmONC	persons	1,035,094	Kotido	54,928	135
G1 Access	people unserved: primary school	persons	2,944,750	Wakiso	86,199	135
G1 Access	people unserved: secondary school	persons	860,840	Alebtong	33,976	135
G1 Access	people unserved: market	persons	2,332,833	Mubende	141,370	135
G1 Access	people unserved: administrative centre	persons	4,223,853	Kasese	254,801	135
G2 Condition	km poor or bad (link level)	km	0	-	0	0
G3 Capacity	links over volume/capacity threshold	links	3,795	Wakiso	1,014	42
G4 Standard	km below the required surface	km	242,068	Wakiso	9,055	135
G4 Standard	km not all-season	km	3,299	Ntoroko	463	123
G5 Connectivity	cut edges isolating ≥ the minimum population	links	3,738	Wakiso	438	135
G5 Connectivity	largest population isolated by one cut	persons	716,686	Sironko	716,686	135
G7 Resilience	population behind flood-exposed critical links	persons	317,140	Ntoroko	298,306	3

G5 'largest population isolated' is the national maximum, not a sum. G6 (sector nodes) is in its own table. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_gap_register.csv).

The register is read row by row, never down a column. The sections below take each dimension in turn.

8.5 G1 Access: who is unserved, and why

8.5.1 The scale of the access gap by class

Unserved numbers differ by an order of magnitude between classes:

- **All-season road within 2 km:** 33.32 million people, 64.1 % of the 52.0 million who require it.
- **District centre:** 4.22 million (8.1 %).
- **Primary health:** 3.21 million (6.2 %).
- **Primary school:** 2.94 million children of primary age (27.0 % of 10.9 million).
- **Market:** 2.33 million (20.2 % of 11.6 million).
- **EmONC:** 1.04 million women of reproductive age (8.5 % of 12.2 million).
- **Secondary school:** 0.86 million young people (11.0 % of 7.8 million).

8.5.2 Why people are unserved

Figure 8.1 and Table 8.3 decompose each class by cause. This decomposition is the most useful single output of TI-05, because it separates the gaps a road can close from those it cannot.

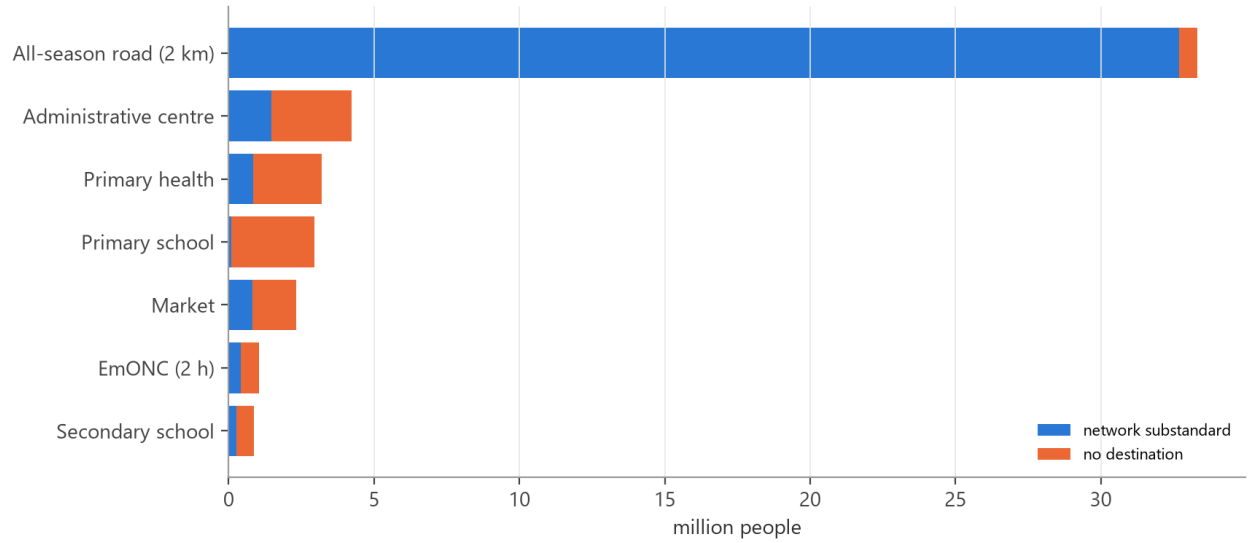

Figure 8.1: Why people are unserved: the G1 access gap decomposed by cause

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 8.3: Access gap by cause (millions of people)

Destination class	network_substandard	no_destination
Administrative centre	1.48	2.74
All-season road within 2 km	32.70	0.62
EmONC facility (2 h)	0.42	0.61
Primary health facility	0.85	2.36
Market	0.82	1.52
Primary school	0.08	2.86
Secondary school	0.27	0.59

Causes are TI-05's decomposition: no destination within reach, network missing, or the network is there but below standard. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_gap_cause.csv).

The two kinds of class behave differently.

- **The all-season road gap is almost entirely a road gap.** Of 33.32 million people beyond 2 km of an all-season road, 32.70 million (98.1 %) would be served if the existing links were raised to their required standard and their crossings fixed. Only 0.62 million would still be beyond 2 km; for them the nearest road of any kind is too far, and only a new link would help. This gap rests heavily on the unknown surface: most of the 32.70 million live along roads whose surface is not recorded and is therefore counted as earth.
- **For services, most of the gap is a sector gap.** Upgrading every road to its required standard would still leave 73.7 % of the primary-health gap (2.36 of 3.21 million), 97.3 % of the primary-school gap (2.86 of 2.94 million), 65.0 % of the market gap, 68.7 % of the secondary-school gap, 65.0 % of the district-centre gap and 59.4 % of the EmONC gap (0.61 of 1.04 million). These people are too far from any facility of that class for road works alone to bring them within the standard. Part of this group would be helped by a *new* road, which the run could not test; the rest needs a new clinic, school or market.

The road-closable parts are still large: 0.85 million people for primary health, 0.42 million women for EmONC, 0.82 million for markets, 0.27 million for secondary school and 1.48 million for the district centre. These are the people that TI-07's upgrades and spot improvements can reach.

The "no destination" findings belong to the sector engines. They are emitted to the Health, Education and Agriculture models (Chapter 22) as facility-siting signals, because building roads to reach them would spend the transport budget on a problem it cannot solve.

8.5.3 Depth: how far beyond the standard

The depth of the gap varies by class. Among people beyond the standard with a wet-season route, the population-weighted median excess is 17 minutes for primary health, 19 minutes for secondary school, 31 minutes for the district centre, 33 minutes for EmONC, 43 minutes for markets, 1.2 km for primary school and 0.6 km for the all-season road. A large share of the unserved have no wet-season route at all: about half of the primary-health, EmONC and district-centre gaps, 36 % of the secondary-school gap and 20 % of the market gap. They are behind a closed crossing, on a flooded road, or in a network component with no destination of that class. For them the depth is infinite, and the response starts with the crossing or the missing link.

8.5.4 All-season road (RAI)

Figure 8.2 maps the 33.32 million rural people beyond 2 km of an all-season road, and Figure 8.3 ranks the 20 largest districts on that one column.

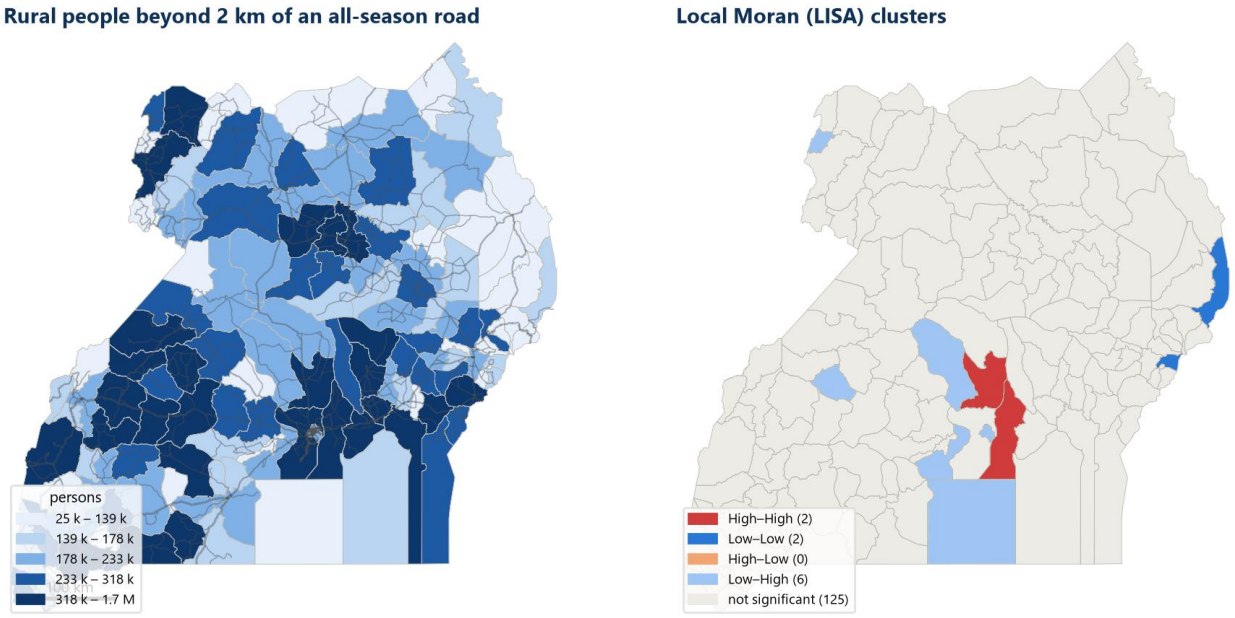

Figure 8.2: G1 Access: rural people beyond 2 km of an all-season road — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

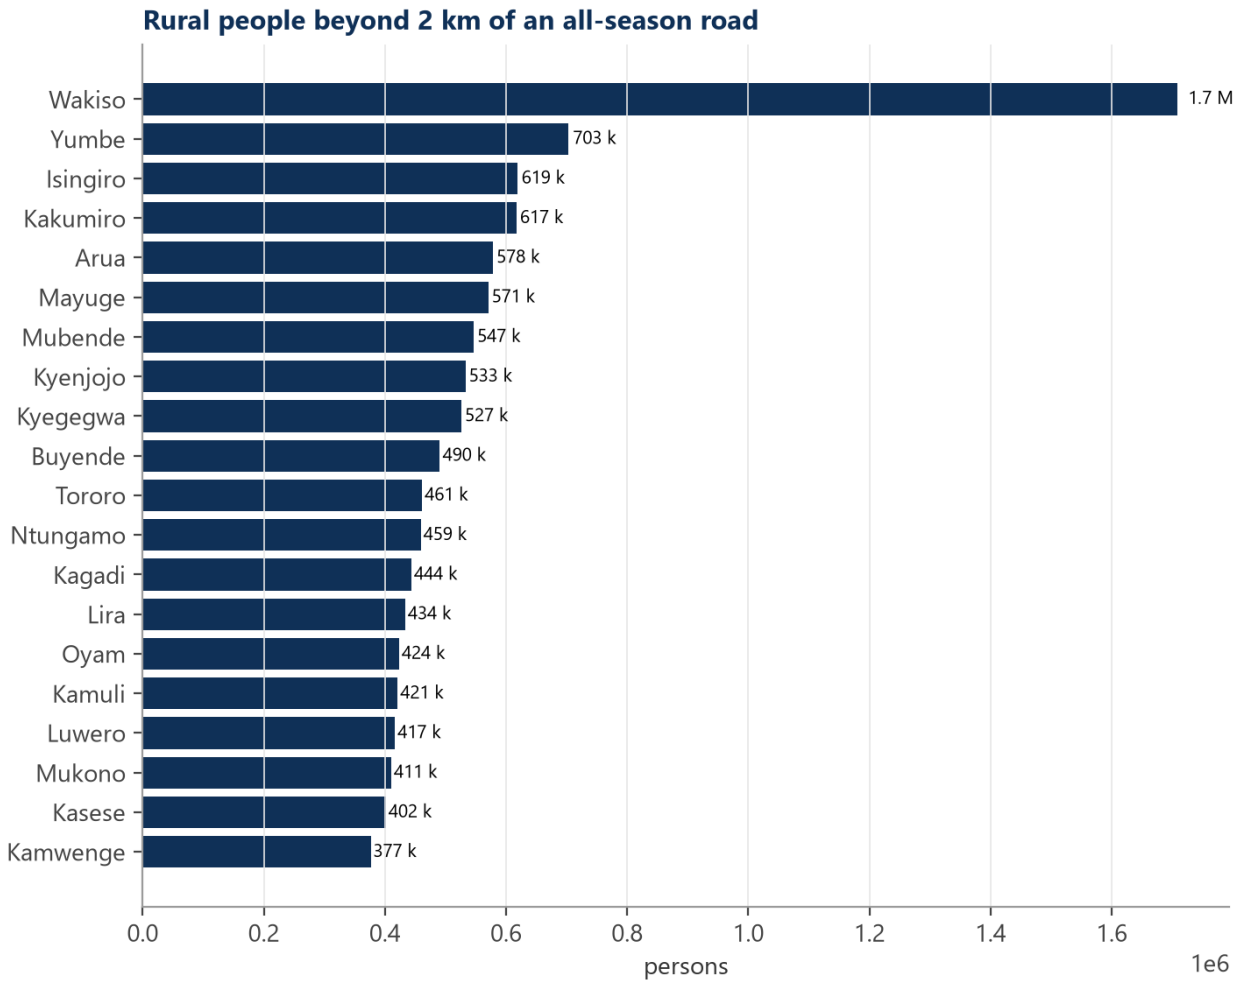

Figure 8.3: Top 20 districts: Rural people beyond 2 km of an all-season road (one dimension, one sort column)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

In absolute numbers the gap follows population. Wakiso has the largest (1.71 million), because it is the most populous district and most of its dense local network has no recorded surface. It is followed by Yumbe (0.70 million), Isingiro (0.62 million), Kakumiro (0.62 million), Arua (0.58 million), Mayuge (0.57 million), Mubende (0.55 million), Kyenjojo (0.53 million), Kyegegwa (0.53 million) and Buyende (0.49 million). The median district has 207,000 people beyond the standard; the smallest gaps are in Obongi (25,000), Kapchorwa (34,000), Manafwa (38,000), Ntoroko (59,000) and Butambala (60,000). By region the gap is spread evenly: 9.1 million in the Western region, 8.4 million in the East, 8.1 million in the North and 7.7 million in the Central region.

Spatial clustering is weak for this column (Moran's I = 0.092, p = 0.033). The only high–high cluster is Luwero and Mukono beside Wakiso; Amudat and Bududa form a low–low pair. The ranking by persons is the right one for "how many people", but the RAI share in Chapter 7 is the right one for "how badly served": Karamoja has few people beyond the standard because few people live there, not because access is good.

8.5.5 Primary health

Figure 8.4 maps the 3.21 million people beyond the primary-health standard.

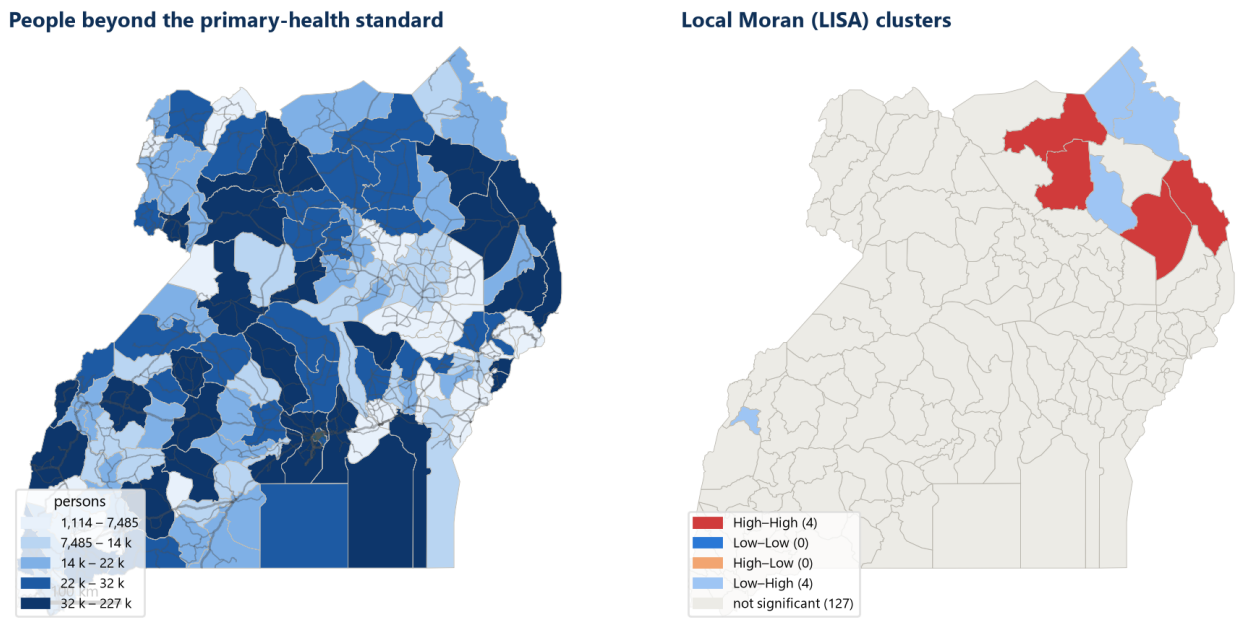

Figure 8.4: G1 Access: people beyond the primary-health standard — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Kotido has the largest gap (227,000), followed by Kasese (106,000), Buvuma (77,000), Amudat (74,000), Bududa (71,000), Mubende (69,000), Napak (66,000), Kakumiro (65,000), Nwoya (65,000) and Wakiso (57,000). The Northern region holds 1.21 million of the 3.21 million. The LISA panel shows one high–high cluster in Karamoja and the north-east: Agago, Kitgum, Moroto and Napak. Abim, Bunyangabu, Kaabong and Karenga are low–high outliers next to it.

8.5.6 EmONC

Figure 8.5 maps the 1.04 million women of reproductive age beyond two hours of EmONC, and Figure 8.6 ranks the top 20 districts.

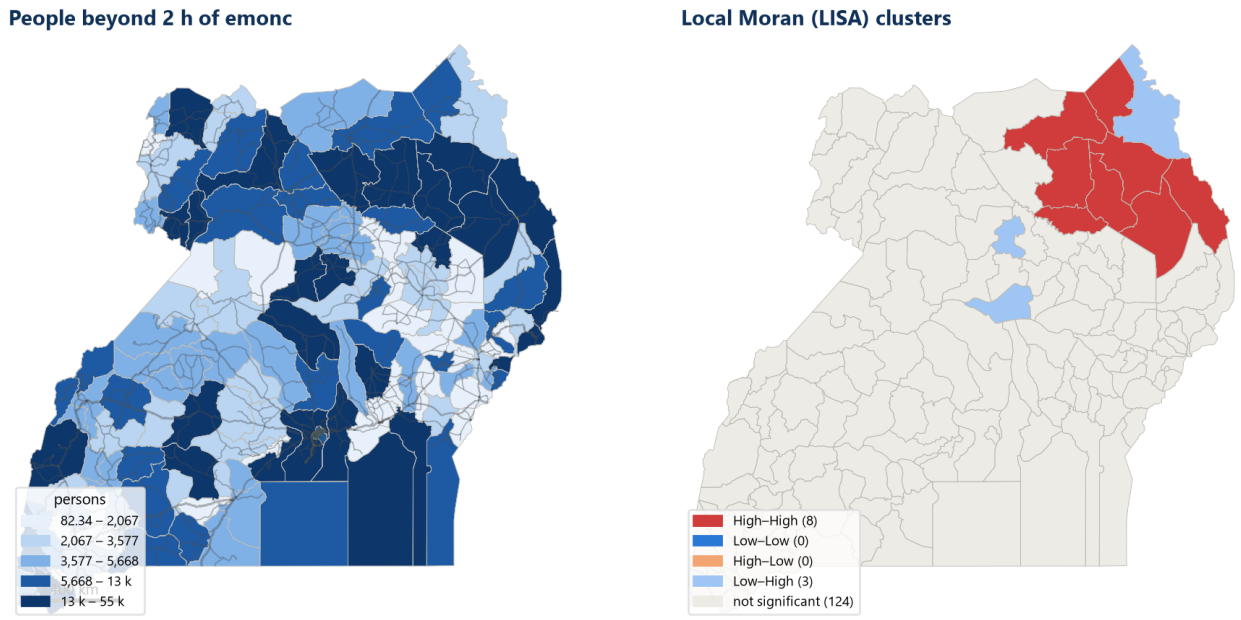

Figure 8.5: G1 Access: people beyond 2 h of EmONC — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

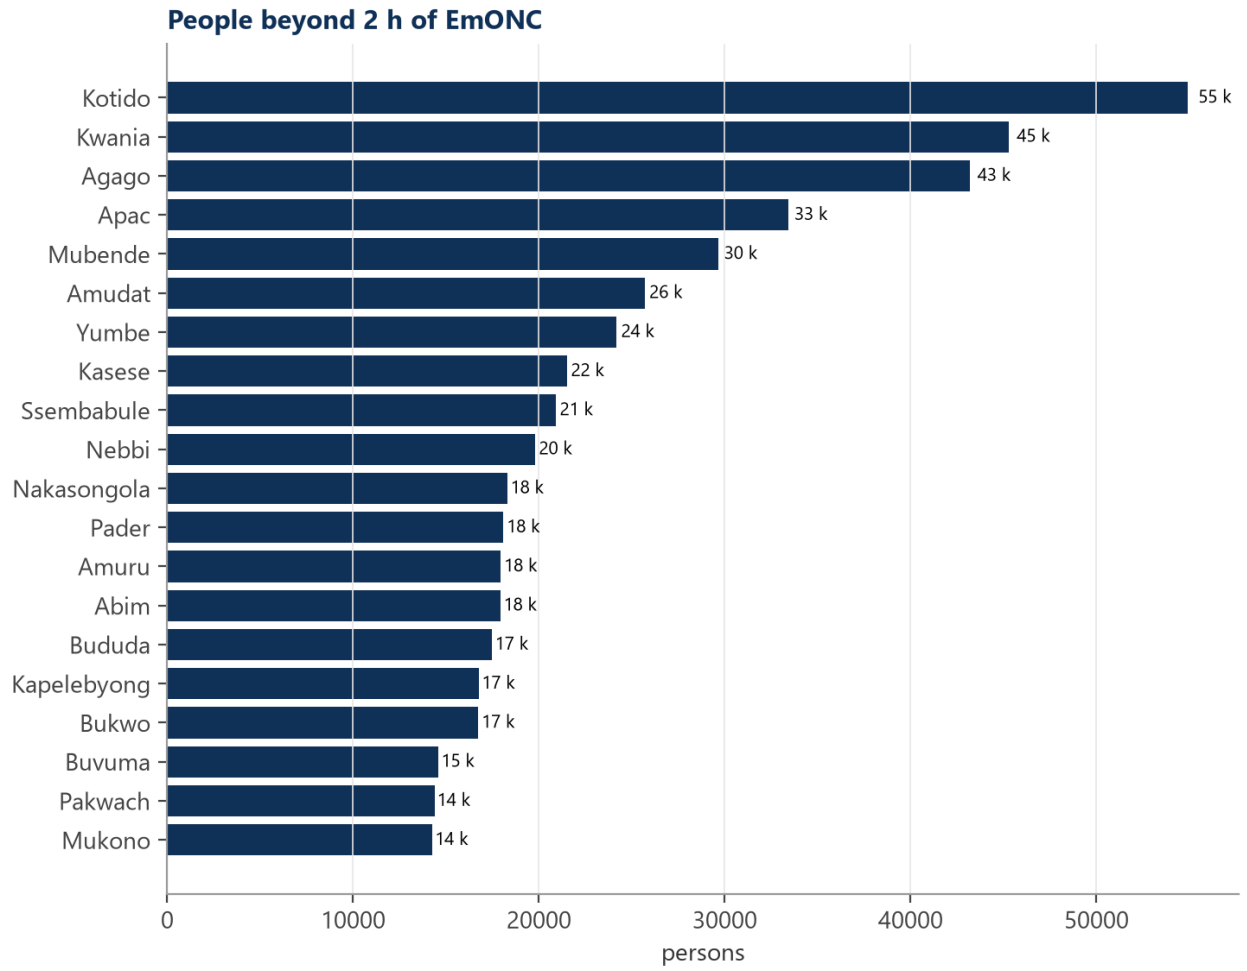

Figure 8.6: Top 20 districts: People beyond 2 h of EmONC (one dimension, one sort column)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The largest gaps are in Kotido (54,900), Kwania (45,300), Agago (43,200), Apac (33,400), Mubende (29,700), Amudat (25,700), Yumbe (24,200), Kasese (21,500), Ssembabule (20,900) and Nebbi (19,800). The Northern region has 465,000, close to half the national total. Clustering is significant (Moran's $I = 0.159$, $p = 0.004$): an eight-district high–high cluster covers Abim, Agago, Karenga, Kitgum, Kotido, Moroto, Napak and Otuke. Amolatar, Kaabong and Kole are low–high outliers. Of the national gap, 0.42 million women could be brought within two hours by road works on existing links; the other 0.61 million need a closer EmONC facility, or a new link the run could not test.

8.5.7 Primary school

Figure 8.7 maps the 2.94 million children of primary age beyond 3 km (rural) of a school by network walking distance.

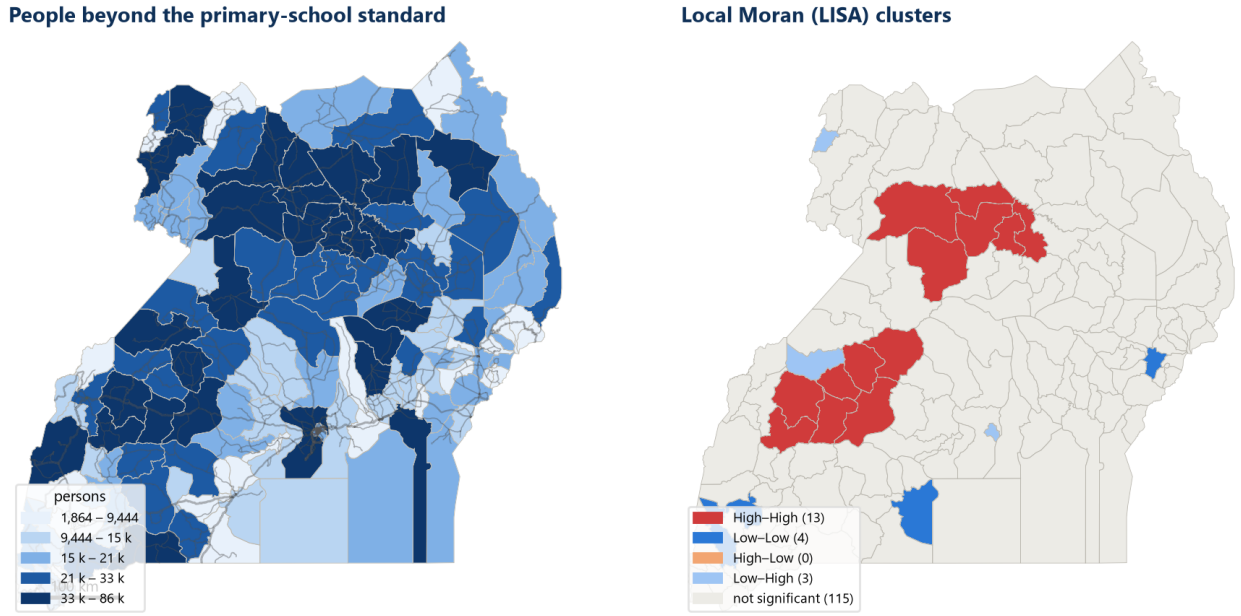

Figure 8.7: G1 Access: people beyond the primary-school standard — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Wakiso (86,200), Kakumiro (71,900), Kyegegwa (66,500), Yumbe (64,800), Mubende (61,200), Isingiro (56,400), Oyam (55,600), Kyenjojo (51,800), Buyende (49,200) and Nwoya (49,200) have the largest gaps. Two high–high clusters appear: a western block of Kakumiro, Kamwenge, Kibaale, Kyankwanzi, Kyegegwa, Kyenjojo and Mubende, and a northern block of Kiryandongo, Kole, Lira, Nwoya, Omoro and Oyam. Bushenyi, Masaka, Mbale and Rukungiri form low–low clusters. Since 97.3 % of this gap has no destination within reach even on an upgraded network, it is overwhelmingly a school-siting question for the Education engine, not a road question.

8.5.8 Secondary school

Figure 8.8 maps the 0.86 million young people beyond the secondary-school standard.

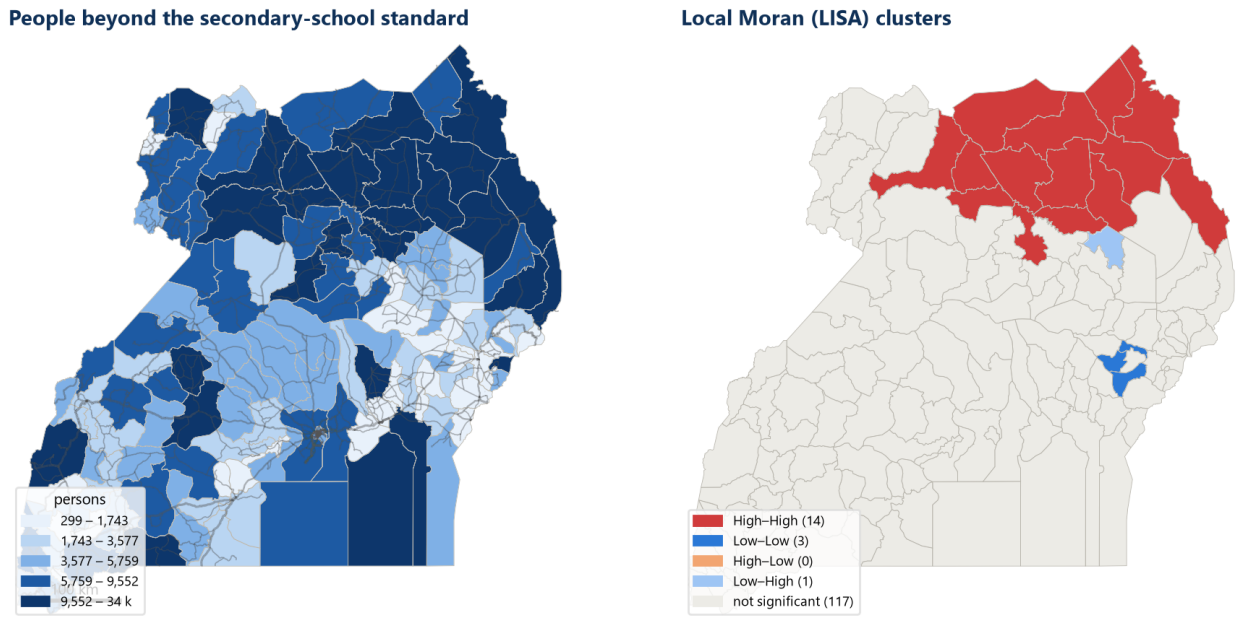

Figure 8.8: G1 Access: people beyond the secondary-school standard — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The gap is the most northern of all classes: 442,000 of the 861,000 are in the Northern region. Alebtong (34,000), Amudat (25,800), Otuke (25,600), Agago (23,200), Kotido (23,200) and Napak (20,600) lead. The high–high cluster is the largest of any G1 layer, 14 districts: Abim, Agago, Amuru, Gulu, Kaabong, Karenga, Kitgum, Kotido, Lamwo, Lira, Moroto, Omoro, Otuke and Pader. Butaleja, Butebo and Kibuku form a low–low cluster in the east. Moran's *I* is 0.291 (p = 0.001).

8.5.9 Market

Figure 8.9 maps the 2.33 million people beyond the market standard.

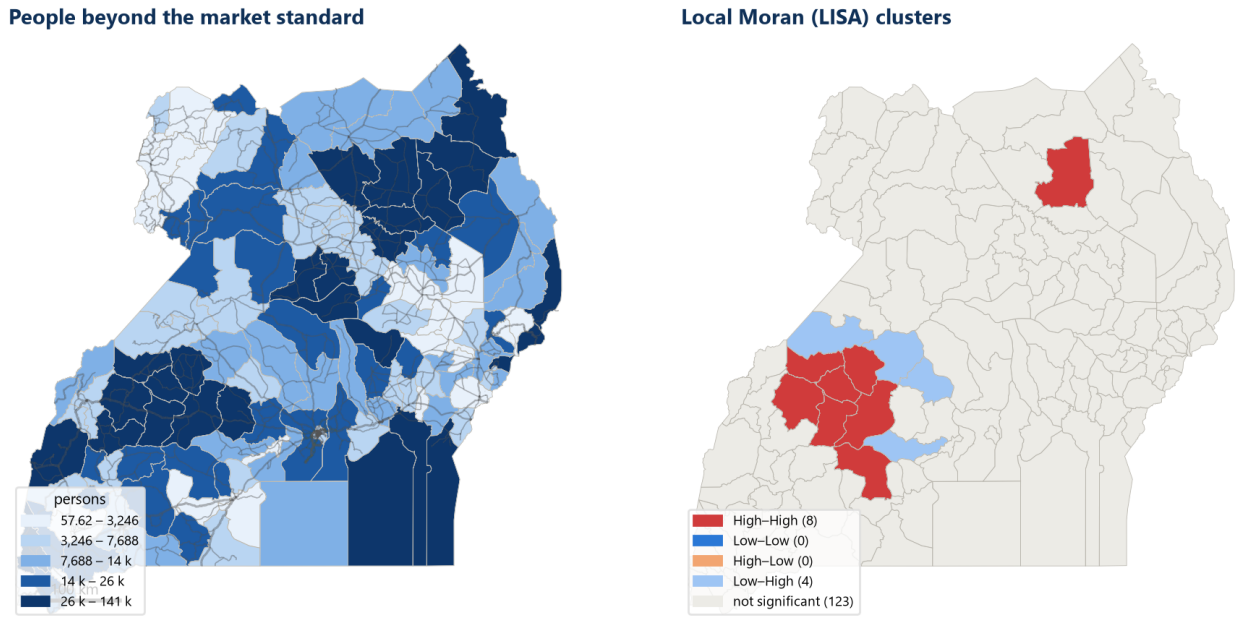

Figure 8.9: G1 Access: people beyond the market standard — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Mubende (141,400) and Kakumiro (140,800) have more than twice the gap of any other district, followed by Agago (67,800), Apac (61,000), Kibaale (56,300), Kagadi (55,500), Kamwenge (53,800), Kotido (51,900), Amolatar (50,300) and Kyegegwa (49,900). The Western (740,000) and Northern (732,000) regions carry most of it. The clustering is the strongest of the G1 layers (Moran's *I* = 0.340): a high–high cluster in the mid-west of Agago, Kagadi, Kakumiro, Kibaale, Kyegegwa, Kyenjojo, Mubende and Ssembabule. As noted in Chapter 7, market access depends on the register of 814 crop markets, so part of this cluster may be a register gap.

8.5.10 District centre

Figure 8.10 maps the 4.22 million people beyond the district-centre standard.

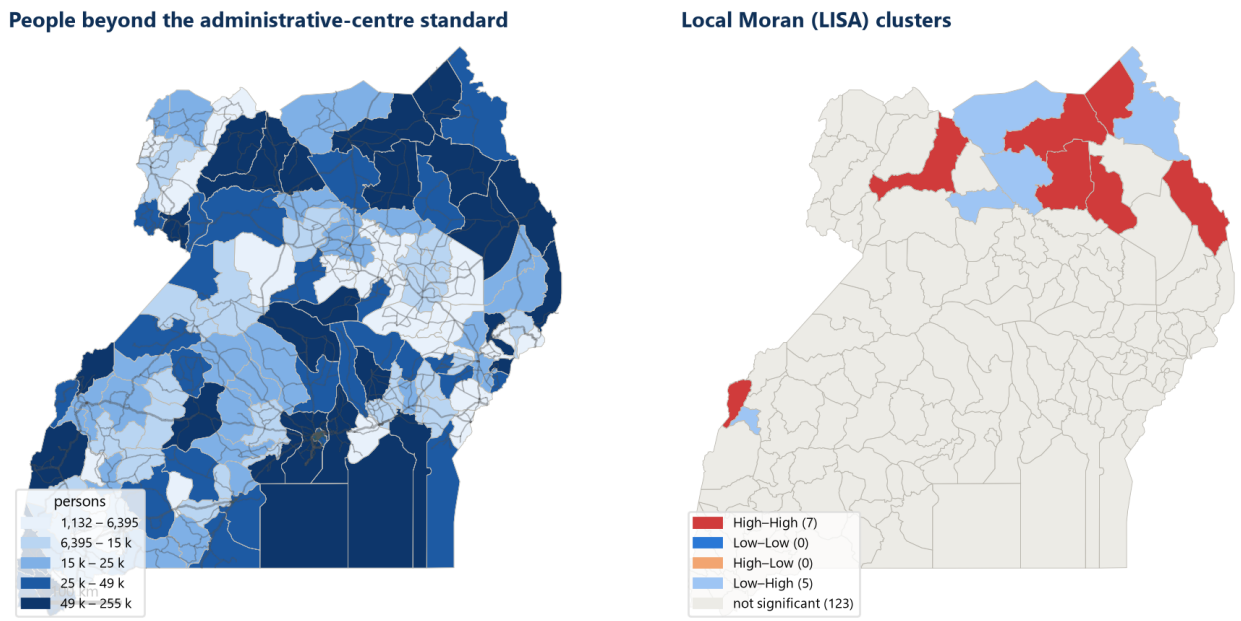

Figure 8.10: G1 Access: people beyond the administrative-centre standard — by district; right: LISA clusters (this dimension only)
Source: AICTEM TIPRAS engine, reference run 014c9cd7

Kasese (254,800), Gulu (239,400) and Kotido (219,100) lead by a wide margin, followed by Mayuge (104,700), Kanungu (98,300), Amudat (96,800), Buvuma (91,400), Moroto (88,000), Mubende (83,600) and Amolatar (82,200). Two-thirds of this gap (2.74 million) is "no destination": the district headquarters is simply too far for the three-hour rural standard, which in a large district no road can fix. The high–high cluster is Abim, Agago, Amuru, Bundibugyo, Karenga, Kitgum and Moroto; global clustering is weak and not significant at 5 % (Moran's I = 0.086, p = 0.051).

8.6 G2 Condition: a gap known only by network

The G2 row of the gap register reads **0 km at link level**, in every district. That zero is not a finding that roads are in good condition. It is the absence of link-level condition data: no link in the snapshot carries an IRI or a visual class, so no link can be counted as poor, and the do-nothing projection has nothing to project.

The condition gap is therefore reported by network class from the agency statistics already shown in Figure 7.13 and Table 7.4 (Chapter 7). In FY2022/23, 139,691 km of public roads were in poor condition:

- 204 km of 6,133 km national paved roads (3.3 %) and 1,657 km of 15,067 km national unpaved roads (11.0 %);
- 35,320 km of 41,194 km district roads (85.7 %) and 70,656 km of 75,404 km community access roads (93.7 %);
- 31,854 km of 35,858 km urban roads (88.8 %), which are reported but not programmed.

The non-urban poor kilometres, 107,837 km, are the G2 gap carried forward. TI-07 turns them into condition programmes in 1,000 km tranches (Chapter 10), TI-06 costs them (Chapter 9), and TI-09 reports the poor km left after each allocation (Chapter 12). Because the statistics carry no geometry, there is no G2 map, no G2 district ranking and no G2 column in the LISA analysis. Any such map would have to invent which roads are poor.

8.7 G3 Capacity

Figure 8.11 maps the links whose required-state design traffic exceeds 0.8 of capacity.

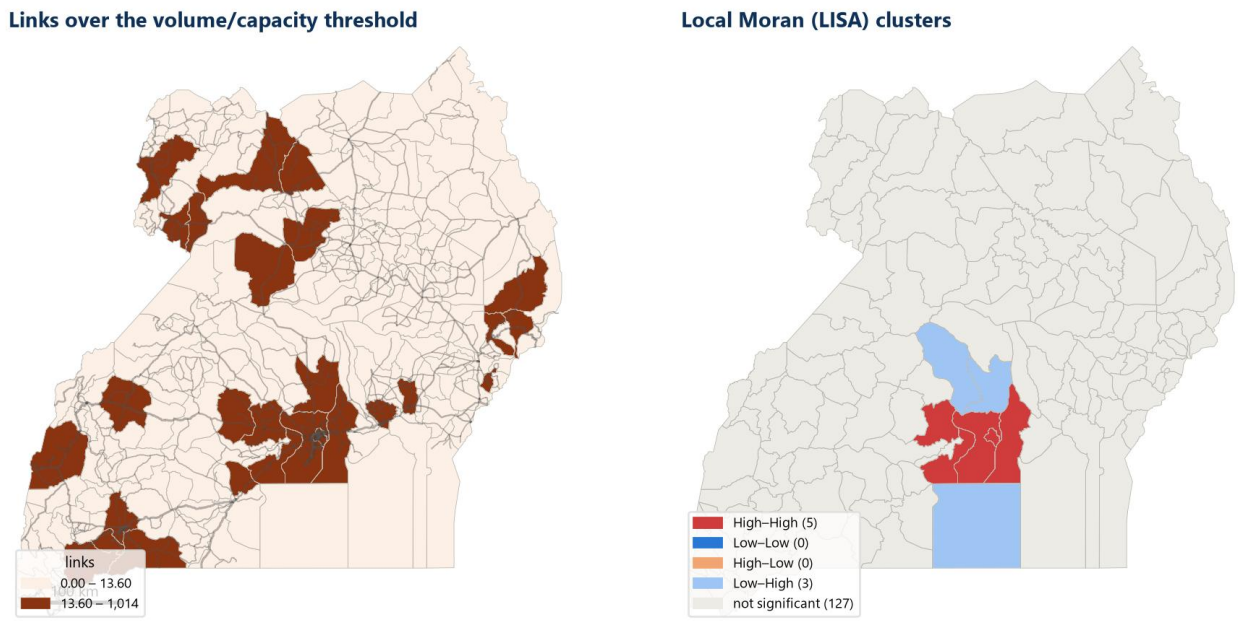

Figure 8.11: G3 Capacity: links over the volume/capacity threshold — by district; right: LISA clusters (this dimension only)
Source: AICTEM TIPRAS engine, reference run 014c9cd7

3,795 links (about 934 km) exceed the threshold, in 42 districts. They are mostly national roads: 2,244 national primary links (624 km), 919 trunk links (210 km) and 504 national secondary links (90 km); only 127 district and community links appear, because lane counts, and therefore capacities, are recorded mainly on the main roads. Of the 3,795 links, 3,372 have a ratio above 1.0 and 3,071 above 1.2; the median ratio is 3.0 and the maximum 47.9.

The gap concentrates around Kampala: Wakiso (1,014 links), Kampala (468), Mukono (402), Mpigi (211) and Mityana (160), which form the high–high LISA cluster, with Amuru (181), Ntungamo (139), Jinja (103), Kiryandongo (97) and Isingiro (92) on the corridors to the north, west and east. G3 is the most clustered layer in the run (Moran's I = 0.393).

Two cautions apply.

- **Capacity is unknown on most of the network.** 276,431 loaded links have no lane count. They are listed as "capacity unknown" and never treated as adequate, so G3 is a lower bound on where capacity binds.
- **The loads are modelled.** Ratios of 10 or more come from the TI-03 gravity and all-or-nothing assignment on the required-state network (Chapter 6), with freight grown at a flagged fallback rate. They locate pressure; they are not traffic counts. Traffic counts from the road agency are the validation they need.

No capacity-expansion works enter the programme in the reference run: the implemented decision rule (Chapter 10) does not yet apply the capacity rule, so G3 is reported for information and validation only.

8.8 G4 Standard

8.8.1 Below the required surface

Figure 8.12 maps the km of links whose existing surface is below the surface TI-03 requires.

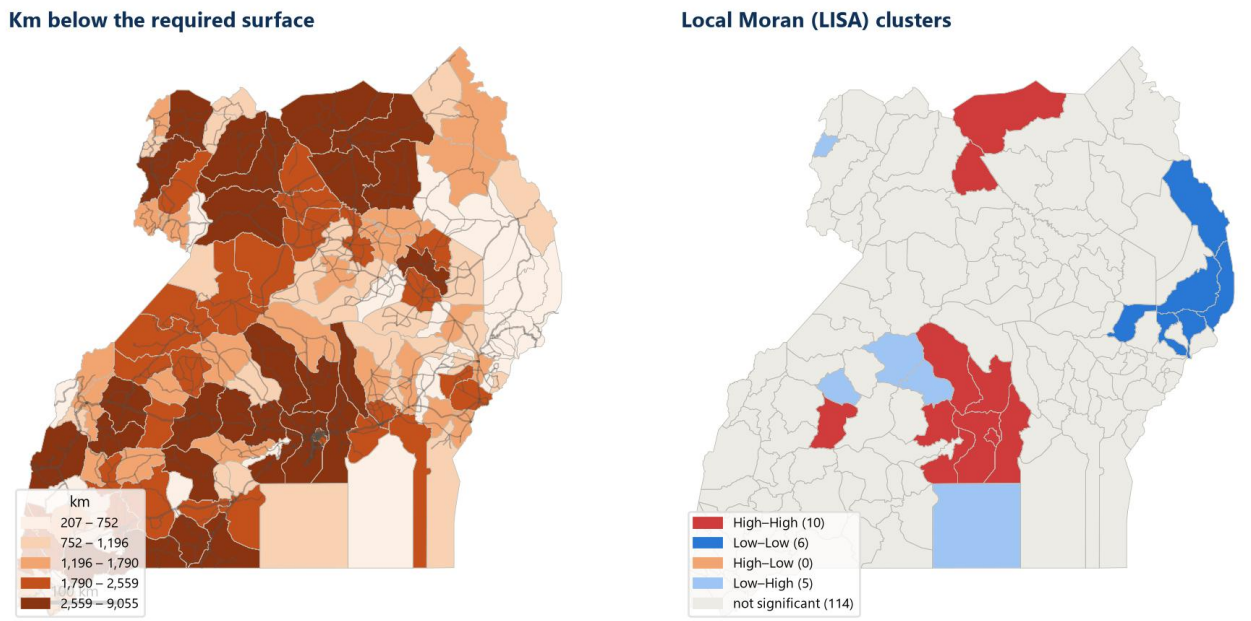

Figure 8.12: G4 Standard: km below the required surface — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

242,068 km are below the required surface. This is 97.6 % of the network, and the number must be read with its composition.

- **72,725 km have a known surface** that is below standard. These links generate upgrading works in TI-07 (Chapter 10).
- **169,343 km have an unknown surface**, which ranks as earth and so falls below any gravel or paved requirement. They are not a works list; they are the survey requirement discussed below.

By binding test (TI-03), 166,633 km are set by the link's function (the class minimum), 75,261 km by traffic and 173 km by a sector node (a mine, farm zone or park that requires a stronger road). By district the km follow network length: Wakiso (9,055 km), Mukono (5,463), Kyenjojo (4,787), Mubende (4,694), Nakaseke (4,503), Yumbe (4,368), Isingiro (4,000), Kasese (3,902), Arua (3,860) and Kyotera (3,668). The high–high cluster is the Kampala region and its northern hinterland: Gulu, Kampala, Kyegegwa, Lamwo, Luwero, Mityana, Mpigi, Mukono, Nakaseke and Wakiso. Karamoja and Elgon form the low–low cluster (Amudat, Bulambuli, Kumi, Kween, Moroto, Nakapiripirit), because their networks are short, not because they meet the standard.

8.8.2 Not passable all season

Figure 8.13 maps the km of links that must be all-season but close in the wet season, and Figure 8.14 ranks the 20 largest districts.

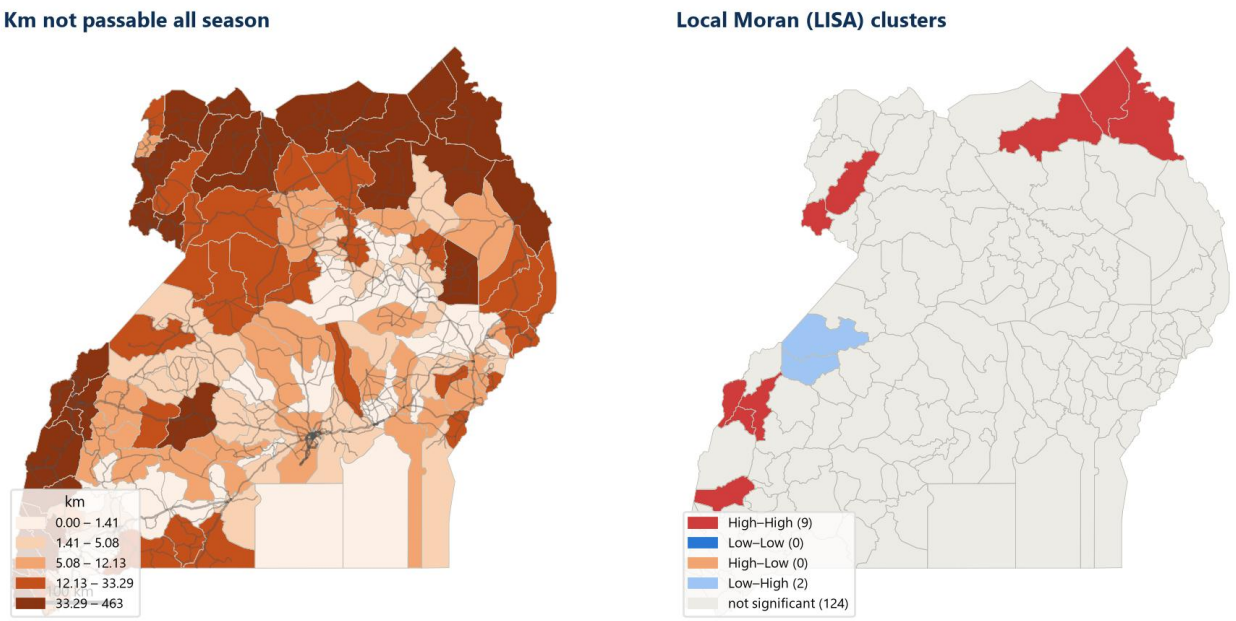

Figure 8.13: G4 Standard: km not passable all season — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

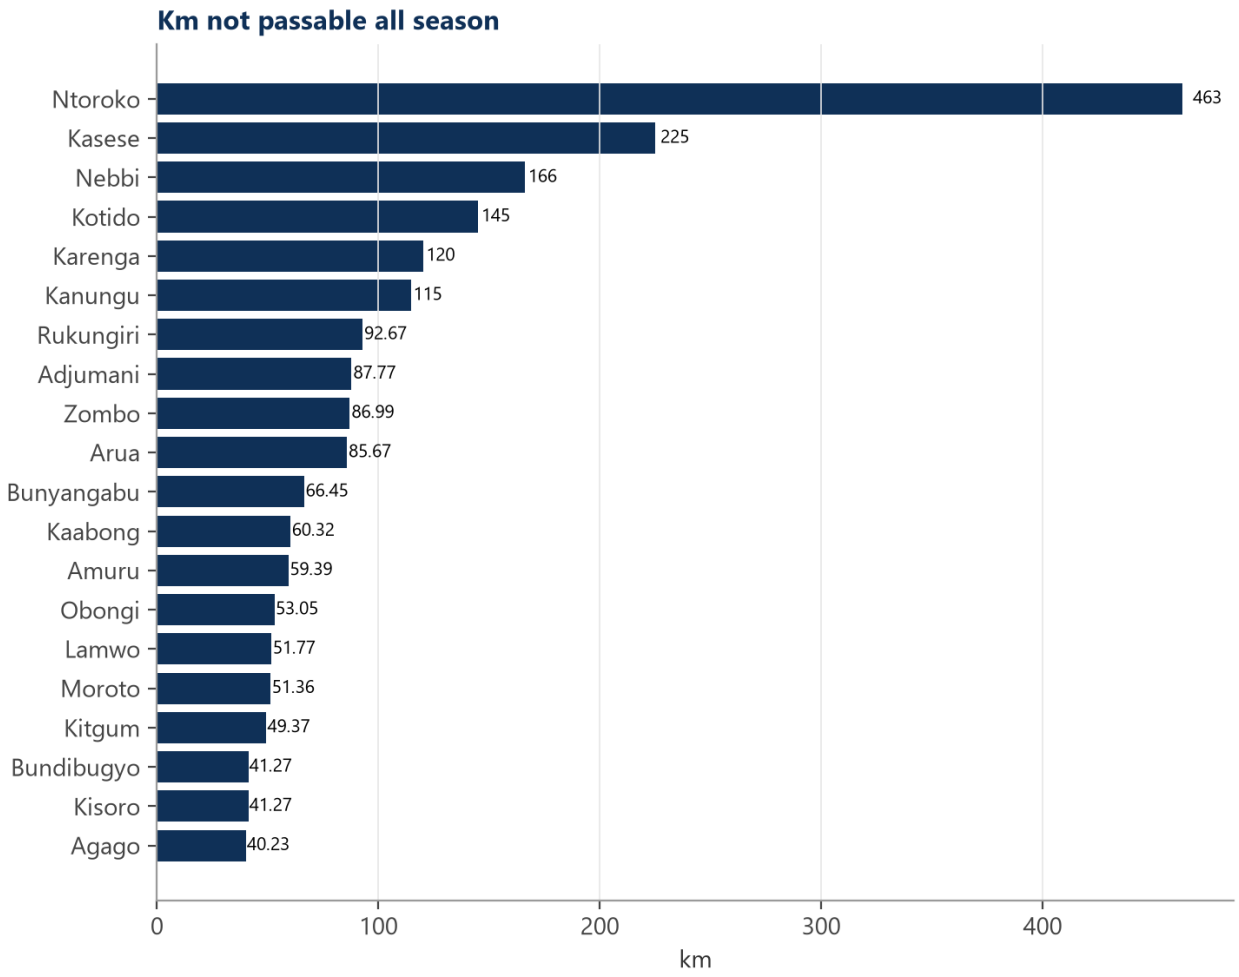

Figure 8.14: Top 20 districts: Km not passable all season (one dimension, one sort column)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

3,299 km on 5,013 links, in 123 districts, are the wet-season closures of Chapter 7 seen as a standard gap. 4,026 of the links are required to be all-season by their function and 987 by their traffic. Ntoroko (463 km) is far ahead,

then Kasese (225), Nebbi (166), Kotido (145), Karenga (120), Kanungu (115), Rukungiri (93), Adjumani (88), Zombo (87) and Arua (86). The Northern (1,417 km) and Western (1,367 km) regions hold 85 % of the total. The LISA high–high cluster joins the Rwenzori districts (Bundibugyo, Bunyangabu, Kabarole, Rubirizi), West Nile (Madi Okollo, Zombo) and the north-east (Kaabong, Karenga, Kitgum).

This is the most actionable standard gap. The closure is usually at one crossing on an otherwise usable link, so the response is a spot improvement, a vented ford or a culvert line, at a fraction of the cost of upgrading the whole link. The 2,248 links with a mapped seasonal crossing are where TI-07's spot-improvement rule applies.

8.9 G5 Connectivity

8.9.1 Cut edges

Figure 8.15 maps, by district, the cut edges whose loss would isolate at least 2,000 people.

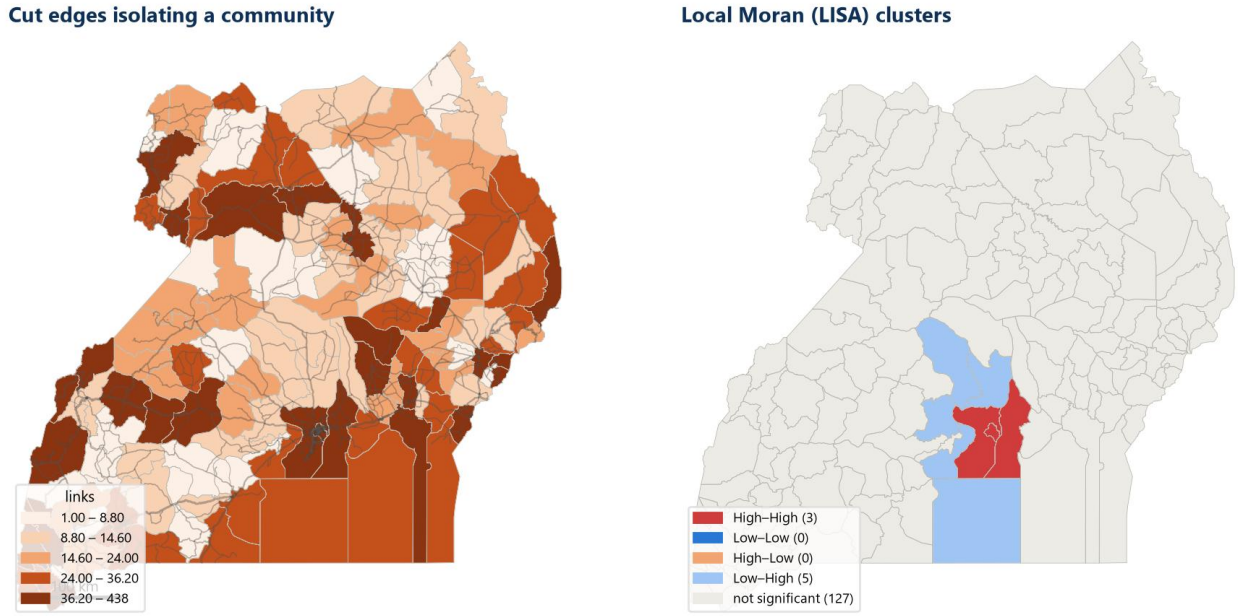

Figure 8.15: G5 Connectivity: cut edges isolating a community — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The wet-season network has **65,548 cut edges with population behind them**. Most isolate a handful of settlement cells. **3,738 isolate at least 2,000 people**, 530 at least 10,000 and 78 at least 100,000. Every district has some. The count is highest in Wakiso (438), Kampala (185), Bududa (113), Kasese (109), Mukono (100), Kanungu (92), Mbale (65), Kyenjojo (59), Ntoroko (57) and Mubende (55). Kampala, Mukono and Wakiso form the high–high cluster; dense urban networks have many short pendant chains, each a cut edge with a few thousand people behind it.

8.9.2 The largest population behind one cut

Figure 8.16 maps the largest population behind a single cut edge in each district, and Figure 8.17 ranks the top 20.

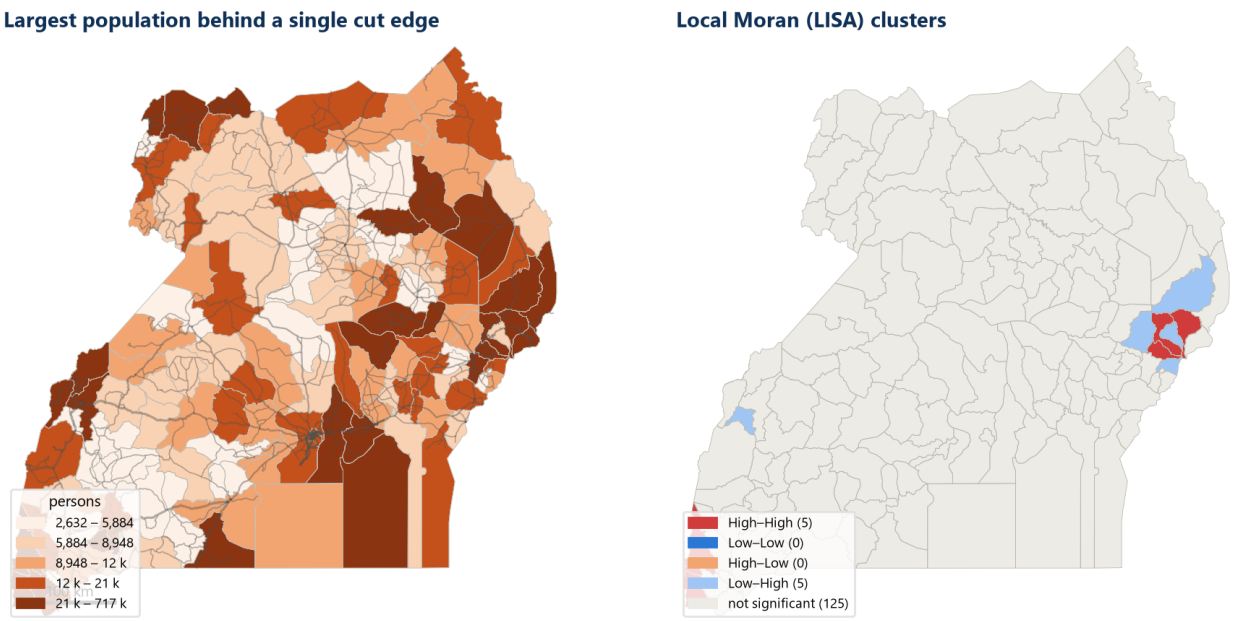

Figure 8.16: G5 Connectivity: largest population behind a single cut edge — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

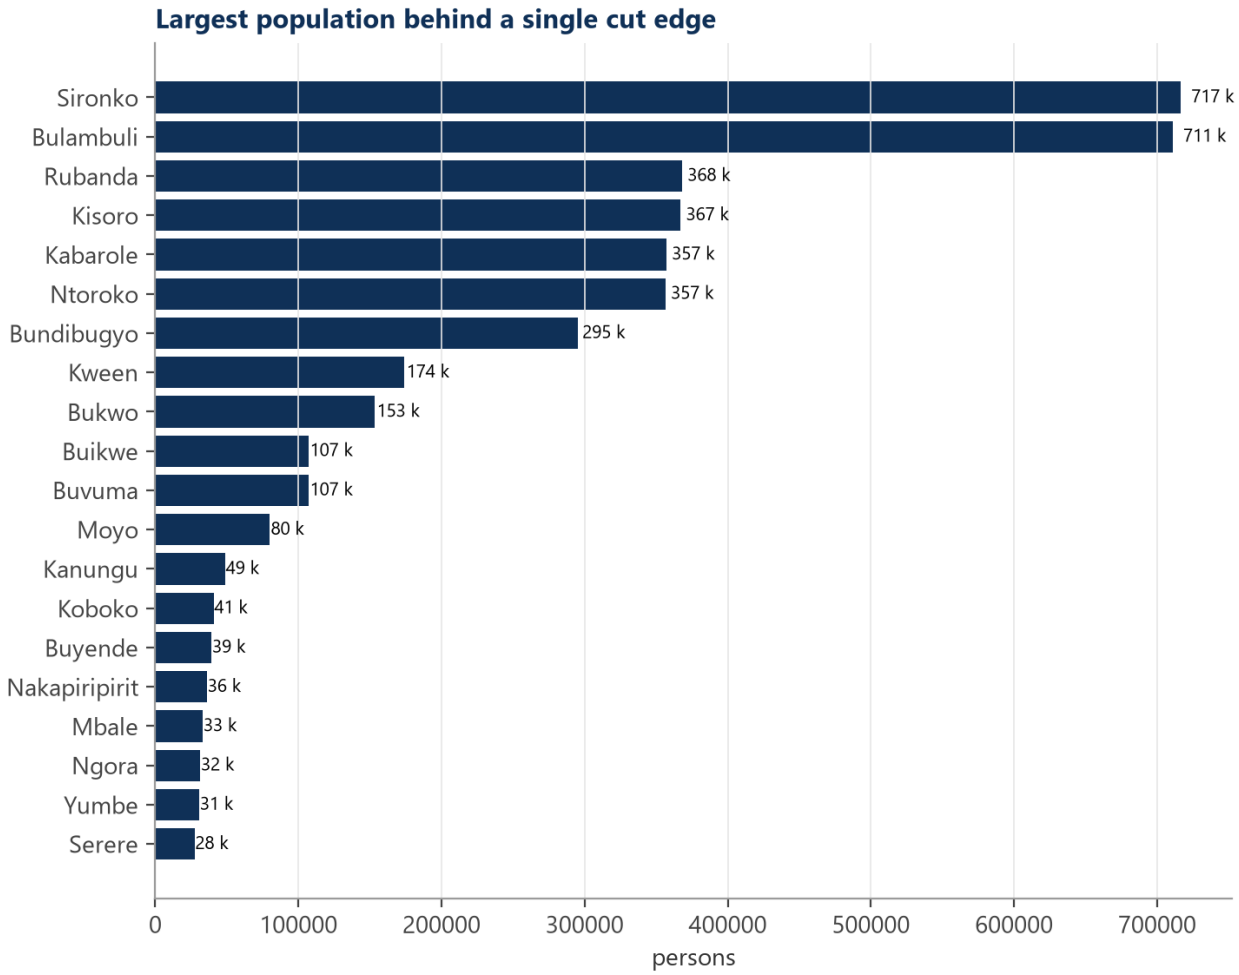

Figure 8.17: Top 20 districts: Largest population behind a single cut edge (one dimension, one sort column)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

This column measures fragility, not frequency. The largest values are on Mount Elgon and in the south-west and Rwenzori highlands: Sironko (716,686), Bulambuli (711,139), Rubanda (368,252), Kisoro (367,166), Kabarole

(357,373), Ntoroko (356,701), Bundibugyo (295,449), Kween (173,795), Bukwo (153,287) and Buikwe (107,012). The population behind a cut is counted on the network, so it can include people in neighbouring districts. Sironko and Bulambuli almost certainly share the same Elgon sub-network: one link in each, if lost in the wet season, would cut off about 0.7 million people from the rest of the network. The high–high cluster is Bulambuli, Kanungu, Kisoro, Kween and Sironko.

These numbers are never added. A pendant chain of ten cut edges has ten nested populations; summing them would count the same people ten times. The register reports the national maximum, 716,686, not a sum.

8.9.3 Unrecorded river crossings

The structure inventory is incomplete. Table 8.4 lists, by road class, the 1,958 river crossings on 1,849 links where a river, stream or canal crosses the road and OSM records no bridge, culvert or ford.

Table 8.4: Structures to survey (G5/G7 evidence gap)

Class	River crossings with unknown structure
NAT_TRUNK	44
NAT_PRIMARY	38
NAT_SECONDARY	91
DISTRICT	156
COMMUNITY	1,090
TRACK	433
SERVICE	22
UNKNOWN	84

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_structure_survey.csv).

Most are on community roads (1,090) and tracks (433), but 173 are on national roads (44 trunk, 38 primary, 91 secondary) and 156 on district roads. They are spread over 116 districts; Zombo (112), Yumbe (106), Kampala (94), Gulu (87), Moroto (77), Mukono (74), Namisindwa (66), Arua (62), Bushenyi (60) and Kaabong (59) have the most. A crossing with no recorded structure is a **survey requirement, not a closure**: it does not change wet-season availability in the run. If a large share turn out to be fords, the closures and cut edges above would grow.

The third G5 sub-measure, *missing links*, cannot be computed: it needs candidate alignments, which have not been supplied.

8.10 G6 Sector connection

Figure 8.18 and Table 8.5 report the sector nodes that lack the connection standard their sector requires.

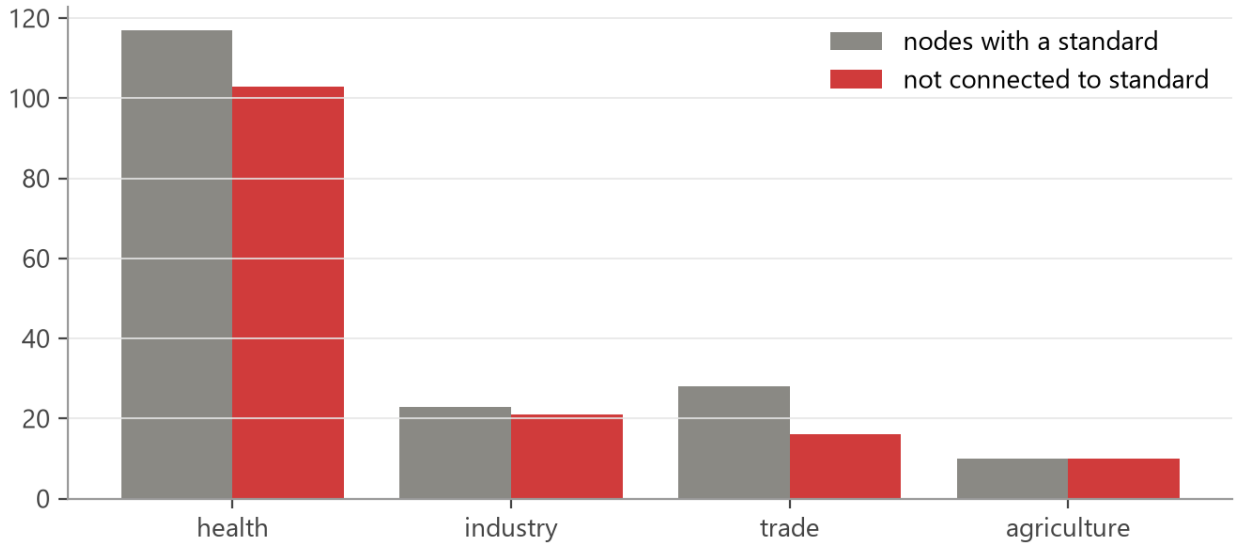

Figure 8.18: G6: sector nodes lacking the connection standard their sector requires

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 8.5: G6: sector nodes that lack their required connection standard

Sector	Nodes with a connection standard	Not connected to standard	Share	Output at stake (t/yr)
health	117	103	88.0%	0
industry	23	21	91.3%	0
trade	28	16	57.1%	0
agriculture	10	10	100.0%	9,764,387

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_sector_connection.csv).

The test is strict, and its results depend on the draft node standards and on surface data.

- **Health: 103 of 117 hospitals (88.0 %).** The 102 general hospitals need a district-class road of gravel or better, and the 15 regional referral hospitals a surfaced one. A path that crosses a link of unknown surface cannot be shown to be gravel and fails the test, so part of this share is a data gap.
- **Industry: 21 of 23 industrial parks and special economic zones (91.3 %).** They need a paved national primary connection with no lower-class connector.
- **Trade: 16 of 28 gateways (57.1 %).** The 22 border posts and 6 airports need a paved trunk connection.
- **Agriculture: all 10 farm zones (100 %)** with output above the 5,000 t threshold lack an all-season connection of the required class. They carry **9.76 million tonnes a year** of output at stake, the only G6 magnitude available.

Output at stake is zero for health, industry and trade because their nodes have no magnitudes: the sector exchanges that would supply patient loads, industrial output and gateway throughput are not yet connected (the minerals and trade feeds are synthetic demonstration data and are not used). The 309 mines enter as nodes without production figures, so no mine has a connection standard and none is tested. G6 findings go outbound to the sector engines (Chapter 22); they show *which* nodes lack a connection, and the sector models must say how much rides on each.

8.11 G7 Resilience

Figure 8.19 maps the people behind flood-exposed critical links, and Figure 8.20 ranks the districts.

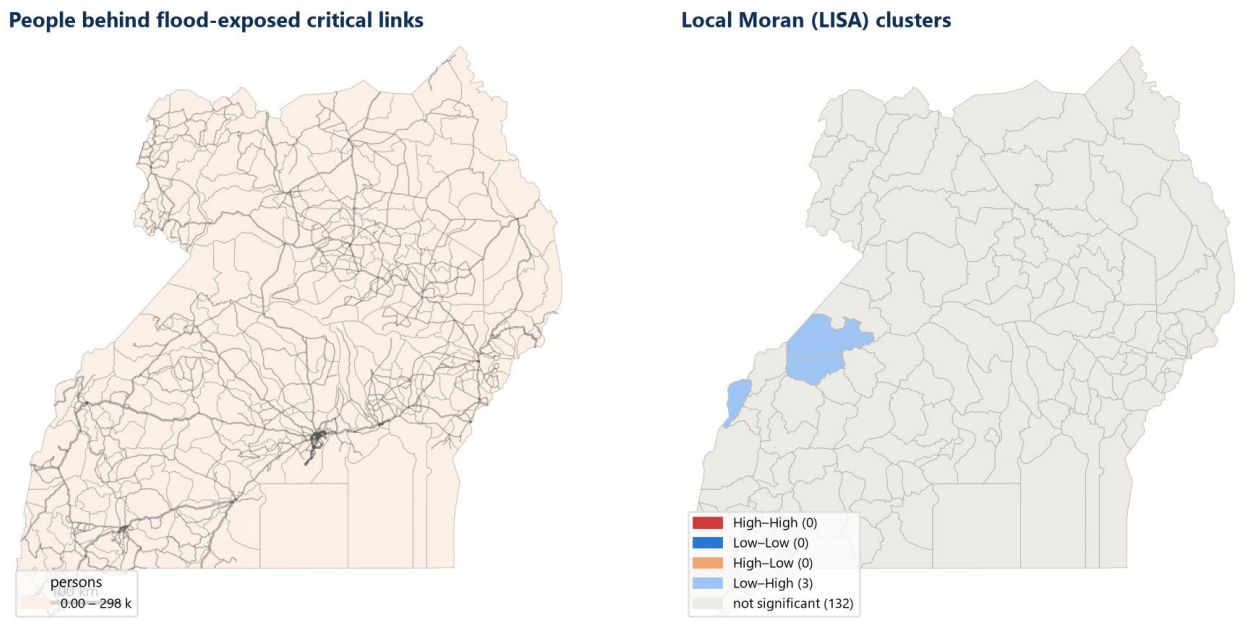

Figure 8.19: G7 Resilience: people behind flood-exposed critical links — by district; right: LISA clusters (this dimension only)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

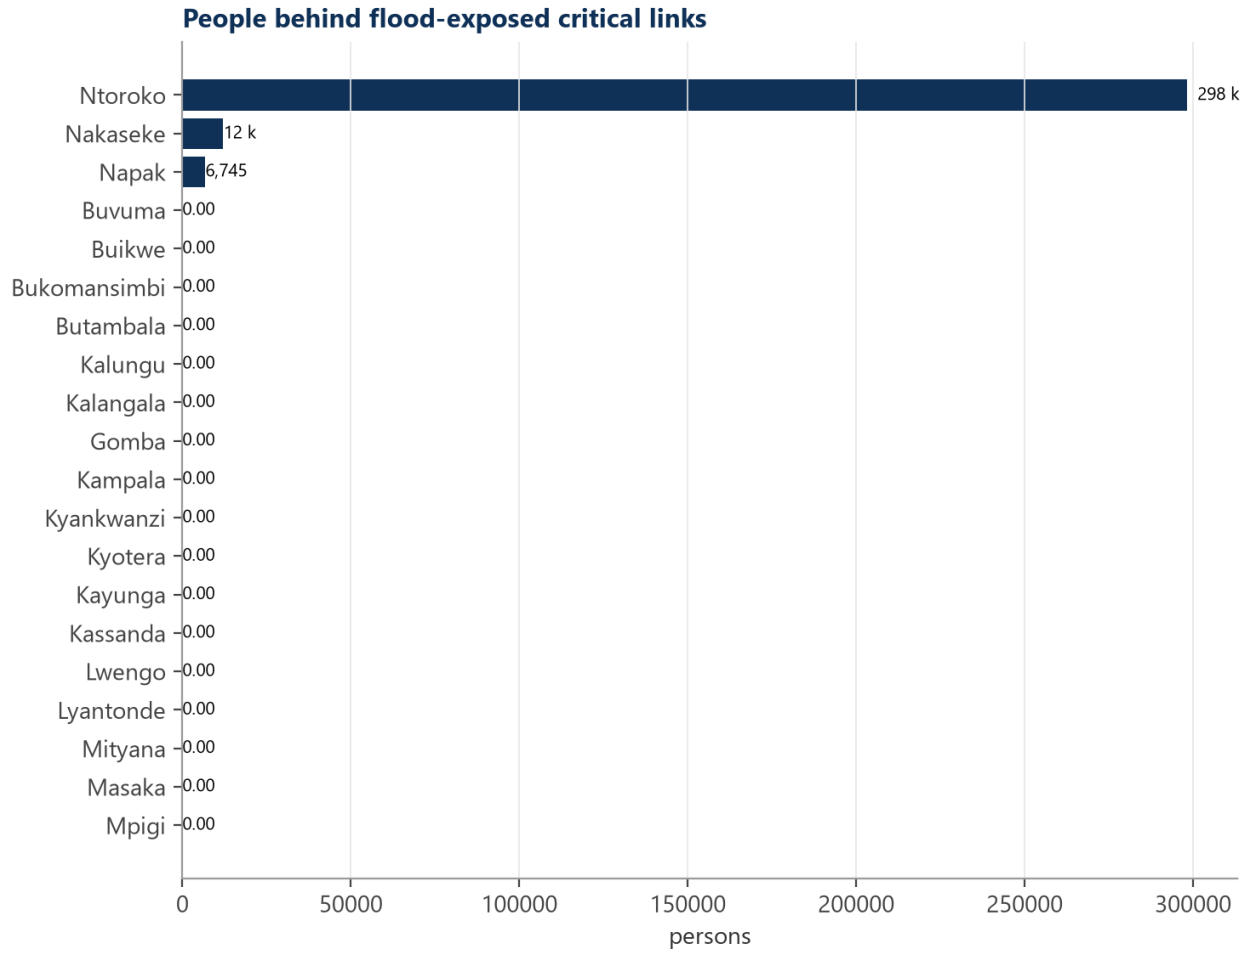

Figure 8.20: Top 20 districts: People behind flood-exposed critical links (one dimension, one sort column)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Only **three links** meet both tests: flood-exposed, and a cut edge with at least 2,000 people behind it. One in Ntoroko isolates 298,306 people, one in Nakaseke 12,089 and one in Napak 6,745; together 317,140 people. Every

other district has zero. The map is almost blank and the clustering statistic is meaningless (Moran's *I* = −0.007, p = 0.793). The Ntoroko link is the outstanding resilience finding: a single flood-exposed link on the Semliki plain whose loss would cut off a population of the size of a city, on a network component that extends beyond the district.

The narrowness of G7 is a data limit. Flood exposure is tested only at link midpoints against the JRC 100-year flood depth; the NASA landslide susceptibility layer is in the estate but not yet joined to links; and criticality is measured only through cut edges, not through detour length on links with an alternative. A fuller G7 would use SDI's hazard overlays and the detour-ratio test of the design.

8.12 The evidence gap: survey before works

Two lists in TI-05 are not works lists. They are the gaps in what is *known*, and they are the first things the plan should close. Figure 8.21 shows both.

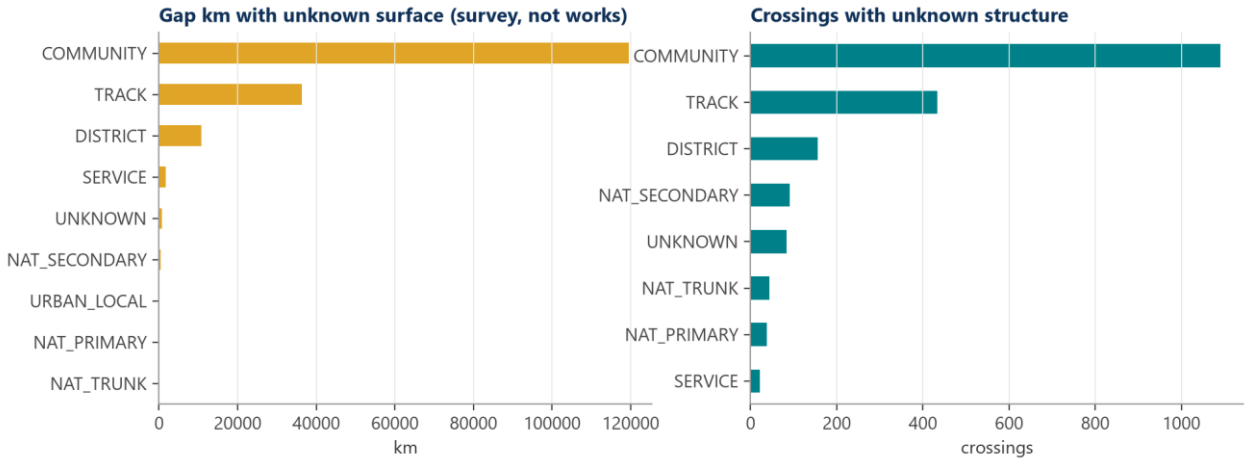

Figure 8.21: The evidence gap: gap km with unknown surface and river crossings with unknown structure, by class

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Gap km with unknown surface. Table 8.6 gives, by class, the 169,343 km of gap links whose surface is unknown. TI-07 lists them as a survey requirement, not as upgrades, because the response to not knowing a road's surface is to find out.

Table 8.6: Survey requirement: gap links whose surface is unknown

Class	Gap km with unknown surface (survey, not works)
NAT_TRUNK	2
NAT_PRIMARY	12
NAT_SECONDARY	330
DISTRICT	10,804
URBAN_LOCAL	44
COMMUNITY	119,458
TRACK	36,284
SERVICE	1,687
UNKNOWN	722

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_survey_requirement.csv).

Community roads account for 119,458 km and tracks for 36,284 km; district roads for 10,804 km; national roads for only 344 km. By district the largest survey lists are in Wakiso (7,436 km), Mukono (5,090), Nakaseke (3,961), Yumbe (3,858), Amuria (3,349), Mubende (3,290), Mityana (3,130), Arua (3,113), Amuru (3,026) and Luwero (3,021); the median district has 923 km to survey.

River crossings with unknown structure. The 1,958 crossings of Table 8.4 need a visit to record the structure (bridge, culvert, drift or none), its span, load rating and condition.

The survey is not an administrative detail. The unknown-surface km move the RAI between 28.7 % and 93.5 % (Chapter 7), set 70 % of the G4 standard gap, and drive most of the G6 failures. A district-by-district surface and structure survey, combined with access to the agencies' road asset management data for condition, would turn the largest parts of this register from "unknown" into either "no gap" or a costed works list.

8.13 Data quality and limitations

- **Two of seven dimensions are not observable at link level.** G2 has no link condition; G5 missing links has no candidate alignments. Both are stated as such, not filled.
- **Unknown surface counts as earth** in G1, G4 and G6. This is conservative for the plan and inflates the gaps; the upper RAI bound shows by how much.
- **"No destination" includes "network missing".** Without alignments, people a new road could serve are counted as needing a facility. The sector share of each service gap is therefore an upper bound.
- **Access standards are mostly assumed.** Only the 2 km RAI distance and the two-hour EmONC time are normed. A national standard for primary-health or market time would change the G1 counts directly.
- **Loads are modelled and capacity is mostly unknown.** G3 is a lower bound in coverage and an uncalibrated estimate in magnitude.
- **Node magnitudes are missing.** G6 shows which nodes fail, but not what is at stake except for agriculture.
- **Hazards are thin.** G7 uses one flood layer at link midpoints.
- **Population and speeds** carry the limits of Chapter 7.

8.14 Implications for planning

- **Each gap needs its own instrument and its own ranking.** The district at the top of one column is rarely at the top of the next: Wakiso leads on rural-access persons, capacity and cut edges; Kotido on primary health and EmONC; Mubende on markets; Kasese on the district centre; Ntoroko on wet-season closures and flood resilience; Sironko on fragility.
- **Most of the service-access gap cannot be closed by roads.** Between 59 % and 97 % of people beyond the health, school, market and district-centre standards remain beyond them on a fully upgraded network. These findings belong with the Health, Education and Agriculture engines.
- **The rural-access gap is a road gap, but mostly an unrecorded one.** 98 % of the 33.3 million people beyond 2 km would be served by raising existing links to standard, yet most of those links have no surface record. Survey first, then works.
- **Wet-season closures are the cheapest access to buy.** 3,299 km close, concentrated in Ntoroko, Kasese, West Nile and Karamoja; spot improvements at mapped crossings address them link by link.
- **A small number of links carry enormous risk.** One flood-exposed link in Ntoroko isolates 298,000 people; single cut edges on Elgon isolate over 700,000. These deserve engineering assessment regardless of any district ranking.
- **Condition must be programmed by network until agency data arrive.** The 107,837 km of poor non-urban roads are known only as totals; their treatment is a network-level programme until UNRA and DUCAR condition data can locate them.
- **Sector nodes need magnitudes.** Without patient loads, industrial output and gateway throughput, G6 can say that 88 % of hospitals and 91 % of industrial parks lack their connection, but not which of them matter most.

CHAPTER 9

9 Costing: unit prices, maintenance and the backlog

This chapter answers the fourth planning question: *what will it cost to close the gaps?* It explains where the unit prices come from, how the engine turns a gap on a link into a priced works item, what the existing network costs to maintain each year, and what it would cost to treat the roads that the agencies report in poor condition. It implements the TI-06 module, design chapter 8 (`docs/08_TI06_COSTING.md`), with the evidence-price method recorded in `data/raw/SOURCES.md` ("Unit costs, round 2") and `IMPLEMENTATION_STATUS.md` ("Unit-cost evidence and run 8"). All results are from the reference run `014c9cd7`, price vintage `UCSN+EV-UGA-2026.09b`, in constant 2026 Uganda shillings (UGX). Every price in this run is an **evidence candidate** (quality tier E5, confidence low), usable only in an exploratory run. The costs are therefore a planning range, not budget estimates.

9.1 What the module does

TI-06 prices three things, each at low, central and high values:

- **Works:** the capital cost of moving a link from its existing standard to its required one, of restoring it, or of fixing a point on it. Works are priced per km, per lane-km, per metre of span or per site.
- **Network maintenance:** the routine maintenance (RM) every maintainable link needs every year, whatever else the plan does.
- **Induced maintenance:** the change in routine maintenance that each new or upgraded asset causes for the rest of its life. Every capital figure in the report carries it.

Two further products follow from the agency condition statistics and from the price file itself:

- **The condition backlog:** the cost of treating the poor kilometres reported by MoWT and UBOS, network by network (`tables/ti06_condition_backlog.csv`).
- **The unit cost catalogue:** every map row from a works key to a price code, with its price and status (`tables/ti06_unit_cost_catalogue.csv`), and the list of works that could not be priced (`tables/unpriced_codes.csv`).

TI-06 owns no prices. Unit costs belong to the Unit Cost Standards Network (UCSN, the UnitCostModelling engine). TIPRAS reads UCSN's coefficients and, where UCSN has no code or only an illustrative one, a TIPRAS **evidence price file** built from documents. A missing price is never set to zero: the item stays in the programme, unpriced, and is never funded.

9.2 Method

9.2.1 Where the prices come from: two price files

A price vintage is the set of price rows loaded under one name. The reference vintage holds two files:

- **UCSN illustrative coefficients** (tier E7, confidence very low): seven transport codes, for paved construction by terrain, gravel construction, bridges per metre and routine maintenance. UCSN itself calls them regional orders of magnitude.
- **The TIPRAS evidence file** `config/prices/UGA-2026.09b-evidence.csv` (tier E5, status `EVIDENCE_CANDIDATE`): 14 codes built by `scripts/build_evidence_prices.jl` from documented Ugandan unit costs.

The norm set `draft-f` maps every works key to an evidence code. The UCSN coefficients stay in the vintage for comparison, but no map row points to them, and the evidence file contains no code that UCSN publishes: TIPRAS never overwrites a UCSN price. The comparison of the two is sent to UCSN (`data/exchange/outbound/ucsn/`), and UCSN decides whether to adopt the evidence.

9.2.2 The evidence base

The evidence file rests on **275 extracted rows** in `data/raw/unit_costs_extracted/`, each with its document, page and a verbatim quote. The sources include:

- MoWT programme performance reports (FY2019/20 to FY2022/23) and statistical abstracts: UNRA average unit costs and project costs;
- Ministry of Finance Budget Monitoring and Accountability Unit (BMAU) road monitoring reports, FY2019/20 to FY2022/23: district force-account and contract costs, Community Roads Improvement Programme (CRIP) contracts, the Gem Farm Bailey bridge and spot improvements;
- Uganda Road Fund budgeting guidelines FY2018/19 to FY2020/21 and its annual report FY2022/23;
- World Bank project appraisal and completion reports (TSDP, RDP2, ARSDP, NERAMP) and the ROCKS road-cost database (Uganda and 11 other East and Southern African countries);
- the IFAD PRELNOR design report, the AfDB CAIIP appraisal, the 2013 low-cost seal study, and Rwandan and Tanzanian project documents as regional comparators.

Of the 275 rows, **89 are used** in the 14 prices. All 89 are tier-1 evidence (Uganda, dated 2011 or later); the regional and older rows are cited in each price's `basis` field as cross-checks only. The rows used, converted to 2026 prices, are in `data/reference/unit_cost_evidence_rows_UGA.csv`.

9.2.3 Carrying each observation to UGX 2026

Every row is converted to 2026 shillings with World Bank WDI price indices (`data/reference/price_indices_UGA.csv`). For a value *p* observed in year *y*, with Uganda's GDP deflator *D(UGA)*, the US GDP deflator *D(USA)* and the exchange rate *FX* in UGX per US dollar, a shilling value is simply deflated:

$$p_{2026}^{UGX} = p_y \cdot \frac{D_{2026}^{UGA}}{D_y^{UGA}} \tag{9.1}$$

A dollar value can reach 2026 shillings by two paths that give different answers when the exchange rate has not tracked relative inflation. The **local path** converts at the exchange rate of the price year and then inflates with Uganda's deflator; the **dollar path** inflates with the US deflator and then converts at the 2026 rate:

$$p^{loc} = p_y \, FX_y \, \frac{D_{2026}^{UGA}}{D_y^{UGA}}, \qquad p^{dol} = p_y \, \frac{D_{2026}^{USA}}{D_y^{USA}} \, FX_{2026} \tag{9.2}$$

The engine keeps both. The row's central value is their geometric mean, and the row's range spans the lower path of its low value to the higher path of its high value:

$$p^{c} = \sqrt{p^{loc} \, p^{dol}}, \qquad p^{lo} = \min(p^{loc}, p^{dol}), \quad p^{hi} = \max(p^{loc}, p^{dol}) \tag{9.3}$$

This replaces the assumed 5 % a year inflation of the first evidence vintage. Figure 9.1 shows the indices, rebased to 2026 = 100.

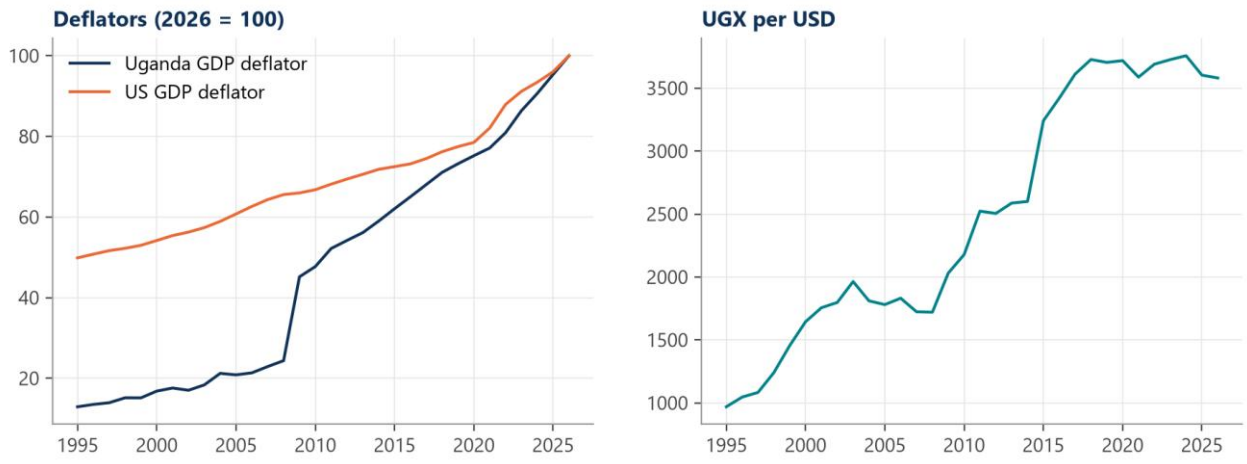

Figure 9.1: The price indices used to carry evidence to 2026 prices

Source: World Bank WDI (NY.GDP.DEFL.ZS, PA.NUS.FCRF); 2026 extrapolated

Three features of the indices matter for the prices.

- **The 2026 values are extrapolated.** WDI ends in 2025; 2026 continues the five-year mean trend: the Uganda deflator at +4.87 % a year, the US deflator at +4.12 % and the exchange rate at −0.63 %, giving UGX 3,580 per dollar in 2026.
- **The Uganda deflator has a step in 2008–09**, visible in the left panel, where the WDI series jumps from about 24 to 45 (2026 = 100). No row used in the prices is dated before 2011, so the step does not reach the prices; it would matter for the ROCKS records of the 1990s and 2000s, which is one reason they are cross-checks only.
- **The shilling depreciated fastest in 2009–2016** (right panel), from about 1,720 to 3,420 per dollar, and has been almost flat at 3,580–3,760 since 2017. For dollar rows dated 2019–2021 the two paths therefore give similar values; for older dollar rows they diverge.

9.2.4 From rows to a price code

Each price code has a central value and a range:

- **Tiers.** Rows are graded tier 1 (Uganda, 2010 or later), tier 2 (Uganda any year, or the East and Southern African region 2010 or later) and tier 3 (anything else). Only the best tier present is used.
- **Central value.** The median of the row centrals in that tier. For paved rehabilitation and for reseal, where MoWT publishes a series, the latest year's value is used instead.
- **Range.** The minimum and maximum of the converted rows in that tier.

Figure 9.2 counts the observations behind each code, and Table 9.1 gives the counts with the years covered.

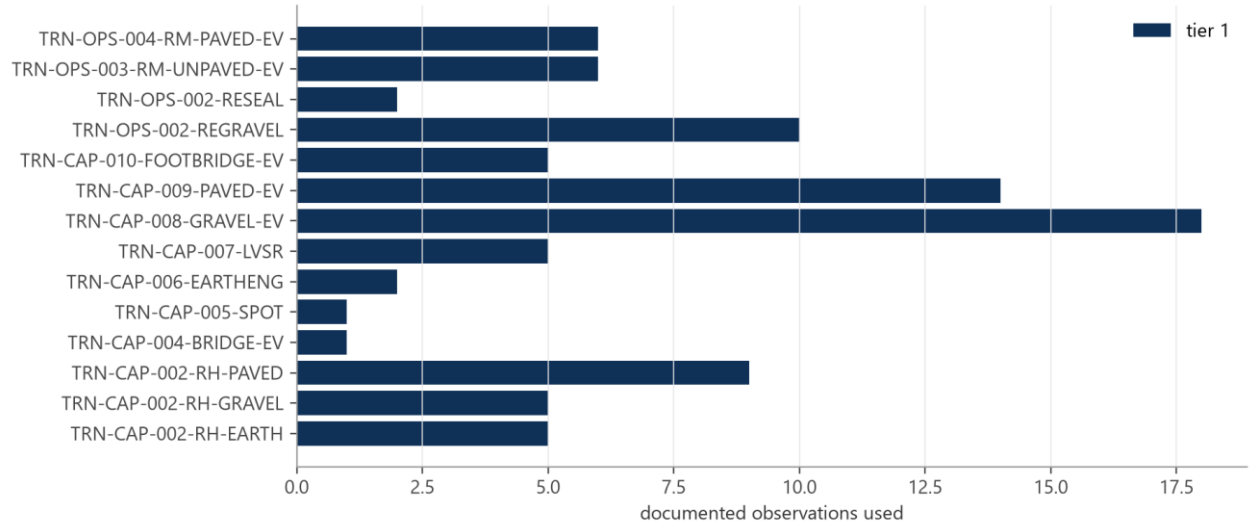

Figure 9.2: Observations behind each evidence price, by evidence tier

Source: data/reference/unit_cost_evidence_rows_UGA.csv

Table 9.1: How many documented observations stand behind each price

Price code	Rows used	Best tier	Years
TRN-CAP-002-RH-EARTH	5	1	2020–2020
TRN-CAP-002-RH-GRAVEL	5	1	2020–2020
TRN-CAP-002-RH-PAVED	9	1	2016–2023
TRN-CAP-004-BRIDGE-EV	1	1	2019–2019
TRN-CAP-005-SPOT	1	1	2020–2020
TRN-CAP-006-EARTHENG	2	1	2020–2020
TRN-CAP-007-LVSR	5	1	2013–2020
TRN-CAP-008-GRAVEL-EV	18	1	2011–2023
TRN-CAP-009-PAVED-EV	14	1	2013–2023
TRN-CAP-010-FOOTBRIDGE-EV	5	1	2023–2023
TRN-OPS-002-REGRAVEL	10	1	2019–2021
TRN-OPS-002-RESEAL	2	1	2020–2021
TRN-OPS-003-RM-UNPAVED-EV	6	1	2014–2021
TRN-OPS-004-RM-PAVED-EV	6	1	2019–2021

Tier 1: Uganda, 2010 or later; tier 2: Uganda any year, or the region 2010 or later; tier 3: other. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_evidence_rows.csv).

The evidence is uneven. Gravel construction rests on 18 observations (2011–2023) and paved upgrading on 14 (2013–2023); regravelling on 10 and paved rehabilitation on 9. At the other end, **the bridge price rests on one Bailey bridge** (Gem Farm, 30.4 m, FY2018/19, including approaches) and **the spot-improvement price on one force-account site** (two bottlenecks on a 29 km road, FY2019/20). Engineered earth has two observations, one contract (Nakitoma–Nabiswera, 10.4 km) and one force-account section, and reseal two. Force-account budgets understate full cost, because MoWT supplies equipment, labour and culverts outside them (BMAU FY2019/20). They are excluded from gravel rehabilitation and kept only where they are the only evidence.

Figure 9.3 plots every observation used, converted to UGX 2026, on a log scale.

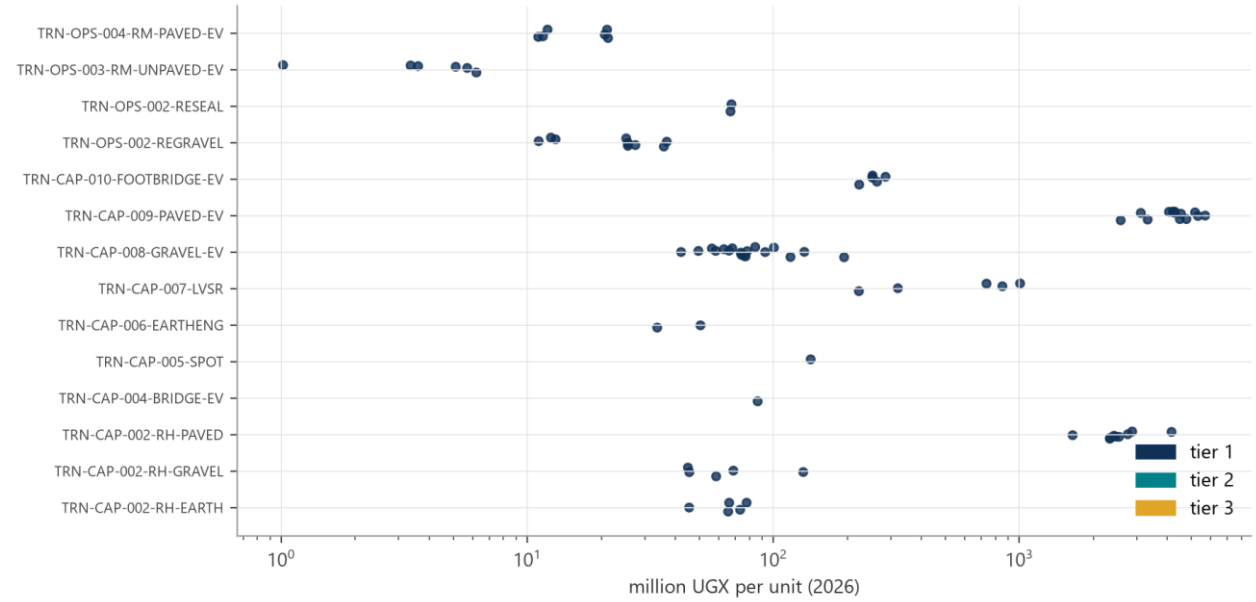

Figure 9.3: Every evidence observation converted to UGX 2026, by price code (log scale)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The scatter shows how tight or wide each price is. Paved rehabilitation (1.6 to 4.2 bn per km) and paved upgrading (2.6 to 5.7 bn per km) cluster within a factor of about two; their observations are UNRA programme averages and named projects. Gravel construction spreads from 42 m to 194 m per km: most CRIP community-access contracts fall between 50 m and 135 m, and the 194 m value is the Bweramule–Kayanja contract in Ntoroko, a 4 km road on the Semliki floodplain. Unpaved routine maintenance spans almost a factor of ten, from 1.0 m per km a year (PRELNOR road gangs paid UGX 100,000 a month) to 6.2 m (mechanised district maintenance). The low-volume seal rows split into two groups: the 2013 low-cost-seal study (220–320 m per km) and three MoWT low-cost sealing contracts of 2020 (735–1,007 m per km), which included swamp raising and base works.

9.2.5 The evidence prices

Table 9.2 lists the 14 evidence prices, and Figure 9.4 plots them beside the UCSN illustrative prices.

Table 9.2: Evidence unit prices, UGX 2026

Price code	Works	Unit	Low (M UGX)	Central	High	Tier
TRN-CAP-002-RH-PAVED	Rehabilitation/reconstruction of paved roads	UGX/km	1,645.7	2,339.0	4,152.4	E5
TRN-OPS-002-REGRAVEL	Periodic maintenance, regravelling (district/DUCAR unpaved)	UGX/km	9.3	25.6	38.4	E5
TRN-OPS-002-RESEAL	Periodic maintenance, paved (reseal)	UGX/km	64.3	67.6	69.8	E5
TRN-CAP-007-LVSR	Low-volume sealed road (low-cost seal incl. preparatory works)	UGX/km	222.7	735.2	1,075.4	E5
TRN-OPS-003-RM-UNPAVED-EV	Routine maintenance, district/community unpaved (manual–mechanised)	UGX/km/year	1.0	4.4	8.9	E5
TRN-OPS-004-RM-PAVED-EV	Routine maintenance, paved roads	UGX/km/year	9.7	16.4	24.3	E5
TRN-CAP-008-GRAVEL-EV	Construction of gravel road / upgrading earth to gravel	UGX/km	42.2	75.1	193.9	E5
TRN-CAP-009-PAVED-EV	Upgrading to paved (bituminous), all terrain (national mix)	UGX/km	2,580.8	4,269.0	5,690.3	E5
TRN-CAP-002-RH-GRAVEL	Rehabilitation/reconstruction of gravel roads	UGX/km	45.0	58.6	132.3	E5
TRN-CAP-002-RH-EARTH	Rehabilitation of earth/community access roads	UGX/km	45.5	66.1	77.9	E5
TRN-CAP-006-EARTHENG	Opening/upgrading to engineered earth (basic access, community access road standard)	UGX/km	33.7	42.1	50.6	E5
TRN-CAP-005-SPOT	Spot improvement on a low-volume road (culvert line, drift or vented ford, fill), per site	UGX/unit	141.0	141.0	141.0	E5
TRN-CAP-004-BRIDGE-EV	Road bridge (district/national), per metre of span	UGX/m	75.4	86.2	86.2	E5

Price code	Works	Unit	Low (M UGX)	Central	High	Tier
TRN-CAP-010-FOOTBRIDGE-EV	Pedestrian/trail bridge, per bridge	UGX/unit	223.4	253.0	285.8	E5

Built by scripts/build_evidence_prices.jl from 275 documented rows; WDI deflators; status EVIDENCE_CANDIDATE (exploratory only). Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_evidence_prices.csv).

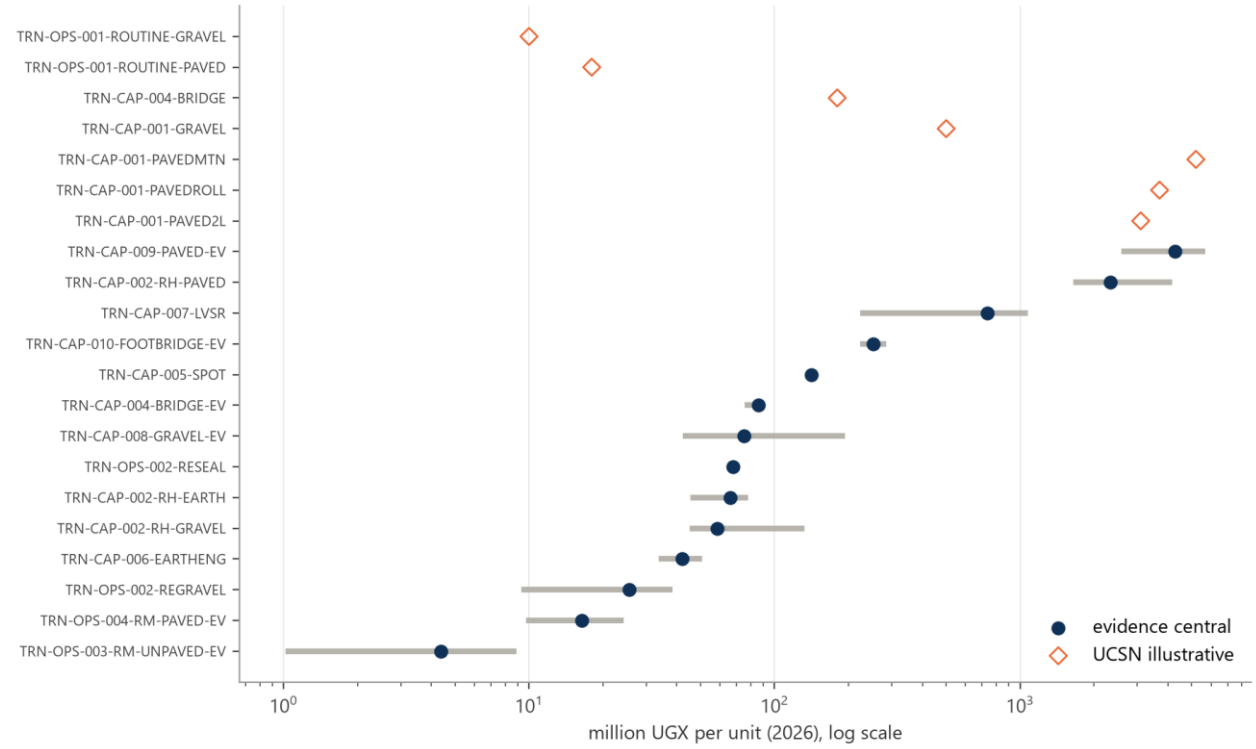

Figure 9.4: Evidence unit prices (low · central · high) against the UCSN illustrative prices

Source: config/prices evidence file; UnitCostModelling UCSN coefficients

The comparison with UCSN is the most important reading of the figure.

- **Gravel construction:** the evidence central of 75 m per km is **0.15 of UCSN's 500 m**. The Ugandan contracts for community and district gravel roads are far below the regional figure UCSN uses.
- **Unpaved routine maintenance:** 4.4 m per km a year, 0.44 of UCSN's 10 m.
- **Paved routine maintenance:** 16.4 m, 0.91 of UCSN's 18 m.
- **Paved upgrading:** 4,269 m per km against UCSN's 3,100 m (flat), 3,700 m (rolling) and 5,200 m (mountainous), ratios of 1.38, 1.15 and 0.82. No source splits paved upgrading by terrain, and ROCKS regional medians show no consistent terrain gradient, so one evidence code covers all terrain.
- **Bridges:** 86 m per metre of span against UCSN's 180 m, but on one observation.

The evidence therefore lowers the cost of the gravel and earth works that dominate rural access, and raises the cost of paving flat roads. Both changes move the programme of Chapter 10.

9.2.6 Resolving works to a price: the unit-cost map

TI-07 assigns each link a works type and a target surface; TI-06 has to find its price. The key is the works type, the existing surface, the target surface and the terrain, and the map `tpl_norm.unit_cost_map` resolves the key to a price code and a quantity basis. A wildcard * matches any value. Table 9.3 is the resulting catalogue, every map row with its code and resolved price.

Table 9.3: Unit cost catalogue applied in the reference run (map key: works, from-surface, to-surface, terrain)

Map key (works|from|to|terrain)	Price code	Basis	Low (M UGX)	Central	High	Tier	Status
NC|none|dbst|flat	TRN-CAP-009-PAVED-EV	per_km	2,580.8	4,269.0	5,690.3	E5	EVIDENCE_CANDIDATE
NC|none|gravel|*	TRN-CAP-008-GRAVEL-EV	per_km	42.2	75.1	193.9	E5	EVIDENCE_CANDIDATE
PM|paved|*	TRN-OPS-002-RESEAL	per_km	64.3	67.6	69.8	E5	EVIDENCE_CANDIDATE
PM|unpaved|*	TRN-OPS-002-REGRAVEL	per_km	9.3	25.6	38.4	E5	EVIDENCE_CANDIDATE
RH|gravel|*	TRN-CAP-002-RH-GRAVEL	per_km	45.0	58.6	132.3	E5	EVIDENCE_CANDIDATE
RH|paved|*	TRN-CAP-002-RH-PAVED	per_km	1,645.7	2,339.0	4,152.4	E5	EVIDENCE_CANDIDATE
RH|unpaved|*	TRN-CAP-002-RH-EARTH	per_km	45.5	66.1	77.9	E5	EVIDENCE_CANDIDATE
RM|paved|*	TRN-OPS-004-RM-PAVED-EV	per_km_year	9.7	16.4	24.3	E5	EVIDENCE_CANDIDATE
RM|unpaved|*	TRN-OPS-003-RM-UNPAVED-EV	per_km_year	1.0	4.4	8.9	E5	EVIDENCE_CANDIDATE
SI|/	TRN-CAP-005-SPOT	per_unit	141.0	141.0	141.0	E5	EVIDENCE_CANDIDATE
ST|/|*	TRN-CAP-004-BRIDGE-EV	per_m_span	75.4	86.2	86.2	E5	EVIDENCE_CANDIDATE
UP|/*ac*/	TRN-CAP-009-PAVED-EV	per_km	2,580.8	4,269.0	5,690.3	E5	EVIDENCE_CANDIDATE
UP|/*dbst*/	TRN-CAP-009-PAVED-EV	per_km	2,580.8	4,269.0	5,690.3	E5	EVIDENCE_CANDIDATE
UP|/*earth_eng*/	TRN-CAP-006-EARTHENG	per_km	33.7	42.1	50.6	E5	EVIDENCE_CANDIDATE
UP|/*gravel*/	TRN-CAP-008-GRAVEL-EV	per_km	42.2	75.1	193.9	E5	EVIDENCE_CANDIDATE
UP|/*lvsr*/	TRN-CAP-007-LVSR	per_km	222.7	735.2	1,075.4	E5	EVIDENCE_CANDIDATE

Tier E5: evidence candidate read from documents; E7: UCSN illustrative. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_unit_cost_catalogue.csv).

The works cost of an item *j* on its links *L(j)*, with quantity *q* (km, lane-km, metres of span or sites), unit price κ resolved from the map and the design-and-supervision and contingency rates ϑ of `tpl_norm.project_overheads`, is:

$$X_j^{cap} = (1 + \theta_w^{DS} + \theta_w^{cont}) \sum_{l \in L_j} q_w(l)\, \kappa_w(S_l \to S_l^*, \text{terrain}_l) \tag{9.4}$$

It is evaluated three times, at the low, central and high price. The overhead rates are feasibility-stage rates, assumed until the agencies confirm them: 8 % design and supervision plus 15 % contingency for upgrading and spot improvements, 8 % plus 12 % for rehabilitation, 5 % plus 8 % for periodic maintenance and 10 % plus 20 % for new construction and structures. Right-of-way cost is a separate line; no source supplies it, so it is reported as missing, not zero.

Four features of the map shape the programme.

- **Terrain does not change the price.** Except for new paved construction (`NC`, flat only), every map row has terrain `*`, because no source separates terrain. Mountain roads on Elgon and in the south-west are priced like flat ones.
- **The target surface sets the code.** Upgrading to asphalt (`ac`) and to surface dressing (`dbst`) both resolve to the single paved-upgrading code, at 4,269 m per km central (5,251 m with overheads). Upgrading to gravel resolves to 75 m (92 m with overheads), to engineered earth to 42 m (52 m), and to low-volume seal to 735 m (904 m).
- **There is no paved-to-paved code.** The map has no price for surface-dressing to asphalt, or for a low-volume seal to a full paved standard. The key `UP` from `dbst` to `ac` matches the row "UP, any surface, to ac", and the link is costed at the **full upgrading price**, as if it were earth. In the reference run 375 such projects cover 2,499 km and cost 13.1 tn UGX central, 15.3 % of all upgrading capital, for a change (asphalt overlay on a sealed road) that should cost a fraction of new paving. The costing studio of the web interface (Chapter 15) flags this and keeps trunk roads at surface dressing by default; the map needs a resurfacing or overlay code, which is a request to UCSN.
- **Spot improvements on gravel and paved links do not resolve.** 23 spot-improvement items, all on gravel or asphalt links, find no map row and stay unpriced (Table 9.4). They are kept in the programme and never funded.

Table 9.4: Works that remain unpriced (never funded)

Unpriced code or map key	Items
SI|gravel|gravel|flat	8
SI|ac|ac|rolling	6
SI|ac|ac|flat	4
SI|gravel|gravel|rolling	4
SI|gravel|gravel|mountainous	1

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_unpriced.csv).

9.2.7 Network maintenance requirement

Every maintainable link needs routine maintenance every year. Tracks, service roads and ferries are not maintainable in the norms. The requirement is the sum over the maintainable links of their length times the routine rate for their surface group, paved or unpaved (unknown surface counts as unpaved):

$$M^{RM} = \sum_{l \in \mathcal{M}} \text{len}_l \cdot rm(\text{surf}_l) \tag{9.5}$$

Periodic maintenance (regravelling every 5 years, resealing every 8 years in the draft norms) has evidence prices but is not added to the annual requirement in the reference run. It enters instead through the condition programmes below, for the kilometres the agencies report as poor.

9.2.8 Induced maintenance

Each upgraded link changes its own routine maintenance from the rate of its current surface to the rate of its target surface. A link that was not maintainable before (a track) adds the full new rate:

$$IM_j = \sum_{l \in L_j} \text{len}_l \left(rm(S_l^*) - rm(S_l) \cdot \mathbf{1}(l \in \mathcal{M}) \right) \tag{9.6}$$

The sign is kept. In the draft-f norms the routine rate depends only on paved versus unpaved, so upgrading gravel to paved raises induced maintenance by 12.0 m per km a year, and upgrading earth to gravel or engineered earth changes it only where the link was a track. Condition programmes induce nothing, because they restore an existing standard.

9.2.9 The condition backlog

Where condition is known only by network class, the backlog is the cost of treating each network's poor kilometres with the treatment the draft norm assigns to that network (`tpl_norm.condition_programme`):

$$BL_n = km_n^{poor} \cdot \kappa(w_n, \text{surf}_n) \cdot (1 + \theta_w^{DS} + \theta_w^{cont}) \tag{9.7}$$

The treatments are illustrative and the road agencies must confirm them: rehabilitation for poor national paved roads; regravelling (periodic maintenance) for national unpaved and district roads; rehabilitation for community roads.

9.2.10 Lifecycle cost and cost bases

For appraisal (Chapter 10), each item's cost stream is discounted over a 20-year period at an assumed 12 % a year. Capital is spread over the construction years with a symmetric beta S-curve (3 years for upgrading, 2 for rehabilitation, 1 for spot improvements and periodic maintenance), and induced maintenance is paid every year after completion:

$$PVC_j = \sum_{t=1}^{n_j} \frac{X_j^{cap} \, s_{w,t}}{(1+r)^t} + \sum_{t=n_j+1}^{20} \frac{IM_j}{(1+r)^t} \tag{9.8}$$

where *s(w,t)* is the cash share of year *t* from the S-curve. The design distinguishes three cost bases, which are never presented as one another: the **requirement** (all gaps closed plus maintenance of the whole network), the **gap** (the incremental capital, induced maintenance and backlog of the programme) and the **allocated** cost (what TI-08 funds; Chapter 11).

9.3 Results: what the existing network costs to maintain

The routine maintenance requirement of the existing network is 954.7 bn UGX a year at the central evidence rates, on 199,221 km of maintainable roads. Paved roads are 7,175 km of the maintainable network and 12.3 % of the requirement; the rest is unpaved and unknown-surface roads at 4.4 m per km. At the low and high rates (1.0 and 8.9 m per km unpaved, 9.7 and 24.3 m paved) the range is wide, because unpaved routine maintenance is the least certain price that matters.

Figure 9.5 maps the requirement by district, with its LISA clusters.

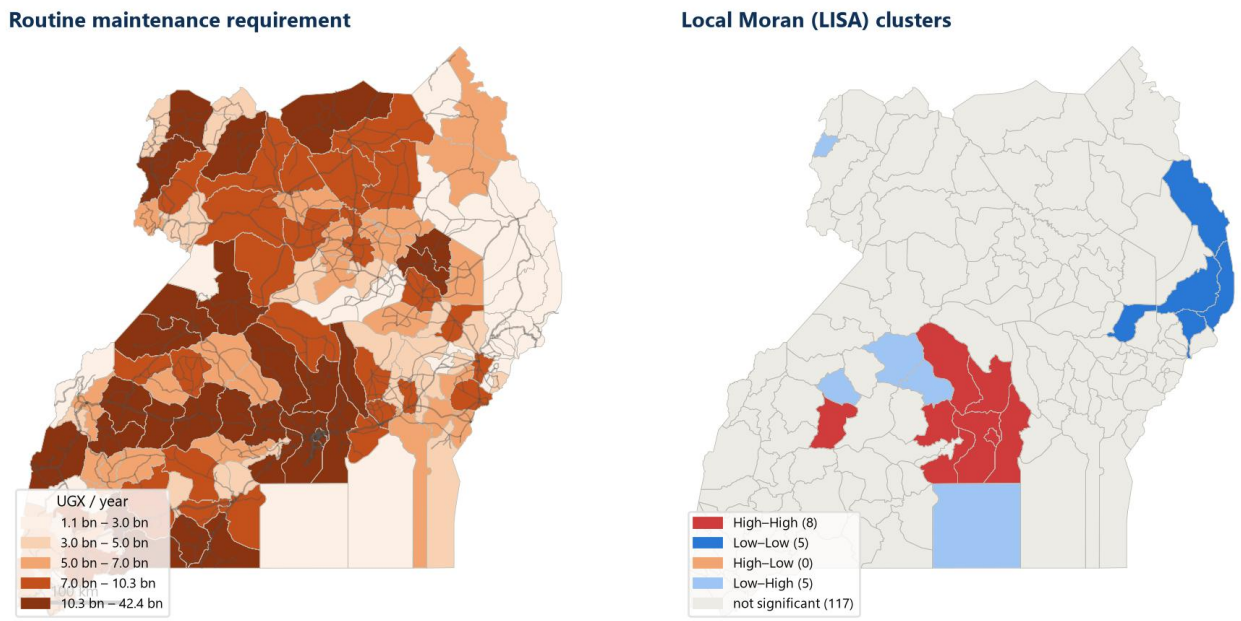

Figure 9.5: Routine maintenance requirement of the existing network, by district; right: LISA clusters

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The figure is drawn from the district layer of simulation S0, which has exactly the reference run's configuration (norms `draft-f`, evidence prices); its district values sum to the reference run's 954.7 bn a year.

- **Largest requirements:** Wakiso (42.4 bn a year), Mukono (23.5 bn), Kyenjojo (20.2 bn), Yumbe (17.5 bn), Arua (17.2 bn), Mubende (16.0 bn), Amuria (15.4 bn), Mityana (14.9 bn), Isingiro (14.8 bn) and Luwero (14.4 bn).
- **Smallest:** Bududa (1.1 bn), Nabilatuk (1.2 bn), Kalangala (1.3 bn), Nakapiripirit and Bulambuli (1.7 bn each). The median district needs 6.2 bn a year.
- **Clusters:** a high–high cluster of eight districts around Kampala (Kampala, Kyegegwa, Luwero, Mityana, Mpigi, Mukono, Nakaseke, Wakiso) and a low–low cluster in Karamoja and Elgon (Amudat, Kumi, Kween, Moroto, Nakapiripirit). Moran's *I* is 0.300.

The requirement follows network length, not need: districts with long mapped local networks need the most maintenance money. Because 169,343 km of the maintainable network have no recorded surface and are priced as unpaved, the requirement on the local networks is as uncertain as their surface. Chapter 13 sets the requirement, plus the induced maintenance of the allocated programme, against the maintenance-eligible resources: on the steady-state path the funding ratio is 0.58 (1,017 bn required against 592 bn available).

9.4 Results: the condition backlog

Figure 9.6 and Table 9.5 cost the treatment of the 107,837 km of non-urban roads that the agencies report as poor.

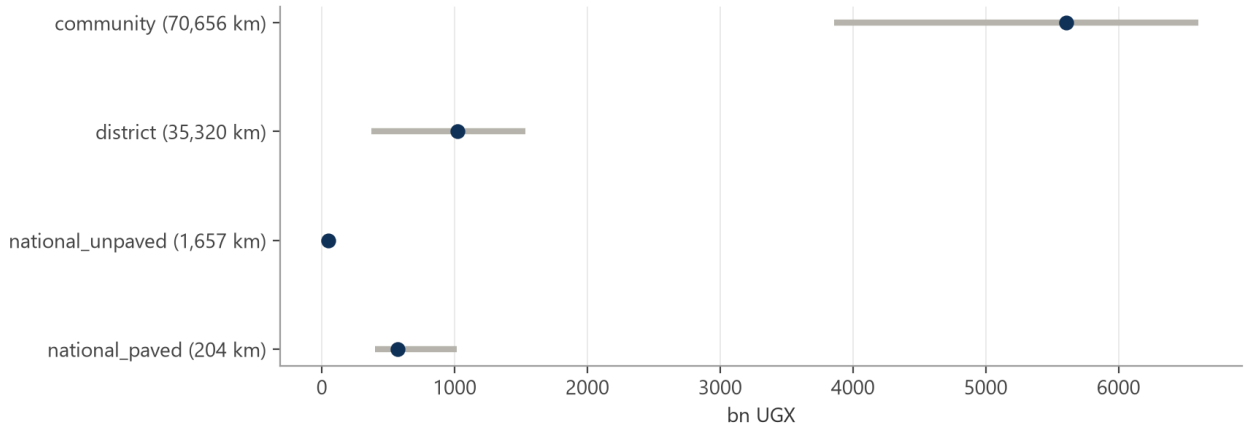

Figure 9.6: The condition backlog: cost of treating the agency-reported poor km (low · central · high)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 9.5: Condition backlog: the cost of treating the agency-reported poor km

Network	Works	Km poor	Low (bn UGX)	Central	High	Price code
national_paved	RH	204	402.9	572.6	1,016.5	TRN-CAP-002-RH-PAVED
national_unpaved	PM	1,657	17.4	48.0	71.8	TRN-OPS-002-REGRAVEL
district	PM	35,320	371.5	1,023.1	1,531.0	TRN-OPS-002-REGRAVEL
community	RH	70,656	3,857.4	5,608.6	6,602.0	TRN-CAP-002-RH-EARTH

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_backlog.csv).

The backlog is 7.25 tn UGX at central prices, with a range of 4.65 to 9.22 tn. It induces no new routine maintenance, because it restores existing roads.

- **Community roads dominate.** Rehabilitating 70,656 km at 66 m per km (plus 20 % overheads) costs 5.61 tn (3.86 to 6.60 tn), 77 % of the backlog. This was unpriced in earlier runs; the five DUCAR rehabilitation contracts of FY2019/20 (45.5 to 77.9 m per km) now price it.
- **District roads** need regravelling of 35,320 km: 1.02 tn (0.37 to 1.53 tn). The wide range reflects the regravelling evidence, from URF force-account rates of 11–13 m per km to its planning example of 36 m.
- **National paved roads** need rehabilitation of 204 km: 0.57 tn (0.40 to 1.02 tn). At 2.3 bn per km this is the most expensive treatment per km, but the smallest in length.
- **National unpaved roads** need regravelling of 1,657 km: 48 bn (17 to 72 bn).

Two cautions. The treatment per network is a norm, not an engineering assessment; a district road reported poor may need rehabilitation rather than regravelling. And the backlog is costed on the agencies' kilometres, which do not match the snapshot's classification (Chapter 7). Urban roads (31,854 km poor) are reported in the gap but neither costed nor programmed.

9.5 Results: which prices carry the programme

The unit cost catalogue has 16 map rows, but a handful of them carry almost all the money. Grouping the reference programme of Chapter 10 (92.8 tn UGX central) by the price code that priced it shows how concentrated the cost is (Table 9.6).

Table 9.6: Programme capital by the price code that priced it, central values, UGX 2026

Price code	Works it prices in the programme	Km	Capital, central (tn UGX)	Share
TRN-CAP-009-PAVED-EV	upgrading to asphalt or surface dressing	15,562	81.72	88.1 %
TRN-CAP-002-RH-EARTH	community-road condition programme	70,656	5.61	6.0 %
TRN-CAP-006-EARTHENG	upgrading earth to engineered earth	37,109	1.92	2.1 %

Price code	Works it prices in the programme	Km	Capital, central (tn UGX)	Share
TRN-CAP-008-GRAVEL-EV	upgrading earth to gravel	20,032	1.85	2.0 %
TRN-OPS-002-REGRAVEL	national unpaved and district condition programmes	36,977	1.07	1.2 %
TRN-CAP-002-RH-PAVED	national paved condition programme	204	0.57	0.6 %
TRN-CAP-007-LVSR	upgrading earth to low-volume seal	21	0.02	0.0 %

One price, paved upgrading at 4,269 m per km central, carries 88 % of the programme. It prices 15,562 km of links that TI-03 requires to be paved because of their traffic or function: 12,150 km from earth, 913 km from gravel and 2,499 km from surface dressing. Its own range, 2,581 to 5,690 m per km, accounts for most of the programme's range of 56.6 to 125.3 tn. Every other code together prices about 165,000 km for 11.0 tn.

This concentration has two consequences for the use of the costs. First, the total is almost entirely a statement about how many kilometres the required standard asks to pave, and at what price. A ministry that lowered the traffic thresholds for paving in `tpl_norm.design_standard`, or adopted low-volume seals where TI-03 now asks for surface dressing, would change the total far more than any refinement of the gravel or earth prices. Second, the rural-access works that close most of the G1 gap are cheap. The 57,141 km of engineered-earth and gravel upgrades cost 3.8 tn together, 4 % of the programme.

9.5.1 A worked example

The prices combine as follows for two typical links, at central values.

- **A 10 km community road of earth, required to be gravel.** The works cost 10 km × 75.1 m = 751 m; design and supervision (8 %) and contingency (15 %) bring the capital to 924 m UGX, with a range of 519 m to 2,385 m from the low and high prices. The road was already maintainable, and earth and gravel share the unpaved routine rate, so induced routine maintenance is zero. Had it been a track, outside the maintainable network, it would add 10 × 4.4 m = 44 m a year.
- **A 10 km gravel road, required to be surface dressing because of its traffic.** The capital is 10 × 4,269 m × 1.23 = 52.5 bn UGX (31.7 to 70.0 bn). Routine maintenance rises from the unpaved to the paved rate, inducing 10 × (16.4 − 4.4) m = 120 m a year for the life of the road.

The second link costs 57 times the first and induces maintenance where the first induces none. Both close a standard gap; whether either closes an access gap is decided in Chapter 10.

9.6 Data quality and limitations

- **Every price is an evidence candidate (E5).** The prices come from documents, not from an agreed national unit-cost standard. UCSN has not yet adopted them. No run using them can be presented as a budget estimate.
- **Some prices rest on one or two observations:** bridges, spot improvements, engineered earth and reseal. A single additional contract could move them by a large factor.
- **No terrain gradient** is applied, for lack of evidence. Mountain construction is likely to be under-priced and flat paving over-priced.
- **No paved-to-paved code** exists, so 2,499 km of sealed-to-asphalt upgrades are costed as new paving, about 13.1 tn central. This is the largest single known distortion in the programme cost.
- **Overheads and treatments are assumed** feasibility-stage rates and network treatments. Right-of-way cost is missing.
- **The 2026 indices are extrapolated** from the 2020–2025 trend.
- **Routine rates do not vary with traffic** in the draft norms, and periodic maintenance is not in the annual requirement.

- **Surface drives cost.** The maintenance requirement and every upgrading cost depend on the recorded surface; on the 169,343 km of unknown surface, cost is as uncertain as access.

9.7 Implications for planning

- **Treat every cost in this report as a range.** The condition backlog alone runs from 4.65 to 9.22 tn; the programme of Chapter 10 from 56.6 to 125.3 tn.
- **Ugandan evidence makes rural access works far cheaper than the regional benchmark.** Gravel construction from CRIP, PRELNOR, CAIIP and NDP III evidence is 15 % of UCSN's illustrative price; engineered earth and earth-road rehabilitation are priced for the first time.
- **A paved-to-paved resurfacing code is urgently needed.** Without it, 13.1 tn of the programme is mis-costed; UCSN should be asked for overlay and resurfacing prices, by terrain.
- **Bridge, culvert and spot-improvement prices need more evidence.** They rest on one observation each, and spot improvements are the cheapest instrument against wet-season closures.
- **Maintenance is the recurrent claim the plan must honour first.** The existing network needs about 955 bn a year in routine maintenance, before any new asset is added.
- **The community-road backlog is the largest condition liability.** 5.6 tn of the 7.3 tn backlog is on community access roads, whose condition the agencies report only in aggregate.
- **The evidence file should go to UCSN for adoption.** Once UCSN publishes estimates rather than illustrative values, runs can leave exploratory mode and the costs can enter budget documents.

CHAPTER 10

10 The investment programme

This chapter turns the gaps of Chapter 8 and the prices of Chapter 9 into a programme of candidate projects. It answers the question *what must be maintained, rehabilitated, upgraded and built, and in what order of logic?* It does not decide what is funded; that is the allocation of Chapter 11. The chapter implements the TI-07 module, design chapter 9 (`docs/09_TI07_INVESTMENT.md`), with the equations of `docs/23_MATHEMATICAL_MODELS.md`. All figures come from the reference run `014c9cd7`, in constant 2026 UGX at evidence-candidate prices (tier E5). The run is exploratory, and **every capital figure in this chapter is given with the routine maintenance it induces**, because a road built without its maintenance is lost within a few years.

10.1 What the module does

TI-07 reads every link with a gap and decides what works would close it.

- **Inputs.** The G2 and G4 gap links of TI-05; the required standard of every link from TI-03; the works costs, induced maintenance and phasing of TI-06; the agency condition statistics; the appraisal parameters of `tpl_norm.appraisal`; and the settlement access results, to measure what each project would do for people.
- **What it computes.**
  - the works type and target surface of each gap link, by an ordered decision rule, with the rule that fired recorded as the link's *rationale*;
  - a list of gap links whose surface is unknown, as a **survey requirement** instead of works;
  - **projects**: contiguous links with the same works, target surface, district and class, packaged up to 60 km;
  - **condition programmes** for the poor kilometres reported by the agencies;
  - each project's capital (low, central, high), induced routine maintenance, lifecycle cost, access gain, sector tags and economic appraisal (net present value, internal rate of return, benefit-cost ratio);
  - the **coverage sets** that tell the optimiser which projects, together, bring each unserved settlement within its standard.
- **Outputs.** The candidate programme `tables/ti07_programme.csv` (41,890 projects, with road numbers where OSM has them), the survey requirement `tables/ti07_survey_requirement.csv`, the coverage sets in the store, and the district layers used in this chapter's maps.

10.2 Method

10.2.1 The ordered decision rule

For each gap link, the rules are tested in order; the first that matches sets the primary works type. The order is data (`tpl_norm.investment_rule`), so a ministry can reorder it:

1. **Routine maintenance (RM)** applies to every maintainable link, every year (always).
2. **Periodic maintenance (PM)** if the condition is fair and a cycle is due (preserve).
3. **Spot improvement (SI)** if the link closes in the wet season at a crossing and is otherwise adequate (the cheapest access fix).

4. **Rehabilitation (RH)** if the condition is poor or bad and the standard is adequate (restore).
5. **Upgrading (UP)** to the required standard if the existing standard is below it (improve).
6. **Structure works, capacity expansion, new construction, climate adaptation** (rules 6 to 9), which need structure ratings, confirmed capacity or candidate alignments.

Formally, with predicates *π(r)* and works types *w(r)* for rules *r*, the primary works of link *l* is the first rule after routine maintenance whose predicate holds:

$$w^{prim}(l) = w_{r^*}, \qquad r^* = \min\left\{r \geq 2 : \pi_r(l) = 1\right\} \tag{10.1}$$

Spot improvement comes before upgrading on purpose. A vented ford that ends a wet-season closure often delivers most of the access an upgrade would, at a small fraction of its cost; the upgrade remains a separate, later candidate.

Three features of the Uganda data decide which rules can fire.

- **Unknown surface is a survey, not works.** A gap link whose surface is unknown generates no works; it enters the survey requirement. Upgrading a road whose surface nobody has recorded could pave a road that is already paved.
- **Condition is unknown on every link,** so the rehabilitation rule (4) and the periodic maintenance rule (2) never fire at link level. Condition works enter only as network-level programmes.
- **No candidate alignments, structure ratings or confirmed capacities** were supplied, so rules 6 to 9 do not fire.

10.2.2 Packaging and condition programmes

Links with the same works type, target surface, district and functional class that meet at a vertex are joined into one project, cut into packages of at most 60 km (an assumed procurement size). Each project lists the OSM road numbers of its links, the only non-geometric key for matching it to an agency's road register.

Condition programmes are added from the agency statistics (Chapter 8): the poor kilometres of each non-urban network, treated as the draft norm assigns (rehabilitation for national paved and community roads, regravelling for national unpaved and district roads), cut into 1,000 km tranches. A programme has no link geometry, because the poor kilometres are not located; its district is `NATIONAL`, it induces no maintenance because it restores an existing standard, and it has no access benefit or NPV.

10.2.3 Coverage sets and access gain

For each unserved settlement and destination class, TI-07 finds the small sets of projects (at most three projects per set, at most three sets, the cheapest retained) that, completed together, would bring the settlement within its standard:

$$\Omega_{i,k} = \left\{\omega \subseteq \mathcal{J} : A_{i,k}(G \oplus \omega) = 1,\ |\omega| \leq 3\right\} \tag{10.2}$$

where $G \oplus \omega$ is the network with the projects of ω applied. The **stand-alone access gain** of a project is the population of the settlement-class groups that it completes on its own (a set of one project). The reference run holds 144,746 settlement-class groups with 566,392 enabling sets. Stand-alone gains overlap between projects and are counted once per destination class, so they are never added across projects or classes; the optimiser of Chapter 11 counts people exactly from the sets.

10.2.4 Economic appraisal

Each priced link project is appraised over 20 years at an assumed 12 % discount rate. Monetised benefits are the road-user savings, evaluated on the network:

- **Travel time:** for each link, the base-year traffic times the time saved, valued at UGX 3,000 per person-hour and an average occupancy of 2.2 (both illustrative). A link closed in the wet season has an infinite current time; the run values the saving at three times the improved time, a flagged placeholder until a closure-cost norm exists.
- **Vehicle operating cost (VOC):** from roughness before and after. It needs a known condition, so in the reference run **no VOC saving is counted on any link**.

$$\Delta TTC_j = \sum_{l \in L_j} 365\, AADT_l \, \frac{t_l^0 - t_l^1}{60}\, occ \cdot VoT \tag{10.3}$$

Costs are the phased capital and the induced maintenance after completion. With benefits *B* starting after construction:

$$NPV_j = \sum_{t=1}^{20} \frac{B_{j,t} - X_j^{cap} s_{w,t} - IM_j \mathbf{1}(t > n_j)}{(1+r)^t}, \qquad BCR_j = \frac{PV(B_j)}{PV(C_j)} \tag{10.4}$$

The internal rate of return solves *NPV(r) = 0* by Brent's method on the bracket −50 % to 500 %, and is left blank when no root lies in it. Access, sector and resilience benefits are **reported beside the NPV and never converted into it**. Their cost-effectiveness is a ratio, not a score:

$$CE_{j,k} = \frac{PVC_j}{\Delta U_k(j)} \tag{10.5}$$

10.2.5 Sector tags

Each project is tagged with the sectors whose gaps it helps close, from the destination classes of its coverage sets: *rural access* (all-season road), *health* (primary health and EmONC), *education* (schools), *agriculture* (markets) and *admin* (district centre), plus the sector of any sector node whose connection standard binds on its links (industry, trade). Tags let Chapter 11 report what each allocation buys for each sector without a composite.

10.3 Results: the candidate programme

10.3.1 Size and composition

The decision rule assigned works to **212,849 links**: upgrading on 212,826 and spot improvement on 23. No link received rehabilitation, because no link has a known condition. A further **169,343 km of gap links with unknown surface** went to the survey requirement instead of works (Table 8.6, Chapter 8). Packaging and the condition programmes give **41,890 projects**, of which 41,867 are priced and 23 are not. Table 10.1 and Figure 10.1 give the programme by works type, each with its induced routine maintenance.

Table 10.1: The candidate programme: capital by works type, always with its induced routine maintenance

Works	Projects	Km	Capital low (bn UGX)	Central	High	Induced RM (bn UGX/yr)
Periodic maintenance	38	36,977	389.0	1,071.1	1,602.9	0.0
Rehabilitation	72	70,860	4,260.3	6,181.2	7,618.5	0.0
Upgrading	41,757	72,725	51,985.7	85,509.5	116,036.3	205.2

Evidence prices (E5 candidates), 2026 UGX; includes design, supervision and contingency. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_capital_works.csv).

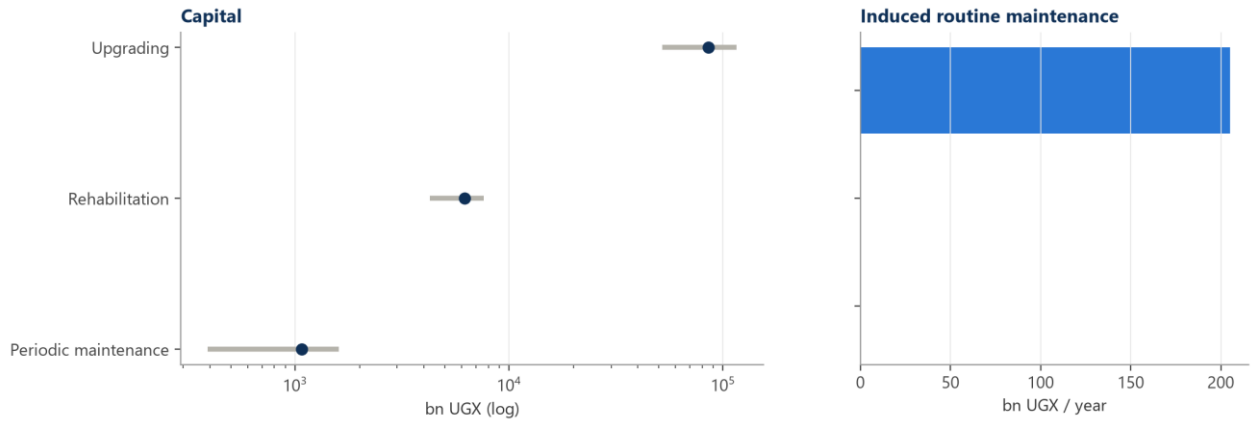

Figure 10.1: Capital by works type (low · central · high) beside its induced routine maintenance

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The priced programme costs 92.8 tn UGX central (56.6 to 125.3 tn) and induces 205.2 bn UGX a year of routine maintenance once complete.

- **Upgrading:** 41,757 projects on 72,725 km, 85.5 tn central (52.0 to 116.0 tn), inducing **205.2 bn a year**: all of the programme's induced maintenance.
- **Rehabilitation:** 72 condition-programme tranches on 70,860 km (70,656 km of community roads and 204 km of national paved roads), 6.18 tn (4.26 to 7.62 tn), inducing **zero**, because it restores roads to their existing standard.
- **Periodic maintenance:** 38 condition-programme tranches on 36,977 km (35,320 km of district roads and 1,657 km of national unpaved roads), 1.07 tn (0.39 to 1.60 tn), inducing **zero**.
- **Spot improvements:** 23 items on 39 km of link, **unpriced** (Chapter 9), so they carry no capital and cannot be funded. Twelve are in Kasese, two in Zombo, and one each in nine other districts.

Figure 10.2 shows the same programme in kilometres.

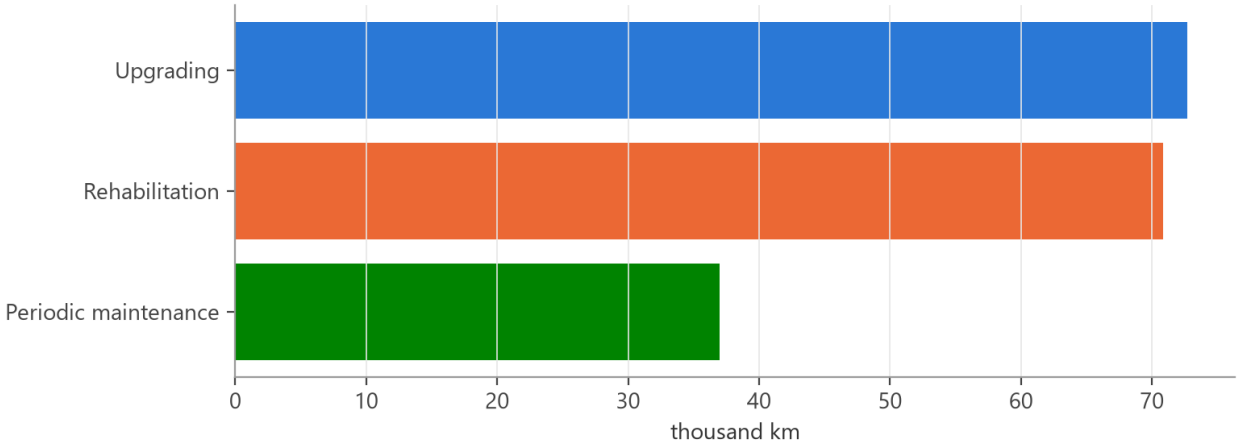

Figure 10.2: Km of works in the candidate programme, by works type

Source: AICTEM TIPRAS engine, reference run 014c9cd7

In kilometres the three priced works types are of similar size: 72.7 thousand km of upgrading, 70.9 thousand km of rehabilitation and 37.0 thousand km of periodic maintenance. In money they are not: upgrading is 92 % of the capital. A kilometre of condition programme costs 67 m UGX on average, a kilometre of upgrading 1.18 bn.

10.3.2 What the upgrading buys

The upgrading programme splits sharply by target surface (Table 10.2).

Table 10.2: Upgrading by target surface: capital beside induced routine maintenance, central, UGX 2026

Target surface	Projects	Km	Capital (bn UGX)	Induced RM (bn UGX/yr)
Asphalt concrete (ac)	1,650	10,765	56,525	100.9
Surface dressing (dbst)	1,602	4,797	25,192	57.6
Engineered earth	32,826	37,109	1,924	46.2
Gravel	5,602	20,032	1,850	0.1
Low-volume seal	77	21	19	0.2
Total	**41,757**	**72,725**	**85,510**	**205.2**

- **Paving** (asphalt and surface dressing) is 15,562 km and 81.7 tn of capital, inducing 158.5 bn a year. It is required mostly by traffic: projects whose binding test is traffic cost 82.9 tn in all.
- **Engineered earth and gravel** cover 57,141 km for 3.8 tn, inducing 46.3 bn a year. The induced maintenance of the engineered-earth upgrades comes from tracks that become maintainable roads; earth and gravel already share the unpaved routine rate.
- **2,499 km of the paving is surface dressing upgraded to asphalt**, costed at the full upgrading price (13.1 tn, inducing 6.0 bn a year) because the unit-cost map has no paved-to-paved code (Chapter 9). This is an over-statement of unknown size.

By binding test, 33,472 projects (39,691 km, 2.42 tn) are required by the link's function, 8,242 (32,987 km, 82.95 tn) by its traffic and 43 (46 km, 0.14 tn) by a sector node. The traffic-bound projects are few in number but carry almost all the money.

Projects are small. The median project is 0.43 km, a quarter are under 0.11 km, and the mean is 1.74 km; packaging joins links only when they meet at a vertex with the same works, target and class, and the local network is finely divided. Only 292 projects carry an OSM road number (national A, B and C roads and some named urban roads), together 6.2 tn. For the rest, matching to the agencies' road registers will need geometry.

10.3.3 Where the works are

Figure 10.3 maps the works the decision rule assigned to each link.

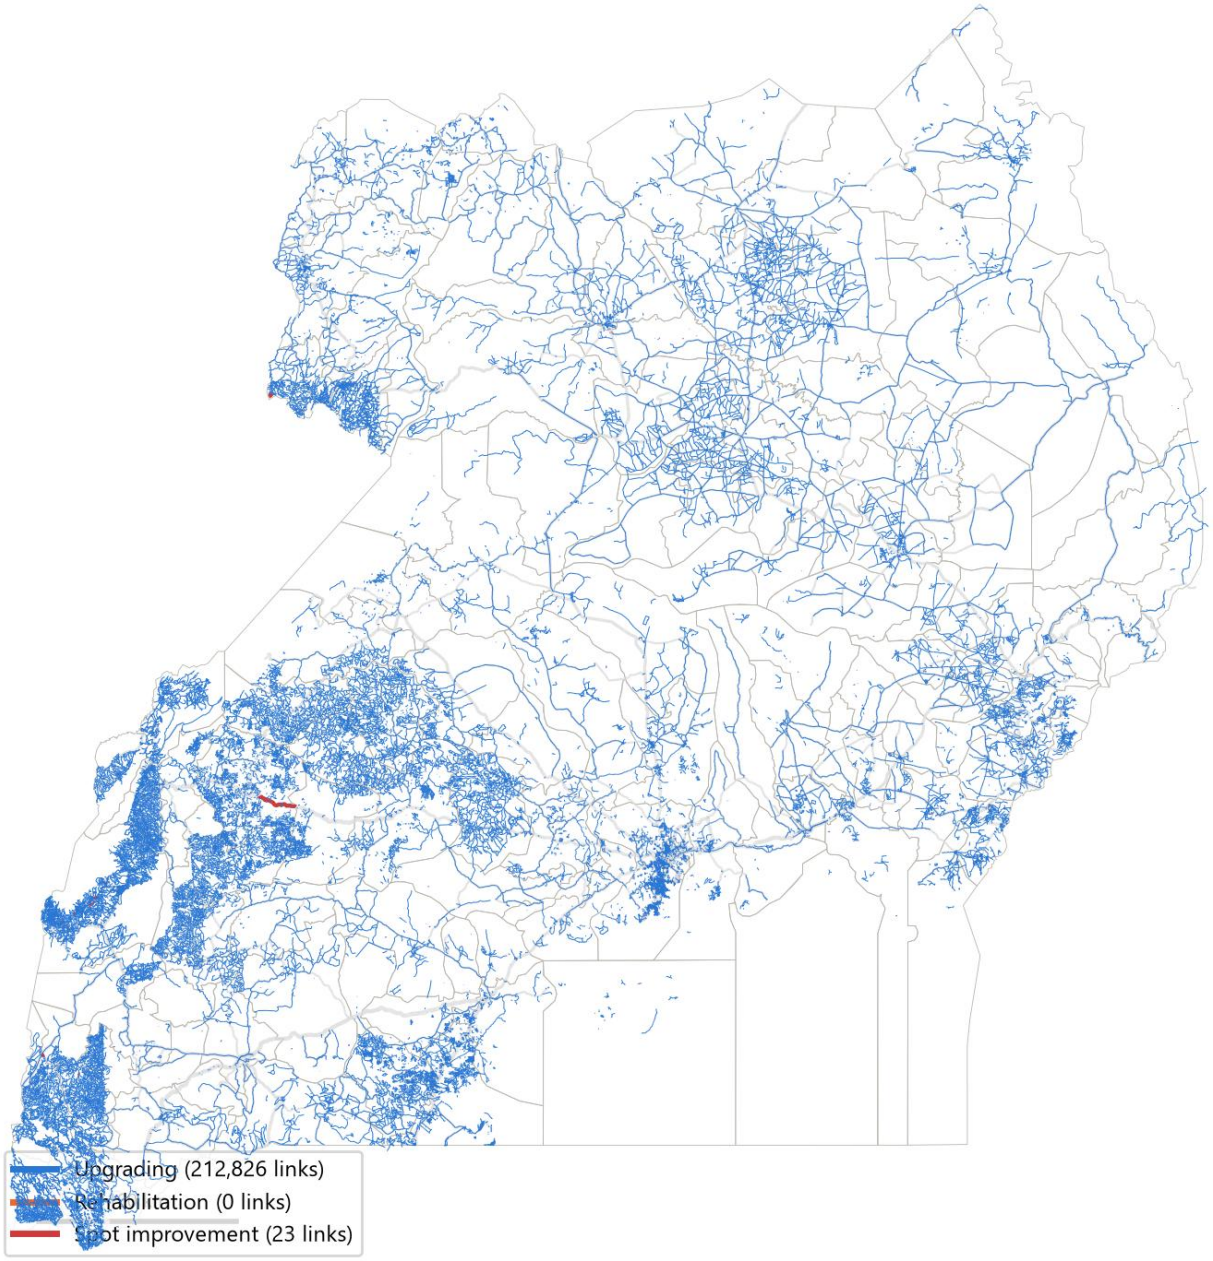

Figure 10.3: The works the decision rule assigns to each link (spot improvement, rehabilitation, upgrading)

Source: TI-07 ordered decision rule

The upgraded links, 212,826 in blue, are densest where the network is both long and *recorded*: the Kigezi highlands of the south-west (Kabale, Rubanda, Kisoro, Kanungu, Rukungiri), the Rwenzori slopes of Kasese, Bundibugyo and Kabarole, the mid-west block of Kibaale, Kakumiro, Kagadi and Kyenjojo, the Kampala metropolitan area, Lango, Elgon and West Nile. The large empty areas of the map, much of the centre and north, are not areas without gaps; they are where the gap links have unknown surface and went to the survey list instead. The legend confirms that no link received rehabilitation. The red spot-improvement links are few; the longest is a 21.9 km link in Kyegegwa.

10.3.4 Capital and induced maintenance by district

Figure 10.4 maps the link-based programme by district (85.5 tn; the condition programmes have no district) beside the routine maintenance it induces, and Figure 10.5 and Table 10.3 rank the largest district programmes.

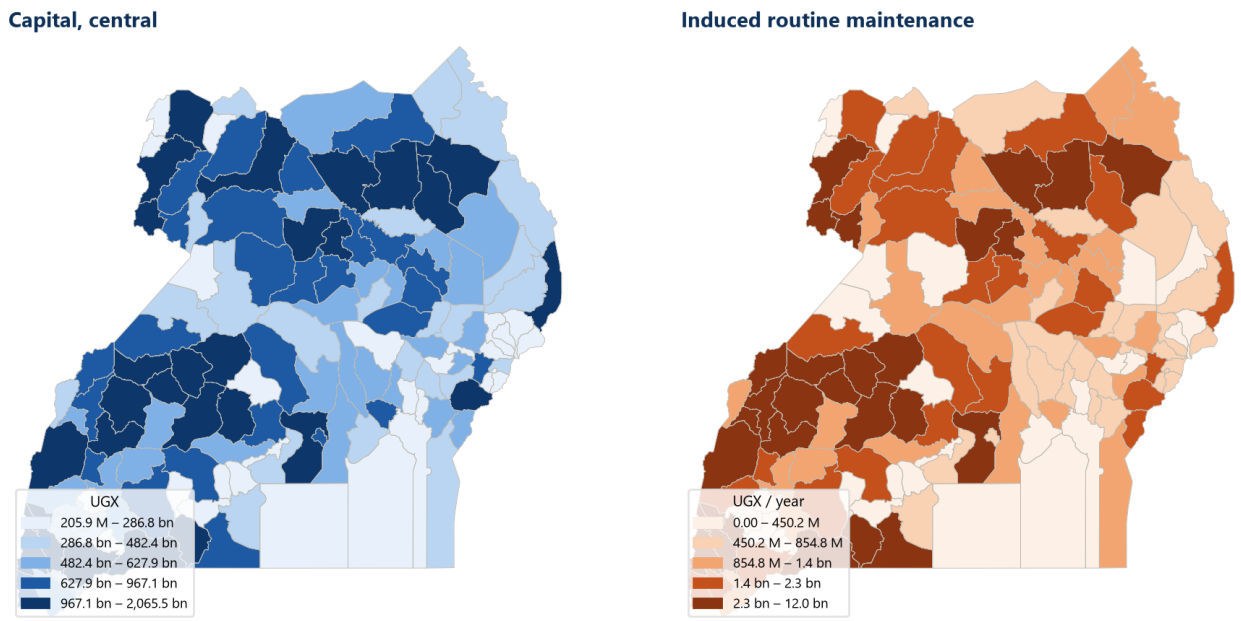

Figure 10.4: Candidate programme capital by district (left) and the routine maintenance it induces (right)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

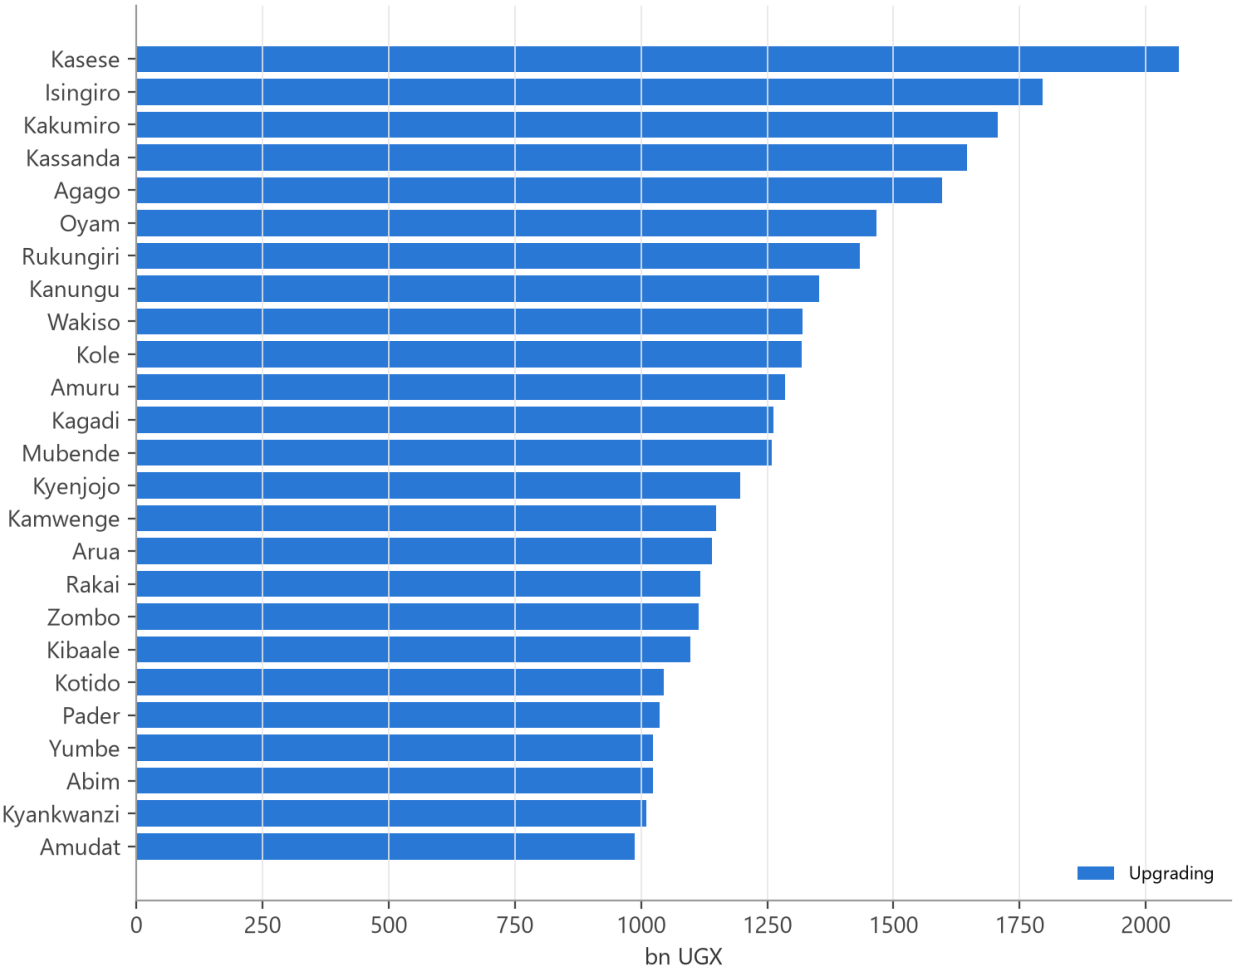

Figure 10.5: The 25 districts with the largest candidate programme, stacked by works type

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 10.3: The 15 districts with the largest candidate programme

District	Projects	Km	Capital (bn UGX)	Induced RM (bn UGX/yr)
Kasese	2,340	3,848	2,065.5	11.97
Isingiro	321	1,009	1,795.9	4.11
Kakumiro	889	1,837	1,706.0	3.97
Kassanda	805	1,944	1,646.5	4.24
Agago	262	1,168	1,596.0	3.75
Oyam	188	777	1,467.1	3.07
Rukungiri	1,053	1,907	1,432.8	5.85
Kanungu	1,319	2,044	1,353.5	6.27
Wakiso	5,297	1,619	1,320.0	3.81
Kole	175	753	1,317.7	2.65
Amuru	115	628	1,284.6	2.11
Kagadi	917	1,818	1,262.5	3.03
Mubende	724	1,404	1,258.8	3.54
Kyenjojo	2,394	2,935	1,195.6	3.08
Kamwenge	1,675	2,284	1,148.2	3.38

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_capital_top.csv).

- **Largest programmes:** Kasese (2,065 bn capital; 11.97 bn a year induced), Isingiro (1,796 bn; 4.11 bn), Kakumiro (1,706 bn; 3.97 bn), Kassanda (1,647 bn; 4.24 bn), Agago (1,596 bn; 3.75 bn), Oyam (1,467 bn; 3.07 bn), Rukungiri (1,433 bn; 5.85 bn), Kanungu (1,354 bn; 6.27 bn), Wakiso (1,320 bn; 3.81 bn) and Kole (1,318 bn; 2.65 bn). Every one of the top 25 is upgrading only.
- **Smallest:** Buvuma (0.2 bn on 4 km), Kalangala (4.3 bn), Kapchorwa (24.3 bn) and Lyantonde (59.3 bn), where the snapshot records little upgradable network.
- **Regions:** Northern 27.8 tn (60.9 bn a year induced), Western 26.7 tn (79.4 bn), Eastern 16.3 tn (29.8 bn) and Central 14.6 tn (35.1 bn).

The two panels of Figure 10.4 do not match. Capital is highest across the north (Agago, Kotido, Abim, Pader, Amuru, Oyam, Kole) and the mid-west, where long earth roads are required to be paved. Induced maintenance is highest in the south-west and Rwenzori (Kasese 12.0 bn, Kanungu 6.3 bn, Rukungiri 5.8 bn, Ntoroko 4.5 bn, Kassanda 4.2 bn, Nebbi 4.2 bn), where many tracks become maintainable engineered-earth roads. Kasese leads on both. Where the induced maintenance is high relative to the capital, the district's recurrent budget, not its development budget, is the binding constraint on the programme.

Figure 10.6 divides the link-based capital by each district's settlement population in the snapshot (42.2 million people, WorldPop base).

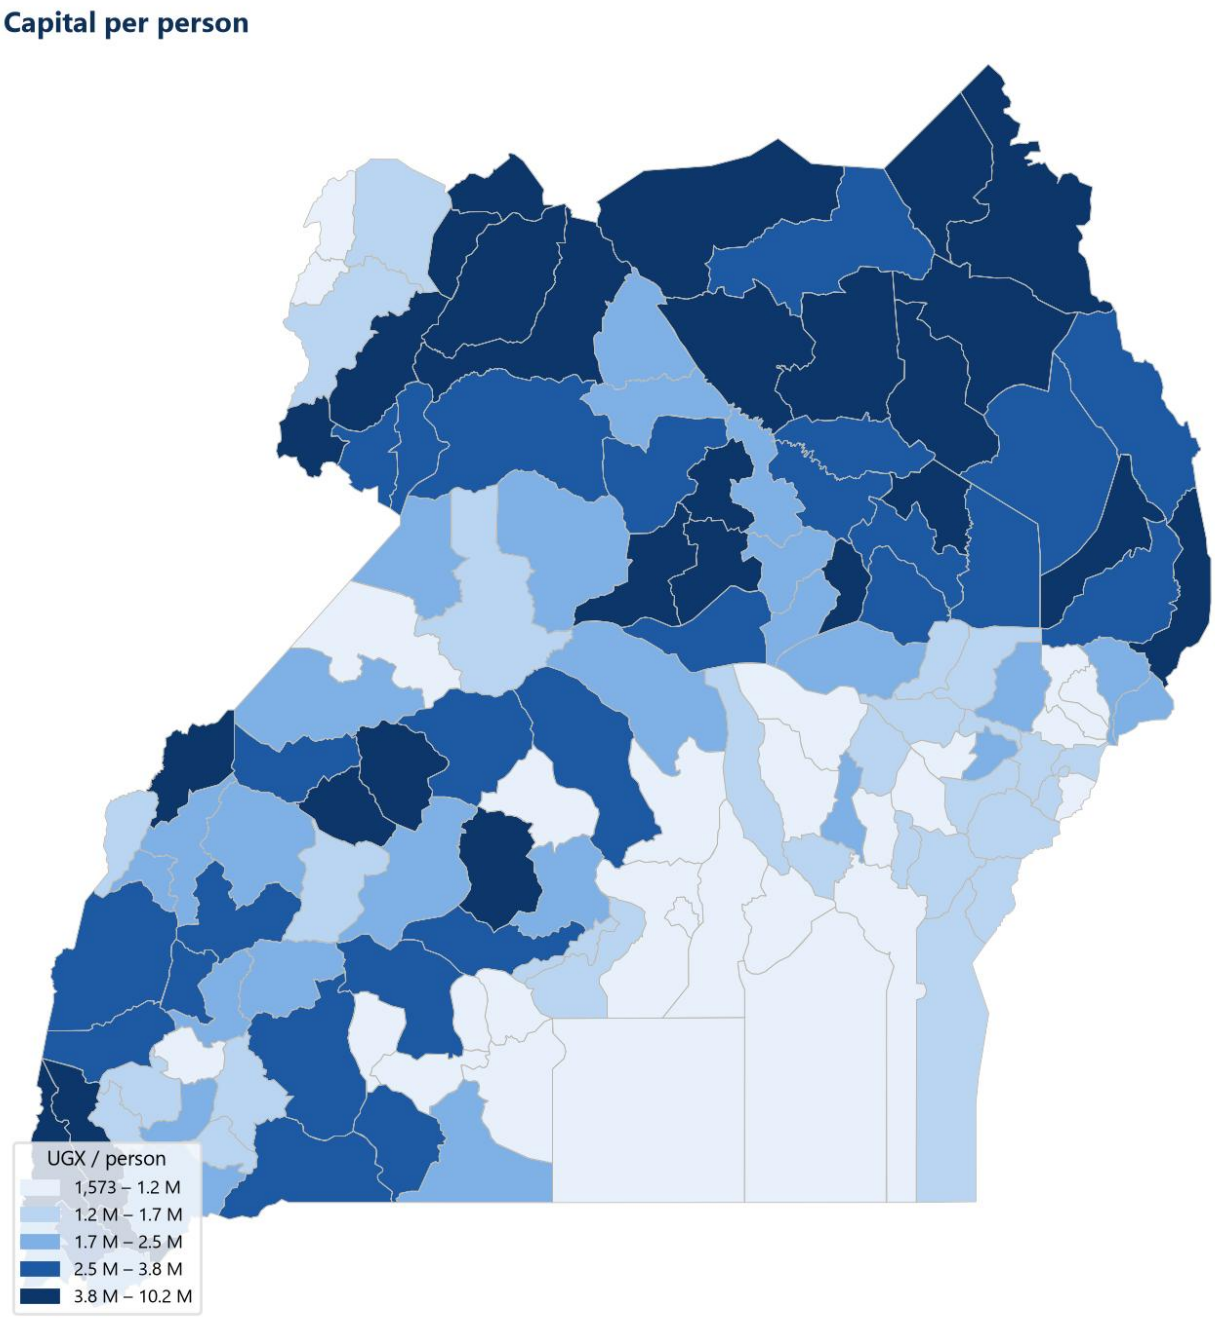

Figure 10.6: Candidate programme capital per inhabitant, by district

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Capital per person is highest in sparsely populated districts with long required upgrades: Karenga (10.2 m UGX per person), Ntoroko (8.7 m), Amudat (6.9 m), Abim (6.5 m), Agago (6.3 m), Amuru (5.8 m), Kibaale (5.6 m), Kapelebyong (5.5 m), Pader (5.1 m) and Obongi (4.9 m). It is lowest in the dense centre and east: Wakiso (443,000), Iganga (424,000), Mayuge (423,000), Kampala (374,000), Lwengo (330,000), Kapchorwa (193,000), Kalangala (65,000) and Buvuma (1,573). The national average is 2.0 m per person and the district median 2.2 m. The map is the mirror of population density: the north and north-east cost most per person to bring to standard, which is the equity trade-off Chapter 11 makes explicit.

10.4 Results: appraisal and cost-effectiveness

10.4.1 Economic return

Figure 10.7 shows the distribution of net present value across the 41,757 appraised upgrading projects.

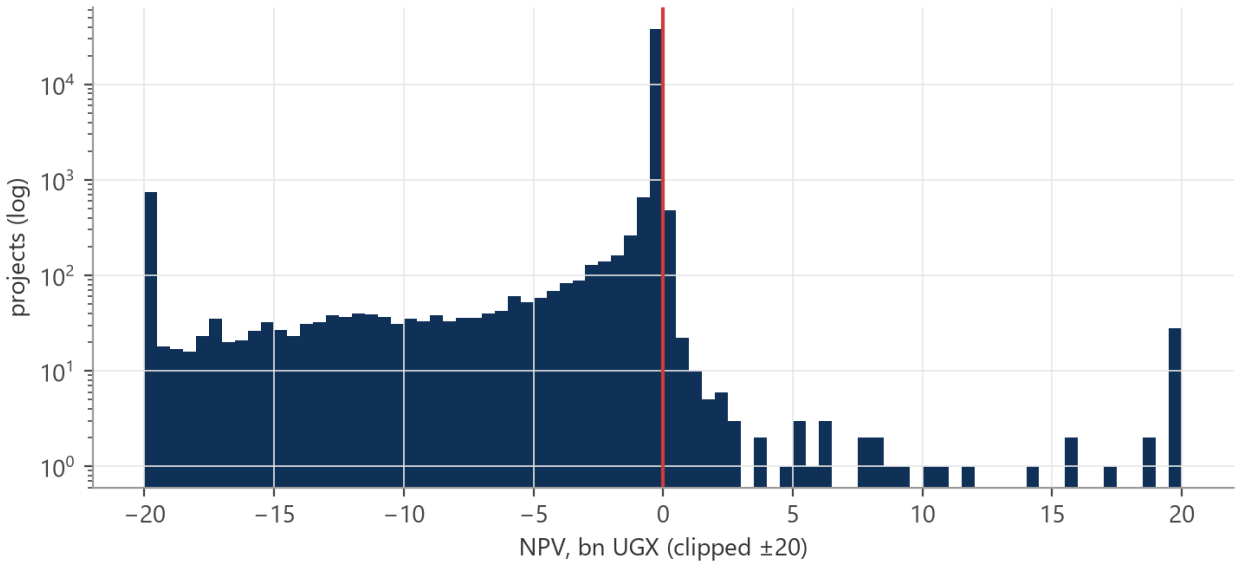

Figure 10.7: Appraisal: net present value per project at a 12 % discount rate (log count)

Source: TI-07 appraisal; VOC uncalibrated

574 projects have a positive NPV at 12 %; 41,183 do not. The distribution has a tall spike just below zero: the median NPV is −22 m UGX, 92 % of projects lie within ±0.5 bn of zero, and the 10th and 99th percentiles are −317 m and +17 m, because most projects are short and carry little traffic. A long tail of 734 projects lies at −20 bn or below (clipped in the figure): long paving projects whose time savings are small beside their cost. On the right, 28 projects exceed +20 bn.

- **The positive-NPV projects** cover 2,595 km and cost 3.64 tn central; their NPVs sum to 2,276 bn. Of the 574, 453 are upgrades to gravel, 114 to asphalt and 7 to engineered earth; 564 are traffic-bound.
- **The largest returns** are on short, heavily trafficked earth links required to be asphalt: a 2.0 km project in Kampala, on Bwaise roads (NPV 363 bn, BCR 44), and two links in Wakiso (1.2 km and 1.1 km, NPVs 128 bn and 89 bn). Longer positive projects include 35.6 km in Kikuube (NPV 222 bn, BCR 2.5), 31.7 km in Adjumani (103 bn, BCR 1.8), 60 km in Yumbe (100 bn, BCR 1.4) and 59.8 km in Kakumiro (94 bn, BCR 1.4).
- **Rates of return** could be solved for 14,700 projects; their median is −8.6 %.

The appraisal is exploratory and biased low. With condition unknown, it counts no vehicle-operating-cost saving, normally the largest benefit of upgrading a rough road. The VOC coefficients are uncalibrated TransportModelling values in assumed dollars, the value of time is illustrative, and the base-year traffic is modelled. A negative NPV here therefore means "time savings alone do not repay the cost", not "the project is uneconomic". The 3,000 UGX per person-hour value of time and the 12 % discount rate await MoFPED guidance.

10.4.2 Access gain against cost

Figure 10.8 plots, for each upgrading project with a stand-alone access gain, its capital against the people who gain access.

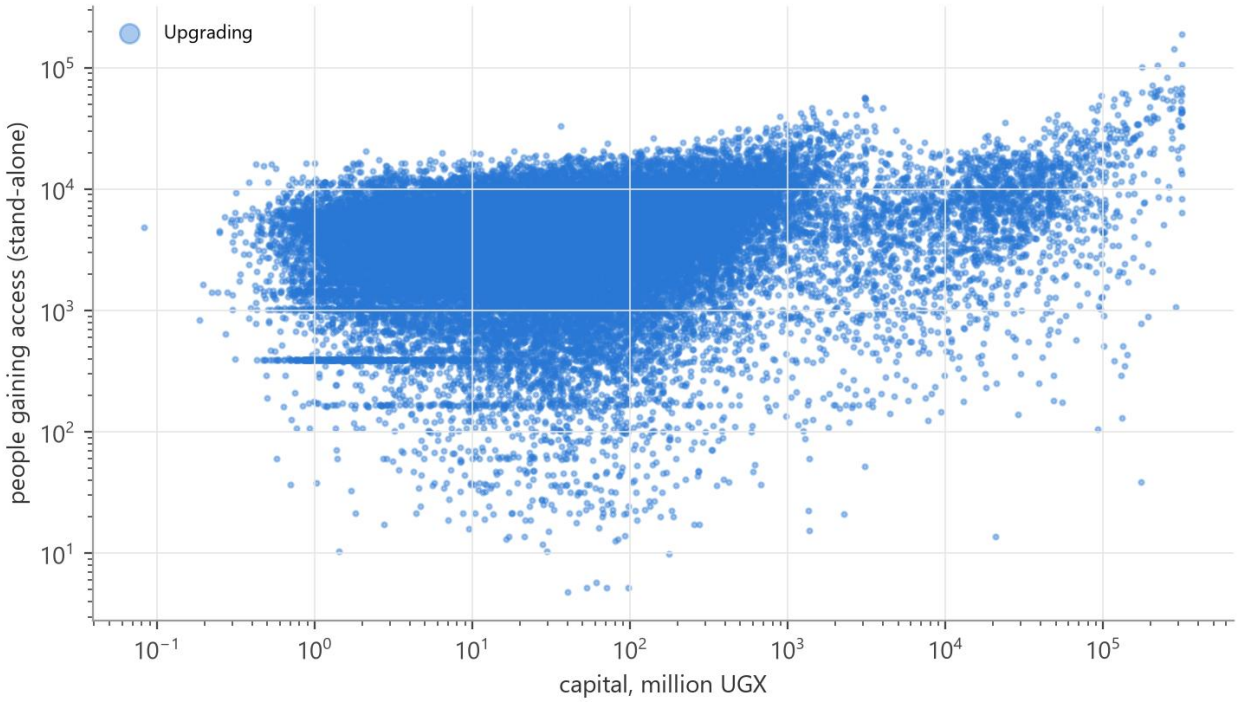

Figure 10.8: Capital against rural-access gain, project by project

Source: AICTEM TIPRAS engine, reference run 014c9cd7

35,444 projects have a stand-alone gain; 6,313 upgrading projects (6.4 tn) have none on their own, either because they help only in combination or because they close a standard gap with no access effect. The scatter spans six orders of magnitude in cost and four in people, and it has a clear structure.

- **A dense cloud of cheap projects** at 1 m to 1 bn UGX each serves 1,000 to 20,000 people each. These are the engineered-earth and gravel upgrades that close the rural-access gap. The horizontal lines near 400 and 160 people are settlement cells of identical population served by many small projects.
- **A second cloud of expensive projects** at 10 bn to 300 bn UGX serves a similar 1,000 to 60,000 people each. These are the paving projects.

Cost per person differs accordingly. Across all projects with a gain, the median capital per person gaining access is about UGX 10,000; the middle half lies between 2,400 and 42,000, and the most expensive tenth exceeds 204,000. By target surface, the ratio of summed capital to summed stand-alone gains is about UGX 15,000 per person for engineered earth, 51,000 for gravel, 2.4 m for surface dressing and 4.4 m for asphalt. These ratios are indicative, because the gains overlap, but the difference of two orders of magnitude between earth and gravel works and paving is robust. It is why the access-first scenarios of Chapter 11 fund thousands of small earth and gravel projects first.

Economic return and access gain point to different projects. The positive-NPV projects are short paving links in busy areas; the cheapest access is on long, lightly used earth roads with negative NPV. The engine reports both, side by side, and leaves the choice to the allocation.

10.4.3 Sector tags

Tags overlap, so their sums exceed the programme total and must be read one at a time. Projects tagged *rural access* are 35,437, with 78.6 tn of capital and 194.2 bn a year of induced maintenance. Projects that also bring people within a service standard are tagged *health* (3,295 projects; 51.6 tn; 99.7 bn a year), *agriculture* (3,782; 51.1 tn; 96.1 bn), *education* (3,089; 46.8 tn; 92.9 bn) and *admin* (1,858; 37.7 tn; 70.9 bn). Only 13 projects (0.04 tn; 0.1 bn a year) carry an *industry* tag and 7 (0.08 tn; 0.2 bn a year) a *trade* tag, because only a few sector-node connection standards bind on links whose surface is known. 6,130 projects, including the 110 condition programmes, carry no tag. Chapter 22 follows the tags into each sector's outbound exchange.

10.5 From gaps to works: what the programme closes and what it leaves

The seven gap dimensions of Chapter 8 are closed by different instruments, and the programme reaches them very unevenly (Table 10.4). Reading the programme against the gap register shows what the 92.8 tn would and would not buy.

Table 10.4: How the candidate programme answers each gap dimension in the reference run

Gap	Size in the register	Works in the programme	Left without works
G1 Access	33.3 m beyond an all-season road; 0.9–4.2 m per service class	upgrading of known-surface links, through the coverage sets	people behind unknown-surface links; the "no destination" share
G2 Condition	107,837 km poor (non-urban)	110 condition programmes, 7.25 tn, no induced maintenance	urban roads (31,854 km poor); location of every poor km
G3 Capacity	3,795 links over threshold	none	all; the capacity rule is not yet applied
G4 Standard, surface	242,068 km below required	upgrading of 72,725 km, 85.5 tn, 205.2 bn a year induced	169,343 km of unknown surface (survey)
G4 Standard, all-season	3,299 km on 5,013 links	upgrading of 2,789 closed links (1,984 km); 23 spot improvements (unpriced)	2,201 closed links of unknown surface (1,275 km)
G5 Connectivity	3,738 critical cut edges; 1,958 unrecorded crossings	none	all; no candidate alignments or crossings supplied
G6 Sector connection	150 of 178 nodes unconnected	upgrades on links whose binding test is a sector node (43 projects, 0.14 tn)	nodes reached only over unknown-surface links
G7 Resilience	3 exposed critical links	none	all; no climate-adaptation works are generated or priced

Three readings follow.

- **The programme is a standard-and-condition programme.** Upgrading closes the G4 surface gap on the links whose surface is recorded, and the condition programmes close the G2 gap in kilometres. Access gains come as a consequence: an upgrade raises a link to an all-season surface, which brings settlements within 2 km of an all-season road or within reach of a service.
- **The wet-season closures are mostly folded into upgrading.** Of the 5,013 closed links, 2,789 are earth links that also fall below their required surface, so the rule upgrades the whole link rather than fixing its crossing; only 23 links meet the spot-improvement rule (adequate surface, one closure point), and those are unpriced. The other 2,201 closed links have no recorded surface and wait for the survey. A rule that offered a spot improvement as a cheaper alternative on closed earth links, mutually exclusive with the upgrade, would let the allocation choose between them; the design allows such exclusions, and the run does not yet generate them.
- **Connectivity, capacity and resilience have no candidate works.** They need inputs the run did not have: alignments for new links and crossings from SDI or TransportModelling, confirmed capacities, structure ratings and a priced climate-adaptation works type. Their gaps remain on the register, reported in their own units, for the next cycle.

Each of these limits is a reason the programme should be read as the first, data-limited version of a plan, not its full extent. In money terms the missing parts are unknown: there is no price for a new crossing, a bridge per metre rests on one observation, and climate adaptation has no unit cost at all.

10.6 Data quality and limitations

- **The programme is what the recorded network allows.** 169,343 km of gap links are a survey requirement, not works; when surveyed, some will need upgrading and the programme will grow.

- **Condition is absent at link level,** so rehabilitation and periodic maintenance appear only as national programmes without location, and appraisal counts no VOC benefit.
- **New links, crossings, structures and capacity works are missing,** because alignments, structure ratings and confirmed capacities were not supplied. The "network missing" part of the G1 gap therefore has no candidate project.
- **The paving cost is overstated** on the 2,499 km of surface dressing upgraded to asphalt, and terrain does not vary the price.
- **Spot improvements are unpriced** on gravel and paved links, so the cheapest closure fixes cannot yet be funded.
- **Appraisal parameters are placeholders:** the discount rate, value of time, occupancy and the 3× value of closed-link time are all assumed; VOC is uncalibrated.
- **The committed pipeline and the ministry pipeline are not loaded,** so the programme is engine-generated only and the pipeline-versus-engine comparison of the design cannot be drawn.
- **All projects have the same earliest start (2027),** because design stages are not recorded.

10.7 Implications for planning

- **Every project carries its maintenance.** The full programme of 92.8 tn (56.6 to 125.3 tn) induces 205.2 bn a year of routine maintenance, on top of the existing network's 955 bn; any allocation must budget both.
- **Paving dominates the cost; earth and gravel dominate the access.** 81.7 tn of the capital paves 15,562 km, while 3.8 tn of engineered-earth and gravel works on 57,141 km does most of the rural-access work at about UGX 15,000 to 51,000 per person.
- **The required paving standard is the main lever on the programme's size.** Reviewing the traffic thresholds that require paving, and adding low-volume seal and overlay options to the map, would change the total more than any other decision.
- **Condition programmes are cheap per km and unlocated.** 107,837 km of poor district, community and national roads can be treated for 7.25 tn, but not mapped or sequenced until the agencies' condition data are available.
- **NPV and access rank different projects.** 574 projects repay their cost in time savings alone, mostly short busy links; the access gains are on lightly used roads. Neither ranking should be used alone.
- **The survey list is part of the programme.** The 169,343 km of unknown-surface gap links, and the 1,958 unrecorded crossings, should be surveyed before the next programme cycle so that their works can be priced and appraised.
- **Pricing spot improvements would unlock the cheapest closure fixes.** 23 spot-improvement projects, 12 of them in Kasese, wait only for a price.

CHAPTER 11

11 Resource allocation: where the money goes

The previous chapters established what Uganda's road network needs, what it has, where the gaps are and what closing them would cost. This chapter answers the fifth planning question: with a fixed envelope of money, split into pools that cannot be freely moved, which works should be funded, in which year, and in which districts? The TI-08 module answers it in two ways. First, it builds seven rule-based scenarios, A to G, each of which a minister can explain in one sentence. Second, it solves a mixed-integer programme (the optimiser) for one declared objective at a time, with any other aims turned into floors. Every allocation is reported side by side, with what it buys and what it gives up. None is recommended. The chapter implements design chapter 10 (`docs/10_TI08_ALLOCATION.md`) and the allocation sections of `docs/23_MATHEMATICAL_MODELS.md`. All results come from the reference run `014c9cd7` (geography `UGA-2026.09-c`, norms `UGA-2026.09-draft-f`, prices `UCSN+EV-UGA-2026.09b`), which is **exploratory**: the norms are drafts and the prices are evidence candidates (tier E5).

11.1 What the module does

TI-08 takes four inputs and produces one allocation per scenario.

- **The candidate programme from TI-07.** The reference run holds 41,890 projects, of which 41,867 are priced and 23 are not. Each priced project carries its works type, its capital cost (low, central, high), its phasing profile, its induced routine maintenance, its sector tags and its appraisal (NPV, EIRR). Chapter 10 describes the programme: 92.8 tn UGX of capital at central prices, with 205.2 bn UGX a year of induced routine maintenance.
- **The coverage sets from TI-07.** For each settlement and destination class, TI-07 records the small sets of projects that, once all completed, make the settlement meet its access standard. The reference run holds 144,746 settlement-class groups with 566,392 enabling sets. These sets let the optimiser count people served exactly, without re-routing the network inside the solver.
- **The maintenance requirement from TI-06.** Routine maintenance (RM) on the maintainable network is 954.7 bn UGX a year. Every funded capital item adds its own induced maintenance.
- **The envelope.** The money available, by year and by pool, built from the medium-term ceiling and the approved budget evidence (described below).

For each scenario TI-08 writes:

- every funded project, its start year and the pool that pays for each year's cash (`tables/ti08_allocations.csv`, 299,016 rows in the reference run);
- the money unallocated, by pool and year, with its reason;
- district allocation layers for the maps;
- for the optimiser runs, every measure's value at the solution and any floor shortfall (`tables/ti08_optimiser_measures.csv`), and the trade-off frontier (`tables/ti08_frontier.csv`).

TI-09 (Chapter 12) then applies each allocation to the network and recomputes the outcomes. TI-10 (Chapter 13) turns one of them into a budget.

11.2 Method

11.2.1 The envelope and its pools

The envelope is not one sum. It is split into **pools**, because money from different sources can legally fund different things. Road fund money is ring-fenced for maintenance; donor project loans fund named capital works; conditional grants go to district and community roads. Pools are not fungible: money left in one pool cannot be moved to cover a shortfall in another.

The envelope for pool π in year t is

$$B_{\pi,t} = C_t \cdot \rho \cdot \sigma_\pi \tag{11.1}$$

where:

- C_t is the final-round transport ceiling for fiscal year t published by MTEFModelling (baseline scenario): 12.27 tn UGX for FY2027/28, 12.28 tn for FY2028/29, 11.98 tn for FY2029/30 and 10.76 tn for FY2030/31;
- ρ = 0.545 is the roads share of transport, the FY2026/27 approved figure (roads 4,813.6 bn of the 8,827.2 bn Integrated Transport Infrastructure and Services programme);
- σ_{π} is the pool share, the three-year amount-weighted split of approved roads spending FY2024/25 to FY2026/27, arrears excluded.

Both ρ and σ_{π} are norms of quality *normed*, taken from `data/reference/budget_roads_evidence_UGA.md` (Chapter 13 describes that evidence). The pool shares and what each pool may fund are:

Table 11.1: Pools, their shares of the roads envelope and the works they may fund

Pool	Share	Envelope 2027–2030 (bn UGX)	Eligible works
`DONOR` (external project finance)	0.478	12,317.2	RH, UP, ST, NC, CE, CA
`GOV_DEV` (domestic development)	0.333	8,580.8	SI, RH, UP, ST, NC, CE, CA, PM
`DISTRICT_GRANT` (transfers to local governments)	0.088	2,267.6	RM, PM, SI, RH
`ROAD_FUND_MAINT` (national road maintenance)	0.080	2,061.5	RM, PM
`GOV_RECURRENT` (wage and non-wage recurrent)	0.021	541.1	RM
Total	1.000	**25,768.2**	

The works codes are those of Chapter 10: RM routine maintenance, PM periodic maintenance, SI spot improvement, RH rehabilitation, UP upgrading, ST structures, NC new construction, CE capacity expansion and CA climate adaptation. The eligibility table is itself a norm (`pool_eligibility`), not code.

The horizon is four plan years, 2027 to 2030, and the envelope totals 25.77 tn UGX. Per year it is 6,684.5 bn in 2027, 6,693.5 bn in 2028, 6,528.3 bn in 2029 and 5,861.9 bn in 2030. The 2027 figure is 39 % above the FY2026/27 approved roads budget of 4,813.6 bn, because the MTEFModelling ceiling for FY2027/28 is higher than the approved FY2026/27 programme. The envelope is therefore a modelled ceiling, not a commitment.

11.2.2 Maintenance first

Before any capital is allocated, every scenario except A funds routine maintenance on the whole maintainable network. The maintenance funding ratio is

$$mfr_t = \frac{\sum_{\pi \in M} B_{\pi,t}}{RM_t + \sum_j IM_{j,t}} \tag{11.2}$$

where M is the set of maintenance-eligible pools, RM_t is the routine maintenance requirement of the existing network and $IM_{j,t}$ is the induced maintenance of project j once it is in service. In the allocation step TI-08 draws the 954.7 bn a year in pool order: road fund first, then recurrent, then district grants. In 2027 the road fund pool (534.8 bn) and the recurrent pool (140.4 bn) together provide 675.2 bn. The remaining 279.5 bn comes from district grants, which leaves 308.7 bn of the 588.2 bn district-grant pool for district capital and periodic maintenance. Routine maintenance on the existing network therefore absorbs all of the road fund and recurrent pools and almost half of the district grants, every year, before any capital is considered. Chapter 13 shows the steady-state version of the same ratio, including induced maintenance, which is 0.58.

Scenario A reproduces current practice and does not apply the maintenance-first rule. As implemented, it funds no routine maintenance from the envelope. It uses 795.5 bn of the road-fund pool for periodic maintenance instead.

11.2.3 The rule-based scenarios

Seven scenarios are built before the optimiser. Each is a benchmark that the optimiser must match or beat on its own objective.

Table 11.2: The seven rule-based scenarios

Scenario	In one sentence	Rule as implemented
A Current trajectory	"Spend as we do now."	The manual's rule is to allocate by the last three years' shares. No historical allocation data are loaded, so as implemented A funds national-class projects first and then the rest in programme order; no maintenance-first rule
B Preserve first	"Keep what we have before we build more."	RM first; then RH and PM (including the network condition programmes PRG_*), then SI, then everything else; NPV breaks ties
C Rural access	"Connect the most people to an all-season road."	RM first; rank by rural persons newly within 2 km of an all-season road per unit cost
D Productive sectors	"Move farm, mine and factory output at least cost."	RM first; projects tagged agriculture, minerals or industry, by NPV
E Social services	"Get people to clinics and schools in all seasons."	RM first; rank by persons newly meeting the health and school standards per unit cost
F Regional equity	"Spend where access is worst."	RM first; district shares by unserved population and coverage; within districts, C's rule
G Climate resilience	"Protect the links we cannot lose."	RM first; adaptation of exposed critical links first, then by NPV

Scenarios B to G use a greedy, dependency-closed knapsack. At each step the project with the highest benefit per unit of present-value cost is funded, if its cash profile fits in the residual pool money of every year it needs:

$$\kappa_j = \frac{ben(cl(j))}{PVC(cl(j))} \tag{11.3}$$

where $cl(j)$ is project j together with any unfunded projects it depends on, ben is the scenario's benefit in its own unit (persons for C and E, NPV for D and G) and PVC is the present value of the cost. After each pick the benefits of projects that share coverage sets with it are updated, so two feeders serving the same village are not both credited with it. The benefit unit differs by scenario, and no scenario adds benefits of different units together. Scenario E adds persons gaining health access and persons gaining school access, which are the same unit.

Scenario F sets district envelopes first:

$$s_d = \frac{U_d(1-c_d)^{\alpha}}{\sum_{d'} U_{d'}(1-c_{d'})^{\alpha}} \tag{11.4}$$

where U_d is the unserved population of district d, c_d its current coverage and α an equity exponent between 0 (by need only) and 2 (strongly pro-poor). The reference run uses α = 1.

11.2.4 The optimiser

The optimiser is a mixed-integer linear programme (MILP) solved with JuMP and HiGHS. A binary variable $x_{j,t}$ is 1 if project j starts in year t. Each optimisation declares **one** objective in its native unit, and turns any other aim into a soft floor. This is the ε-constraint method. It replaces a weighted index, which the engine never builds. The model solved is

$$\max \quad \frac{f_0}{f_0^{max}} - \Lambda \sum_q r_q \tag{11.5}$$

$$f_q \geq v_q(1-r_q), \quad 0 \leq r_q \leq 1, \quad \Lambda = 10 \tag{11.6}$$

where:

- f_0 is the declared objective (for example, persons newly within 2 km of an all-season road) and f_0^{max} its attainable maximum, the optimum of that measure alone under the same envelope;
- f_q is the value of floor measure q, and v_q its target, written either in absolute units or as a fraction of the measure's own attainable maximum (for example `served:HEALTH_EMONC>=0.8max`);
- r_q is the relative shortfall on floor q; any positive value is written to the run log as a finding;
- Λ = 10 is a penalty large enough that any feasible way of meeting a floor is worth more to the solver than giving it up. Because the objective is normalised to lie between 0 and 1, the penalty does not depend on the objective's unit.

The constraints include the one-start rule, phased cash drawn from eligible pools, and the pool budgets after routine maintenance:

$$\sum_t x_{j,t} \leq 1, \qquad \sum_j f_{j,\pi,t} + h_{\pi,t} + u_{\pi,t} = B_{\pi,t} \tag{11.7}$$

where $f_{j,\pi,t}$ is project j's cash drawn from pool π in year t (zero if the pool cannot fund the project's works type), $h_{\pi,t}$ is maintenance funded from the pool and $u_{\pi,t} \ge 0$ is unallocated money. A start whose cash profile would run past the horizon is forbidden.

Access is counted through coverage variables. A coverage set ω is complete only if every project in it is complete, and a settlement group g is served only if at least one of its sets is complete:

$$z_\omega \leq y_j \quad \forall j \in \omega, \qquad s_g \leq \sum_{\omega \in \Omega_g} z_\omega, \qquad served_k = \sum_{g \in G_k} P_g s_g \tag{11.8}$$

Here y_j is 1 if project j is complete by the horizon, Ω_g the enabling sets of group g, P_g the group's population and G_k the groups of destination class k. This keeps the access objective exact: two parallel feeders that serve the same village count it once.

Three practical features keep the national problem solvable:

- **Candidate pre-screen.** For each measure involved, projects are ranked by score per unit cost and taken until their cost reaches 1.5 times the capital envelope. The lists are merged round-robin across measures and capped at 20,000 projects plus the warm-start projects. The score only screens; it never enters the objective.
- **Warm start.** The rule-based allocation that scores highest on the declared objective is given to the solver as a starting solution. The optimiser can therefore always reproduce it, and the test that the optimiser matches or beats every rule-based scenario on its own objective holds by construction as well as by search.
- **Tolerances.** Money is scaled to bn UGX inside the solver. HiGHS runs with a 1 % MIP gap and a 900-second limit per solve.

Dependencies, exclusions, contracting capacity and per-district access floors are specified in the manual (constraints C10–C12 and C15–C16) but are not yet in the model.

11.2.5 The four optimiser runs and the frontier

The Uganda configuration (`config/scenarios/optimiser_UGA.yml`) declares four optimisations and one frontier. Each is named by its objective and floors:

- **OPT**: maximise persons newly within 2 km of an all-season road; no floor.
- **OPT_RAI_EMONC80**: the same objective, with the floor that at least 80 % of the attainable gain in access to emergency obstetric care (EmONC) is kept.
- **OPT_RAI_COND50**: the same objective, with the floor that at least 50 % of the attainable poor-condition km are treated by the network condition programmes.
- **OPT_NPV_RAI50**: maximise the total NPV of funded projects, with the floor that at least 50 % of the attainable rural-access gain is kept.

The **frontier** maximises rural access with a condition floor at 0 %, 25 %, 50 %, 75 % and 100 % of the attainable poor km treated, and reports both measures at each point.

11.3 Results: the envelope and what is left of it

Figure 11.1 and Table 11.3 show the capital each allocation commits over the horizon, by pool. Routine maintenance (954.7 bn a year, 3,818.8 bn over four years) is drawn first in every scenario except A and is not shown in the figure.

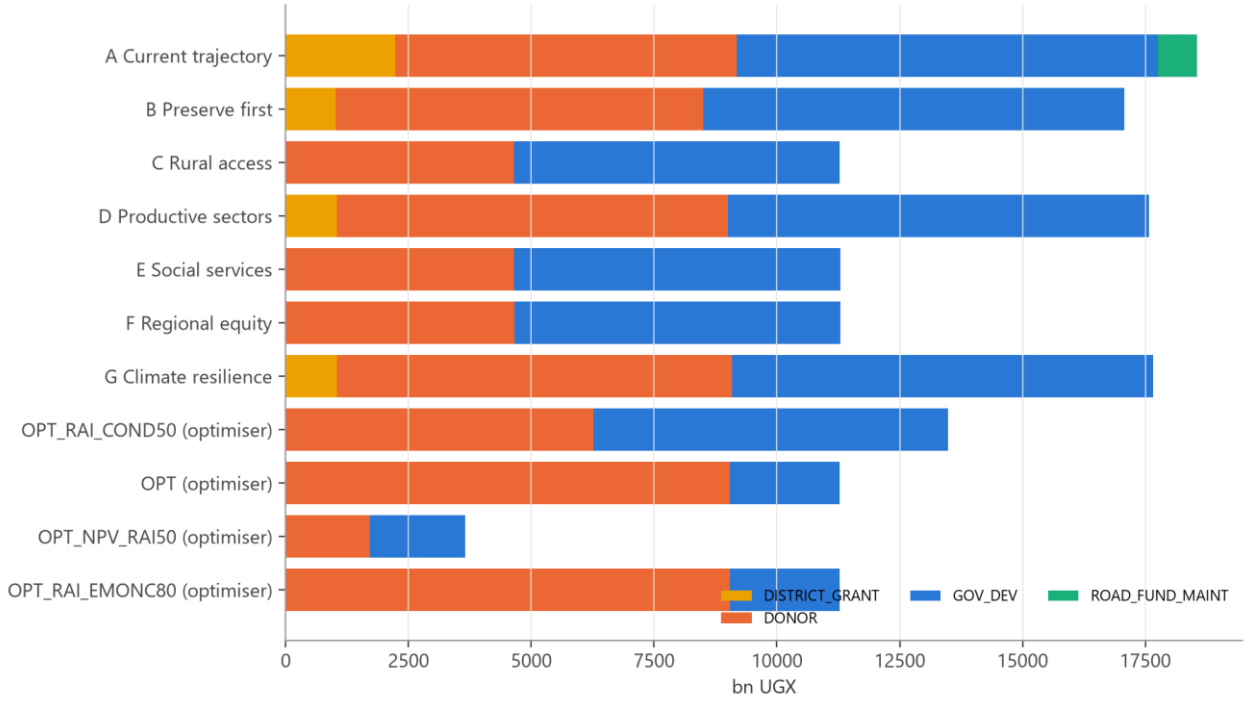

Figure 11.1: Capital allocated by funding pool in each scenario (horizon total)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 11.3: Capital allocated by funding pool and scenario, horizon total 2027–2030 (bn UGX)

Scenario	DISTRICT_GRANT	DONOR	GOV_DEV	ROAD_FUND_MAINT	Total (bn UGX)
A Current trajectory	2,229.0	6,948.9	8,580.8	795.5	18,554.2
B Preserve first	1,011.1	7,484.2	8,580.8	0.0	17,076.1
C Rural access	0.0	4,646.9	6,630.5	0.0	11,277.4
D Productive sectors	1,047.8	7,947.3	8,580.8	0.0	17,576.0
E Social services	0.0	4,651.1	6,635.9	0.0	11,287.0
F Regional equity	0.0	4,652.6	6,637.9	0.0	11,290.5
G Climate resilience	1,046.0	8,032.5	8,580.8	0.0	17,659.2
OPT_RAI_COND50 (optimiser)	0.0	6,268.0	7,215.9	0.0	13,483.9
OPT (optimiser)	0.0	9,048.5	2,228.9	0.0	11,277.4
OPT_NPV_RAI50 (optimiser)	0.0	1,713.3	1,945.6	0.0	3,658.9
OPT_RAI_EMONC80 (optimiser)	0.0	9,048.5	2,228.9	0.0	11,277.4

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_alloc_pool.csv).

Three groups of scenarios stand out.

- **The broad scenarios A, B, D and G commit 17.1 to 18.6 tn UGX of capital.** They exhaust the domestic development pool (8,580.8 bn in each) and draw 6.9 to 8.0 tn from the donor pool. They also use 1.0 to 2.2 tn of district grants, because they fund periodic maintenance and rehabilitation, which district grants may pay for. A is the largest at 18,554.2 bn, partly because it is the only scenario that does not fund routine maintenance first; it spends 795.5 bn of the road-fund pool on periodic maintenance instead.
- **The access scenarios C, E and F, and OPT, commit about 11.3 tn.** C commits 11,277.4 bn, E 11,287.0 bn and F 11,290.5 bn. They use only the development and donor pools, because almost every project they fund is an upgrade (UP), which district grants cannot fund.
- **The two optimiser runs with floors sit apart.** OPT_RAI_COND50 commits 13,483.9 bn; OPT_NPV_RAI50 commits only 3,658.9 bn.

OPT and C fund exactly the same 8,611 projects, with the same cash in every year. They differ only in which pool pays: OPT draws 9,048.5 bn from donors and 2,228.9 bn from domestic development, while C draws 4,646.9 bn and 6,630.5 bn. Both are feasible. The solver is indifferent between them because the objective counts people, not pools. A finance ministry is not indifferent: the OPT financing leaves about 4.4 tn more domestic development money unspent. The choice of pool is therefore a decision the allocation leaves open, not one it settles.

11.3.1 The unallocated money is a horizon effect

Every allocation leaves money unspent. The run log records, for the rule-based scenarios:

Table 11.4: Money left unallocated over 2027–2030, by scenario (bn UGX)

Scenario	Capital	Routine maintenance	Unallocated
A Current trajectory	18,554.2	0	7,214.0
B Preserve first	17,076.1	3,818.8	4,873.4
C Rural access	11,277.4	3,818.8	10,672.0
D Productive sectors	17,576.0	3,818.8	4,373.5
E Social services	11,287.0	3,818.8	10,662.4
F Regional equity	11,290.5	3,818.8	10,659.0
G Climate resilience	17,659.2	3,818.8	4,290.2

By the same arithmetic, OPT and OPT_RAI_EMONC80 leave 10,672.0 bn, OPT_RAI_COND50 about 8,465.6 bn and OPT_NPV_RAI50 about 18,290.6 bn.

This is **not** a pricing gap. In run 5 (illustrative UCSN prices), 32,849 projects were unpriced and could not be funded, and the leftover money reflected that. With evidence prices, 41,867 of 41,890 projects are priced. The leftover now has two causes, both tied to the four-year horizon.

- **C, E and F stop when no project adds access.** Once every project that brings a settlement within the standard has been funded, the remaining candidates have zero benefit under these rules. Their leftover of about 10.7 tn has nothing further to buy within the scenario's own objective.
- **B, D and G are limited by phasing.** These scenarios fund almost every priced project they can, B funding 40,486 of them. Their leftover of 4.3 to 4.9 tn is mostly donor money in 2027 and 2030. Works use an S-curve: a three-year project draws about a quarter of its cost in its first year, about half in its second and a quarter in its third. So a project started in 2027 puts little cash into 2027, and a three-year project cannot start late enough to finish by 2030. In scenario B, for example, the donor pool provides 3,195.2 bn in 2027 but only 695.8 bn is drawn; in 2030 it provides 2,802.0 bn and 494.2 bn is drawn.

Allowing commitments that complete after the horizon would absorb most of this money. That is a modelling choice, not evidence that the money is not needed.

11.4 Results: preserve or expand?

Figure 11.2 splits each scenario's capital by works type. It shows the central choice of the plan: maintaining and rehabilitating the existing network, or upgrading it.

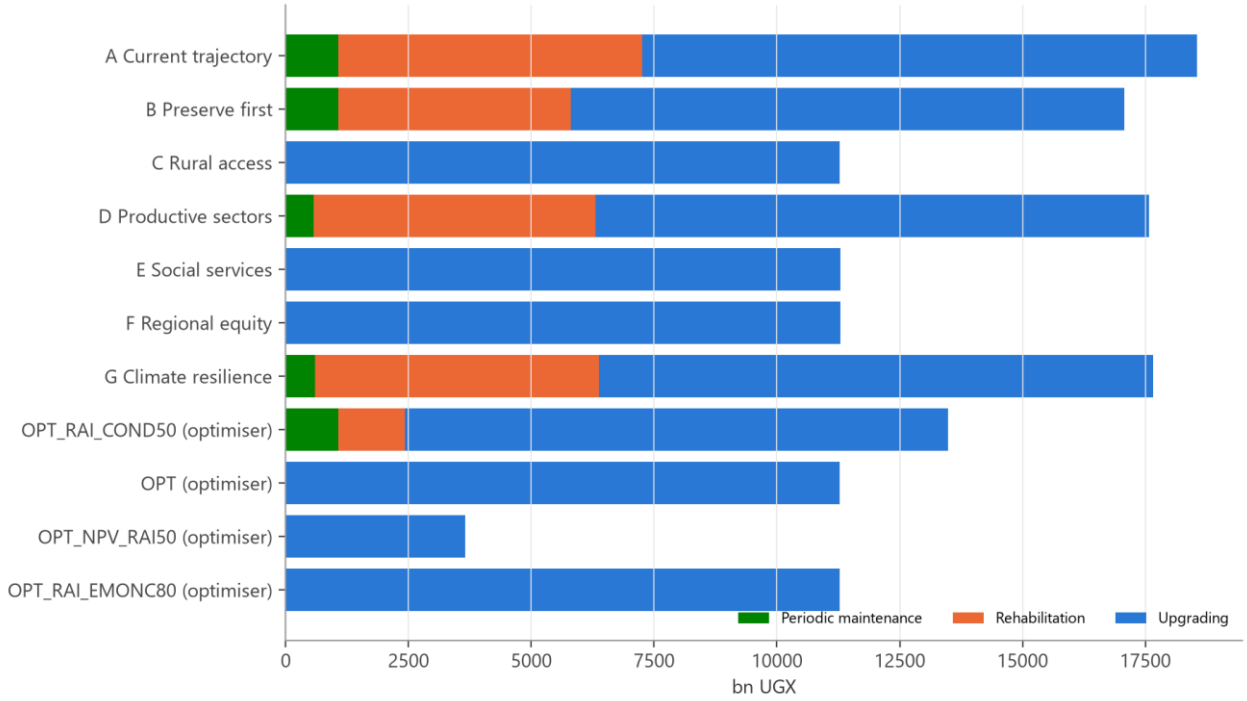

Figure 11.2: Capital allocated by works type in each scenario: preserve or expand?

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **Upgrading dominates every scenario.** Every scenario spends 11.06 to 11.30 tn UGX on upgrading (UP), except OPT_NPV_RAI50, which spends 3.66 tn. Upgrading is 85.5 tn of the 92.8 tn programme, so any allocation of a 25.8 tn envelope is mostly an upgrading allocation.
- **The condition programmes are what separates the scenarios.** A funds all of them: 1,071.1 bn of periodic maintenance and 6,181.2 bn of rehabilitation. B funds all the periodic maintenance and 4,735.4 bn of rehabilitation. D and G fund 0.57 to 0.59 tn of periodic maintenance and 5.73 to 5.78 tn of rehabilitation. C, E, F and OPT fund none: rehabilitating a road that is already all-season does not bring anyone new within 2 km of one.
- **OPT_RAI_COND50 is the compromise the floor forces.** It funds 1,071.1 bn of periodic maintenance and 1,349.4 bn of rehabilitation beside 11,063.3 bn of upgrading, and treats 53,977 km of poor road.

Every one of these capital figures carries induced routine maintenance. Over the whole programme, induced maintenance is 205.2 bn UGX a year at central prices. For scenario B, the assets funded add 62.7 bn a year to the routine maintenance requirement once in service (Chapter 13).

11.4.1 The condition programmes rest on agency statistics

The rehabilitation and periodic maintenance in these figures are network-level programmes (`PRG_*`) built from agency statistics of poor km by network, in tranches. Condition is not known for any individual link: no link-level survey exists in the estate. The engine therefore does not invent link condition. It funds condition programmes against the agency totals (107,837 poor km outside urban areas) and reports km treated, not a projected condition per link. Chapter 12 reports the poor km left untreated in each scenario.

11.5 Results: how many projects, and when

Figure 11.3 counts the projects funded.

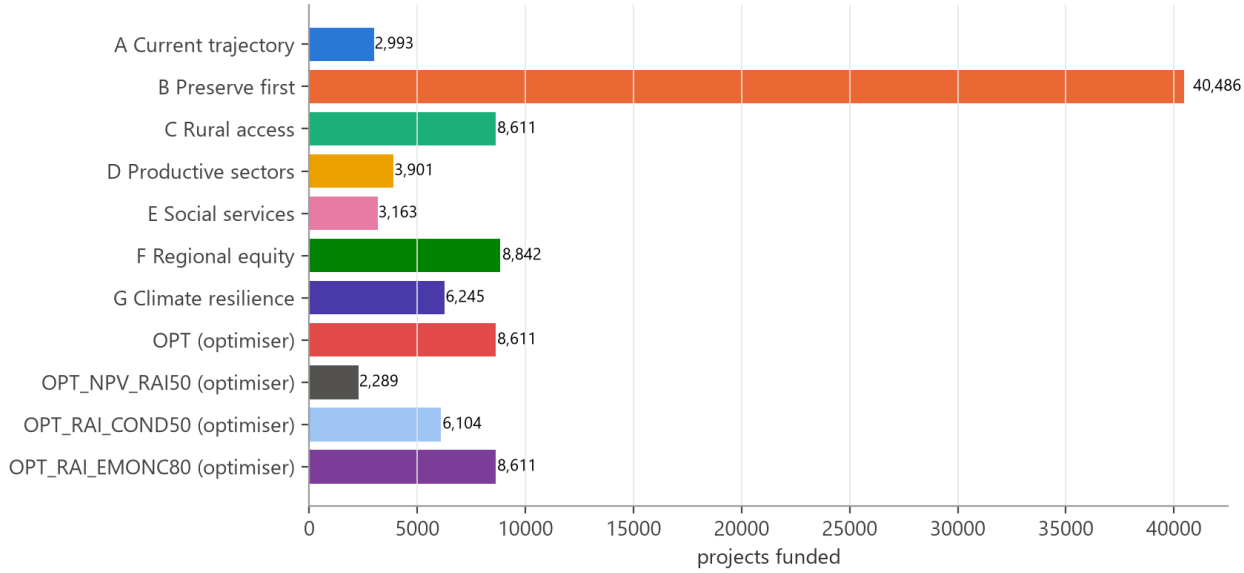

Figure 11.3: Projects funded in each scenario

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The counts differ by a factor of 18:

- B funds **40,486** projects, almost the whole priced programme. Its rule ranks every rehabilitation and maintenance programme first and then works down the upgrades by NPV.
- F funds 8,842; C and OPT 8,611; OPT_RAI_EMONC80 the same 8,611; G 6,245; OPT_RAI_COND50 6,104.
- D funds 3,901, E 3,163 and A 2,993.
- OPT_NPV_RAI50 funds 2,289, the fewest. Only 574 projects have a positive NPV at the 12 % discount rate, so maximising NPV quickly runs out of projects that add value; it funds the rest only to meet its access floor.

The count is not a measure of value. A and B commit similar capital (18.6 and 17.1 tn) but B funds 13 times as many projects, because its ranking reaches deep into the long list of small upgrades once the condition programmes are funded. The same capital buys very different numbers of projects.

Figure 11.4 shows when the money is drawn.

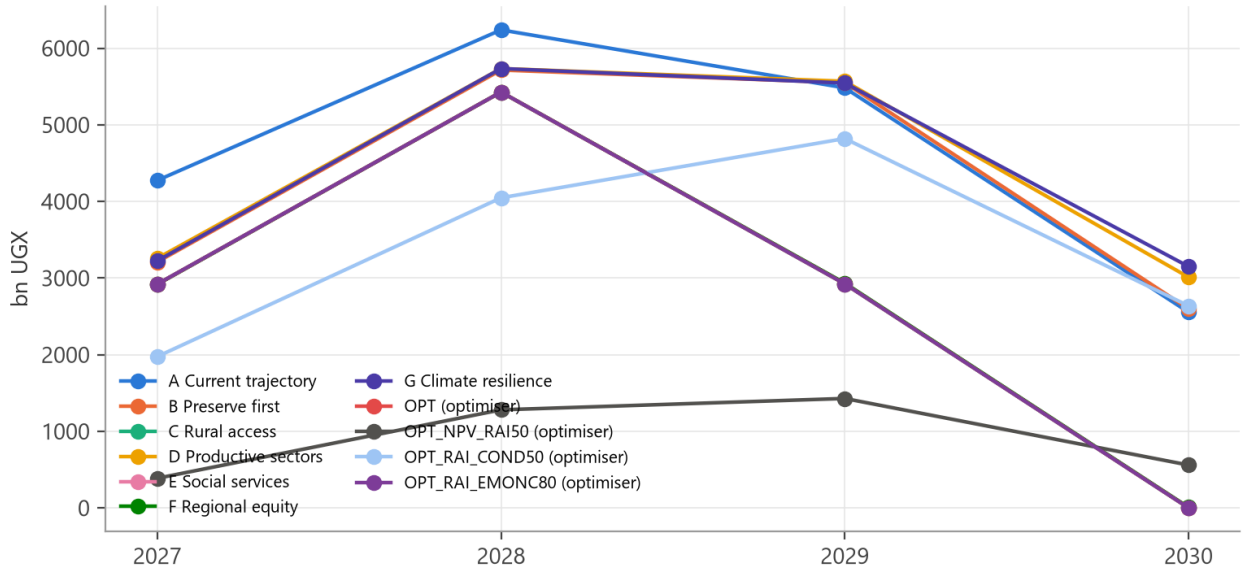

Figure 11.4: Capital disbursed by year within the horizon, by scenario

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **Almost all projects start in 2027.** In C, 8,609 of 8,611 projects start in 2027; in B, 40,399 of 40,486. Starting later would put cash beyond the horizon.
- **Cash peaks in 2028.** C draws 2,922.0 bn in 2027, 5,428.4 bn in 2028, 2,925.3 bn in 2029 and only 1.7 bn in 2030. The broad scenarios draw 5.5 to 5.7 tn in 2028 and 2.6 to 3.2 tn in 2030, because they include longer programmes.
- **OPT_RAI_COND50 is the most even**, with 1,975.9 bn, 4,048.1 bn, 4,821.6 bn and 2,638.2 bn. Its condition floor pulls in programmes that start in 2028 (3,431 of its 6,104 projects).
- **The 2030 trough of the access scenarios** is the horizon effect again: nothing started after 2027 can finish in time.

A real programme would not front-load in this way. Contracting capacity would limit the km started in one year, but that constraint (C12) is not yet in the model. The profile should be read as the cash requirement of the selected portfolio if it could all be procured at once.

11.6 Results: where the money goes

Figure 11.5 maps the allocation per inhabitant by district for every scenario on one shared colour scale. The same colour means the same UGX per person in every panel.

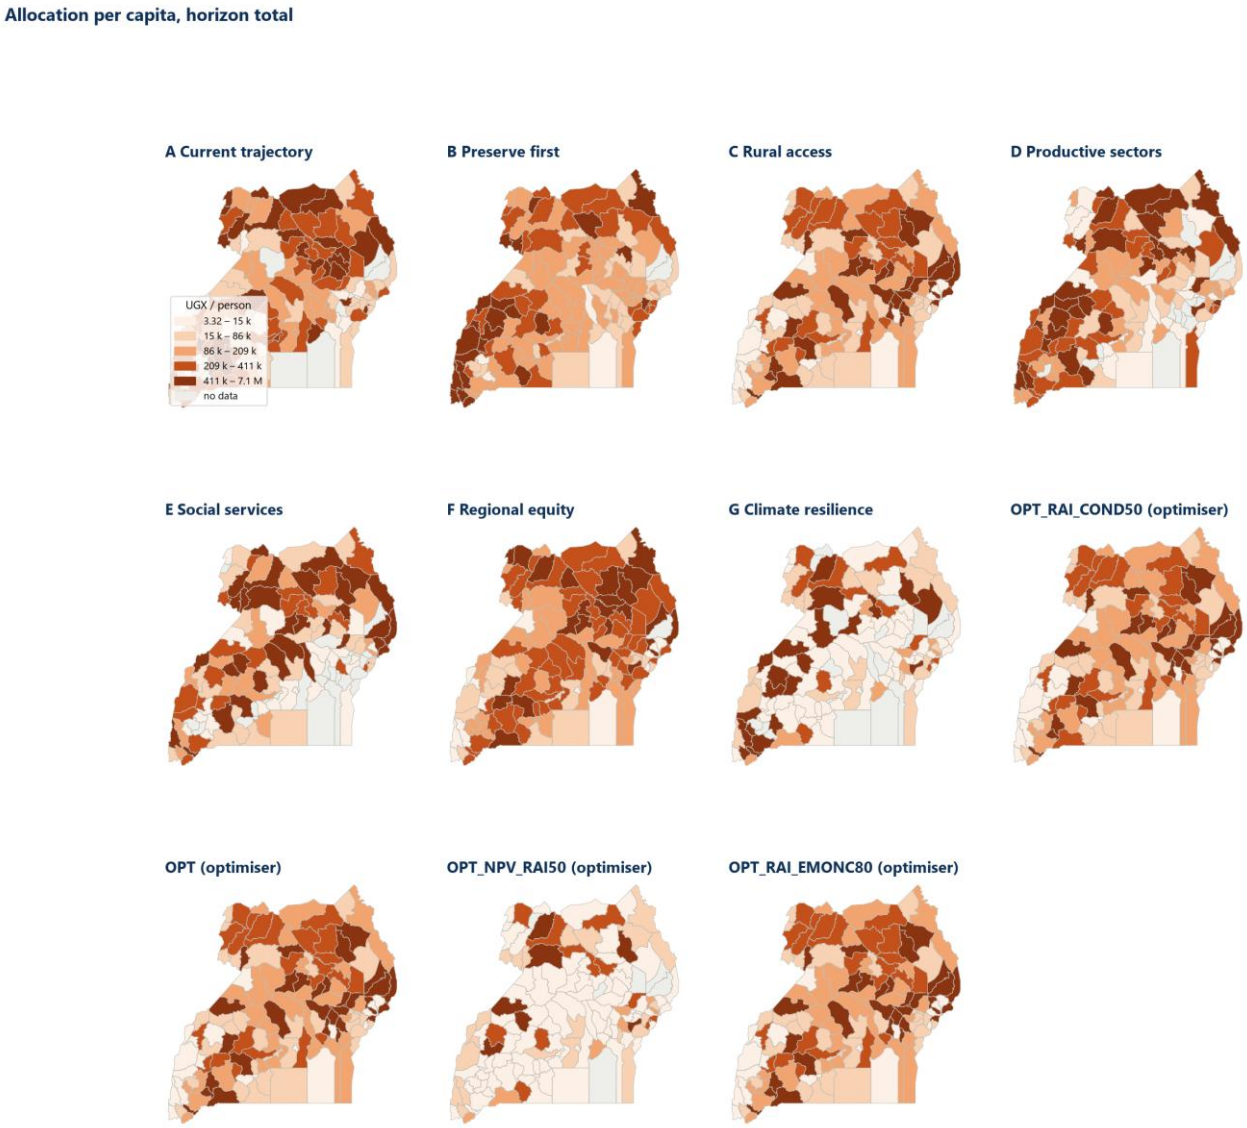

Figure 11.5: Capital allocation per inhabitant by district in every scenario — one shared colour scale

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The panels show distinct geographies:

- **A (current trajectory)**, which in this run is a fallback ordering rather than a reproduction of historical shares, reaches 127 of 135 districts, with a median of 214,000 UGX per person. It is concentrated in the north and north-east. Kapelebyong (1.29 M UGX per person), Lamwo (1.15 M) and Butambala (1.05 M) receive the most.
- **B (preserve first)** reaches 133 districts, median 158,000 UGX per person. It is heaviest in the west and south-west, where the network to be preserved is longest. Ntoroko (850,000), Nebbi (650,000) and Kanungu (649,000) lead.
- **C (rural access) and OPT** reach all 135 districts, median 121,000 UGX per person. They are concentrated in the north-east and east, where few people live near an all-season road. Abim (2.98 M), Amudat (2.54 M) and Bukwo (1.72 M) receive the most per person.
- **D (productive sectors)** reaches 120 districts, median 182,000, heavily in the agricultural west: Kibaale (1.35 M), Kakumiro (1.24 M) and Kapelebyong (1.06 M).
- **E (social services)** reaches 113 districts, median 156,000, and is the most concentrated in Karamoja: Amudat (3.87 M), Kotido (3.32 M) and Apac (2.58 M).

- **F (regional equity)** reaches 134 districts with the highest median, 248,000 UGX per person, and the flattest distribution: its maximum is only 669,000 (Kotido). This is what its district envelope rule (Equation (11.4)) is designed to do.
- **G (climate resilience)** reaches 116 districts but is extremely concentrated: the median is 16,000 UGX per person, and Ntoroko receives 7.10 M per person, Rukiga 3.73 M and Rukungiri 3.22 M. These are the districts with flood-exposed critical links.
- **OPT_NPV_RAI50** reaches 129 districts, with a median of only 6,000 UGX per person; Adjumani (610,000), Kikuube (506,000) and Kakumiro (487,000) lead.
- **OPT_RAI_COND50** looks like C, with a median of 119,000 UGX per person and the same three leading districts.

The scenarios with the largest totals are not the ones that spread money widest. F, with 11.3 tn, gives a higher median per person than B with 17.1 tn.

Figure 11.6 isolates the difference between C and B, the two scenarios that bring rural access closest to 70 %.

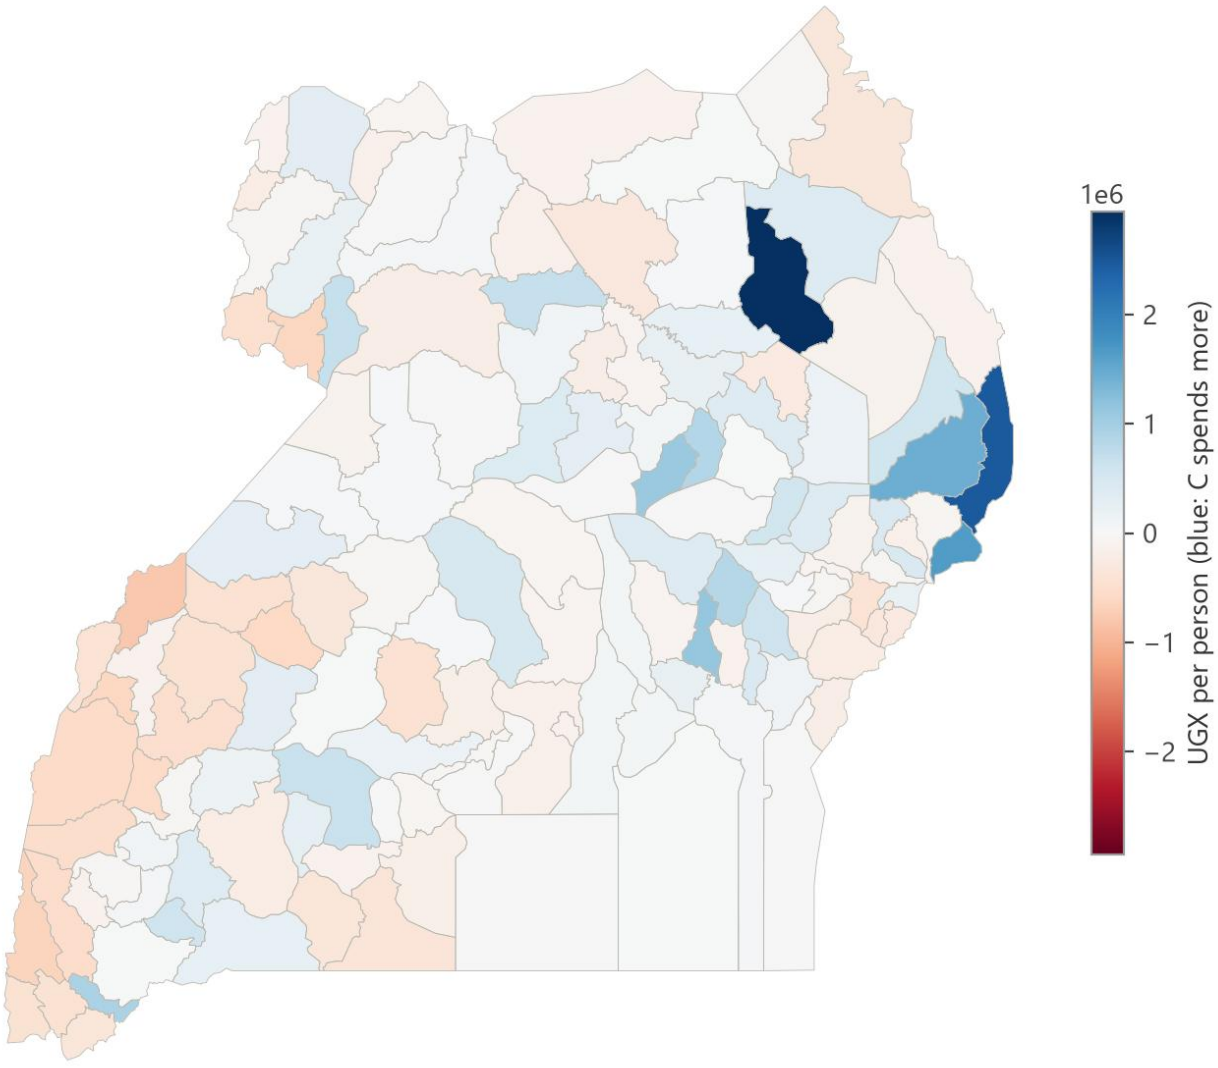

Figure 11.6: Where scenario C (rural access) and scenario B (preserve first) differ: allocation per capita, C − B

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **C spends more per person in 59 districts, B in 75.** The median district receives 15,000 UGX per person less under C than under B, because C's total is 5.8 tn smaller.

- **C's extra money goes to the north-east and east** (blue): Abim (+2.94 M UGX per person), Amudat (+2.46 M), Bukwo (+1.65 M), Nakapiripirit (+1.44 M), Luuka (+1.14 M) and Kaberamaido (+1.08 M). These are districts where few people are within 2 km of an all-season road today.
- **B's extra money goes to the west and south-west** (red): Ntoroko (−0.80 M), Kanungu (−0.64 M), Nebbi (−0.63 M), Bunyangabu (−0.62 M), Kibaale (−0.58 M) and Kasese (−0.57 M). These districts carry long existing networks with condition programmes.

The map is the spatial form of the choice between preserving and expanding. It is not a ranking of districts.

11.7 Results: the optimiser

Table 11.5 lists every measure at each optimiser solution and any floor shortfall.

Table 11.5: The optimiser runs: one objective each, other measures as ε-floors

Optimisation	Measure	Value	Floor shortfall
OPT	served:ALL_SEASON_ROAD	20,898,928	0.0%
OPT_RAI_EMONC80	served:ALL_SEASON_ROAD	20,898,928	0.0%
OPT_RAI_EMONC80	served:HEALTH_EMONC	17,789	92.3%
OPT_RAI_COND50	served:ALL_SEASON_ROAD	20,873,142	0.0%
OPT_RAI_COND50	condition_km	53,977	0.0%
OPT_NPV_RAI50	served:ALL_SEASON_ROAD	10,446,005	0.0%
OPT_NPV_RAI50	npv	2,265	0.0%

A positive shortfall is a finding: the floor could not be met within the solver limit. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_optimiser.csv).

- **OPT (maximise rural access, no floor)** serves 20,898,928 persons newly within 2 km of an all-season road, equal to the attainable maximum. The solver proved this within a gap of 0.0. The solution is the warm start, scenario C: within the candidate set, the greedy access ranking cannot be improved. This differs from run 5, where OPT reached an RAI of 0.657 against 0.645 for C.
- **OPT_RAI_COND50 (rural access, treating ≥ 50 % of attainable poor km)** treats 53,977 km of poor road, above its floor of 53,918.5 km, and serves 20,873,142 persons. Meeting the condition floor costs 25,786 persons of rural access, 0.12 % of the unconstrained value, at a solver gap of 0.12 %.
- **OPT_NPV_RAI50 (maximise NPV, keeping ≥ 50 % of attainable rural access)** reaches an NPV of 2,265.2 bn UGX, against an attainable maximum of 2,276.5 bn, and serves 10,446,005 persons. The access floor (10,449,464) is missed by 0.03 %, within the solver's 0.34 % gap. The run log flags it as a shortfall nonetheless, as it flags every positive shortfall.
- **OPT_RAI_EMONC80 (rural access, keeping ≥ 80 % of the attainable EmONC gain)** did **not** meet its floor. The attainable EmONC gain is 288,278 persons and the floor 230,623; the solution reaches 17,789, a shortfall of 92.3 %. The solver hit its 900-second limit with a relative gap of 112 % and returned the warm start, scenario C.

The EmONC result is a **known limitation of the solve, not a finding about Uganda**. The warm start is chosen as the rule-based allocation that scores highest on the *objective* (rural access), so it is C, which does nothing deliberate for emergency obstetric care. Within the time limit the solver could not move from C to a solution that also meets the floor. A warm start that already meets the floor (scenario E for EmONC) would fix this. The obstacle is size: a warm start from B, needed for the analogous condition case, brings about 40,000 projects into the candidate set, and at a 60,000-project cap every solve hit the time limit in an earlier trial run. This is open. The result should be read as "not solved", not as "EmONC access cannot be had at the price of 20 % of rural access".

The two access measures in this chapter and the next differ slightly. The optimiser counts persons in coverage groups whose enabling sets are complete (20.90 M for OPT). TI-09 recomputes the whole network after the allocation and counts persons meeting the standard against doing nothing (22.31 M for the same allocation). The

two are computed differently, and the difference, about 6.7 %, has not yet been decomposed. The optimiser's measure is the one it optimises; TI-09's is the one reported as the outcome.

11.7.1 The optimiser fix

Run 8 included a correction to `src/Optimise/Portfolio.jl`. With 41,867 priced projects, the old pre-screen took the per-measure lists in order and stopped at 14,000 projects. The access lists filled the cap, every condition programme was dropped, and `condition_km` came out as 0 whatever the floor. The attainable maximum was also computed without the warm start and could fall below OPT. The fix merges the per-measure lists round-robin, raises the cap to 20,000 plus the warm-start projects, and uses the same rule-based warm start for the attainable maxima and the frontier. The condition results in this chapter depend on that fix. TI-08 took 4,569 seconds (76 minutes) in the reference run, because it solves four optimisations, their attainable maxima and five frontier points.

11.8 Results: the trade-off between access and condition

Figure 11.7 and Table 11.6 trace the frontier between rural access and the poor km treated by the condition programmes.

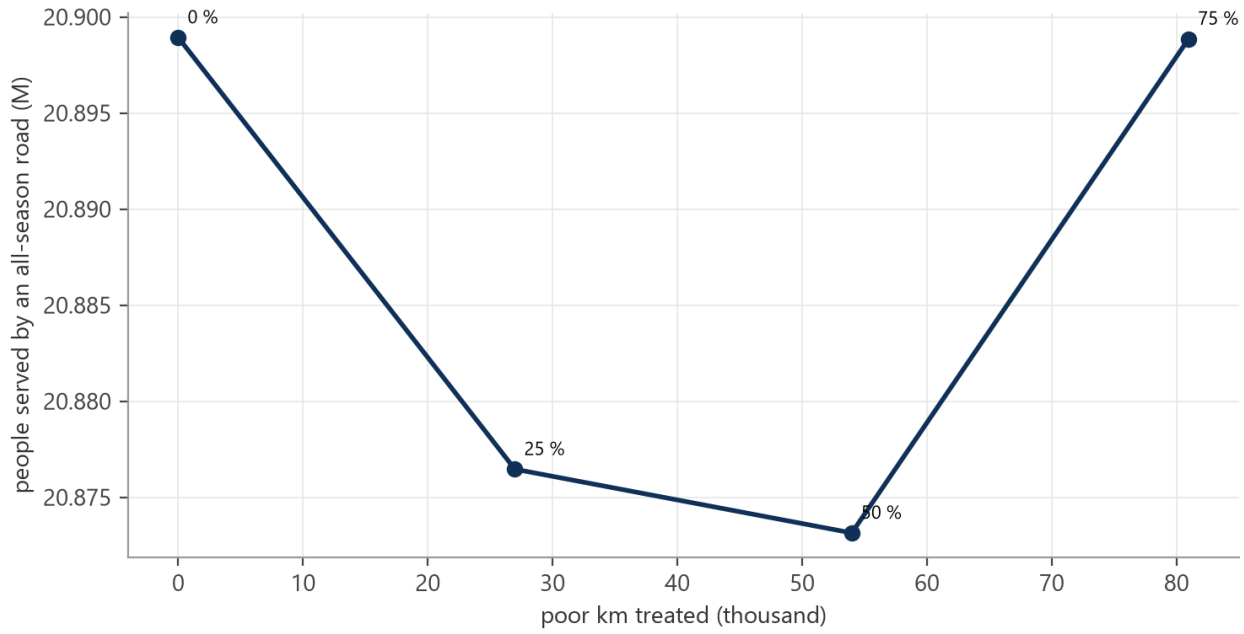

Figure 11.7: The trade-off frontier: people served against poor network treated

Source: TI-08 optimiser (JuMP/HiGHS), ε-constraint

Table 11.6: Trade-off frontier: rural access against poor network treated

Floor (share of attainable)	Floor (km)	People served (M)	Poor km treated	Shortfall	Solver status
0%	0	20.899	-	0.0%	OPTIMAL
25%	26,959	20.876	26,977	0.0%	OPTIMAL
50%	53,918	20.873	53,977	0.0%	OPTIMAL
75%	80,878	20.899	80,977	0.0%	OPTIMAL
100%	107,837	20.899	0	100.0%	TIME_LIMIT

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_frontier.csv).

- **Treating poor roads costs almost no rural access.** With no condition floor, 20.899 M persons are served and no poor km are treated. At 25 %, the optimiser treats 26,977 km and serves 20.876 M; at 50 %, 53,977 km and 20.873 M; at 75 %, 80,977 km and 20.899 M.

- **The dip is within the solver's tolerance.** The largest loss, at 50 %, is 25,786 persons or 0.12 %, well inside the 1 % MIP gap. The 75 % point, solved to a gap of 0.0, serves only 69 fewer persons than the unconstrained case. The V shape in the figure is therefore not a real trade-off curve; the frontier is flat within the precision of the solve. Read numerically, about 81,000 km of agency-reported poor road can be treated with no measurable loss of rural access within this envelope.
- **The 100 % point did not solve.** With the floor at 107,837 km, the solver hit its time limit (gap 111 %) and returned the warm start, C, which treats 0 km. Like the EmONC case, this is a solver limitation, not evidence that treating the last 25 % of poor km is impossible.

The reason the frontier is flat is the envelope. The access scenarios leave about 10.7 tn unspent (see above). Condition programmes can use that money, from the development pool and, for rehabilitation, from donors, without displacing any access project. Within the 25.8 tn envelope the choice between access and condition is not a choice of money. It is a choice of what the leftover money is allowed to buy. The figure and table make that plain without combining the two measures.

11.9 Results: capital against access

Figure 11.8 plots each allocation's capital against the Rural Access Index (RAI) at the horizon. The dashed line is doing nothing (28.7 %).

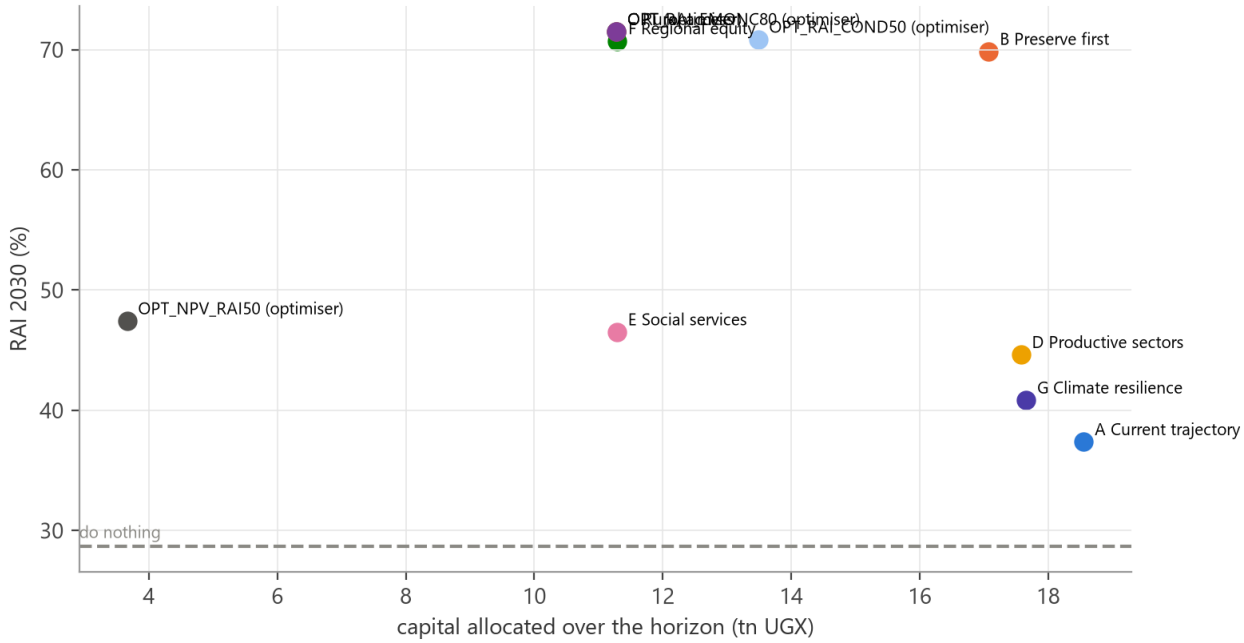

Figure 11.8: Capital allocated against rural access achieved: every scenario

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **C, OPT, OPT_RAI_EMONC80 and F reach 70.7 to 71.5 % with about 11.3 tn.** They add 42.1 to 42.8 points of RAI.
- **B reaches 69.8 % with 17.1 tn.** It adds 41.2 points, close to C, while also treating condition: it leaves 11,204 poor km, against 107,837 under C.
- **OPT_RAI_COND50 reaches 70.9 % with 13.5 tn** and leaves 53,860 poor km.
- **E reaches 46.5 %, D 44.6 %, G 40.8 % and A 37.4 %** with 11.3 to 18.6 tn. These scenarios pursue other objectives (social services, freight, resilience, current shares) and RAI is not their measure.
- **OPT_NPV_RAI50 reaches 47.4 % with only 3.7 tn.** It is the cheapest per person gaining access, at about 359,000 UGX of capital per person, against about 506,000 under C and about 4.0 M under A. Its capital is

small because it funds only the projects that pay their way plus the cheapest access projects needed for the floor.

The figure shows the comparison; it does not rank. A point to the upper left buys more RAI per shilling, but RAI is one measure among several. The scenarios to the lower right buy condition, freight, social-service access and resilience, which are reported in their own units in Chapter 12. All capital figures here carry induced maintenance, which TI-10 adds to the outer-year maintenance requirement.

11.10 Data quality and limitations

The allocation results are exploratory. The main limits are:

- **The envelope is a modelled ceiling.** The transport ceiling comes from MTEFModelling's baseline scenario, not from an approved MTEF. The roads share (0.545) and the pool split are taken from approved estimates, which are not outturn; external financing typically under-executes. The envelope for 2027 is 39 % above the FY2026/27 approved roads budget.
- **Pools are national.** Donor earmarks to named corridors or districts are not represented, so the cost of an earmark in access terms, which the manual asks for, is not yet reported.
- **Missing constraints.** Dependencies, exclusions, contracting capacity and per-district access floors (C10–C12, C15–C16) are not in the MILP. The front-loaded cash profile is the main consequence.
- **The horizon cuts off spending.** Projects that would finish after 2030 cannot start, which leaves 4.3 to 10.7 tn unallocated in the rule-based scenarios.
- **Solver limits.** OPT_RAI_EMONC80 and the 100 % frontier point reached the 900-second limit and returned the warm start, which does not meet their floors. All other solves finished within the 1 % gap.
- **Prices.** The 14 evidence prices are tier E5 candidates. Engineered earth, spot improvement and bridges rest on one or two observations each. Twenty-three spot-improvement projects remain unpriced and are never funded.
- **Network inputs.** Surface is known on only about 32 % of network km (31.7 %); unknown surface is treated as earth for access, so the access gains each scenario buys are upper-bounded by that convention. Condition is known only as agency statistics by network, not per link, so condition programmes are network tranches.
- **Benefits.** The appraisal behind NPV uses uncalibrated vehicle operating cost coefficients and a placeholder value for wet-season closures (Chapter 10), so OPT_NPV_RAI50 inherits those assumptions.

What would change the result most: a link-level condition survey (to replace the network tranches), commitments allowed beyond the horizon (to absorb the unallocated money), and a warm start that meets each floor (to solve the EmONC and 100 % condition cases).

11.11 Implications for planning

- **Money is not the binding constraint within this envelope; the horizon and the pool rules are.** Every scenario leaves 4.3 to 10.7 tn of the 25.8 tn unallocated, mostly because works cannot finish by 2030. A multi-year commitment framework that extends past the MTEF period would change what can be bought.
- **Routine maintenance absorbs the whole road-fund and recurrent pools and almost half of district grants before any capital is considered.** Any plan that adds assets adds to that claim; Chapter 13 shows the steady-state funding ratio of 0.58.
- **Rural access and network condition need not compete for money in this envelope.** The frontier shows about 81,000 poor km treated with no measurable loss of rural access. The decision is whether the money left after access is spent on condition, not whether access is sacrificed.

- **The same project set can be financed in very different ways.** OPT and C fund identical projects but differ by 4.4 tn in their call on domestic development money. Pool assignment should be decided explicitly with the finance ministry.
- **The geography of each scenario is distinct.** Access-led allocations concentrate on the north-east and east; preservation-led allocations on the west and south-west. District leaders will see these differences, and they should be presented as consequences of stated rules.
- **The EmONC floor result should not be quoted.** It is an unsolved case, not a finding. Fixing the warm start is a prerequisite for any health-floor analysis.
- **Contracting capacity must be added before the year profile is used.** All scenarios start almost every project in 2027, which no procurement system could deliver.

CHAPTER 12

12 Outcomes: what each allocation achieves

Chapter 11 showed what each allocation buys: which projects, from which pools, in which districts. This chapter asks what those purchases achieve. If the money is spent as each scenario proposes, how many more Ugandans will live within 2 km of an all-season road by 2030? Who gains: the poorest households or the richest? How many more people will reach a clinic, an emergency obstetric unit, a secondary school or a market? How much of the network reported in poor condition will still be poor? The TI-09 module answers these questions by applying each allocation's works to the network and recomputing access, rather than adding up the benefits claimed by individual projects. It implements design chapter 11 (`docs/11_TI09_OUTCOMES.md`) and the outcome equations of `docs/23_MATHEMATICAL_MODELS.md`. All results are from the reference run `014c9cd7` and are **exploratory**. No scenario is recommended: the outcomes are compared, each in its own unit, and never combined into a score.

12.1 What the module does

TI-09 takes each allocation from TI-08 (scenarios A–G and the four optimiser runs) and one more: **doing nothing**, the zero-capital allocation. For each, it:

1. applies the funded works to the network: every upgraded link takes its target surface and standard;
2. recomputes wet- and dry-season availability of each link;
3. recomputes travel times from every settlement cell to every destination class, on the full noded network;
4. counts the population meeting each access standard and the Rural Access Index (RAI);
5. splits the RAI by household wealth quintile;
6. subtracts the network condition programmes funded from the agency's poor-km statistics;
7. compares the recomputed gain with the sum of the gains each project claimed on its own (the interaction adjustment).

All outcomes are reported **against doing nothing**, not against the base year. In this run the do-nothing state equals the base-year network, because no deterioration is modelled: condition is not known per link, so the engine does not project it. The outcome is evaluated at the horizon year, 2030, with the population carried forward from UBOS 2022 district totals at an assumed 3 % a year.

The module writes `tables/ti09_scenario_comparison.csv` (one row per allocation, no ranking column), district layers of RAI at the horizon for every scenario, and the run log, which states the interaction adjustment for each scenario.

12.2 Method

12.2.1 The Rural Access Index

The RAI is the indicator of Sustainable Development Goal 9.1.1: the share of the rural population living within 2 km of an all-season road. In TIPRAS:

$$RAI = \frac{\sum_{i \in R} P_i A_i}{\sum_{i \in R} P_i} \tag{12.1}$$

where R is the set of rural settlement cells, P_i the population of cell i at the horizon and A_i equals 1 if the cell is within 2 km of a link that is passable in the wet season and has an all-season surface. The wet-season rule matters: a gravel road cut by a seasonal ford is not all-season. The surface rule matters more. Surface is known on only about 32 % of network km, and unknown surface is treated as earth. Chapter 7 shows that the base-year RAI lies between 28.7 % (unknown surface counted as earth) and 93.5 % (counted as gravel), against the published benchmark of 53.1 % for 2015. Every RAI in this chapter is computed under the lower convention. Gains from upgrading are therefore real gains in the engine's terms, but part of them may reflect roads that are already gravel in fact and simply unrecorded.

12.2.2 Equity: RAI by wealth quintile

The same index is computed separately for each wealth quintile:

$$RAI_q = \frac{\sum_{i \in R_q} P_i A_i}{\sum_{i \in R_q} P_i} \tag{12.2}$$

where R_q is the set of rural cells in quintile q. Each settlement cell carries the International Wealth Index (IWI) of the nearest survey point within 5 km. The quintile cut-points are the 20, 40, 60 and 80 % points of IWI across rural cells that have a value, not weighted by population. Quintile 1 is the poorest. IWI is an equity weight, not an official poverty statistic, and cells without an IWI value are excluded. The equity gap reported below is $RAI_5 - RAI_1$, in percentage points.

12.2.3 Persons served per destination class

For each destination class k, the outcome is the population meeting its access standard at the horizon, minus the same under doing nothing:

$$\Delta S_k = \sum_i P_i \cdot 1[tt_{i,k} \le \tau_k]_{alloc} - \sum_i P_i \cdot 1[tt_{i,k} \le \tau_k]_0 \tag{12.3}$$

where $tt_{i,k}$ is the travel time (or distance) from cell i to the nearest destination of class k and τ_k the class's standard from the draft norms (for example, 2 km to an all-season road; two hours to an EmONC facility). The classes are never added together.

12.2.4 Network condition

Link condition is unknown, so network condition is taken from the agency statistics (Ministry of Works and Transport sector performance report FY2022/23 and the UBOS Statistical Abstract 2024). The outcome is

$$KP = \sum_{n \neq urban} km_n^{poor} - \sum_{j \in PRG} km_j \tag{12.4}$$

where the first sum is the poor km on the national paved, national unpaved, district and community networks (107,837 km) and the second is the km of the network condition programmes the allocation funds. This is the poor km **treated**, not a projected condition: no deterioration is modelled at network level.

12.2.5 The interaction adjustment

TI-07 appraises each project on its own. Those gains do not add up. Two feeders serving the same village do not double its access, and a link upgraded in one project may be needed by another to reach a destination. TI-09 measures the difference:

$$IA = \frac{\sum_j B_j^{iso} - \Delta Z^{net}}{\sum_j B_j^{iso}} \tag{12.5}$$

where B_j^{iso} is project j's stand-alone gain in rural access and ΔZ^{net} the recomputed network gain for the whole allocation. If IA exceeds 15 %, the per-project figures are materially misleading as a guide, and the run log says so.

12.3 Results: rural access at the horizon

Figure 12.1 shows the RAI in 2030 for every allocation, and Table 12.1 gives the full comparison.

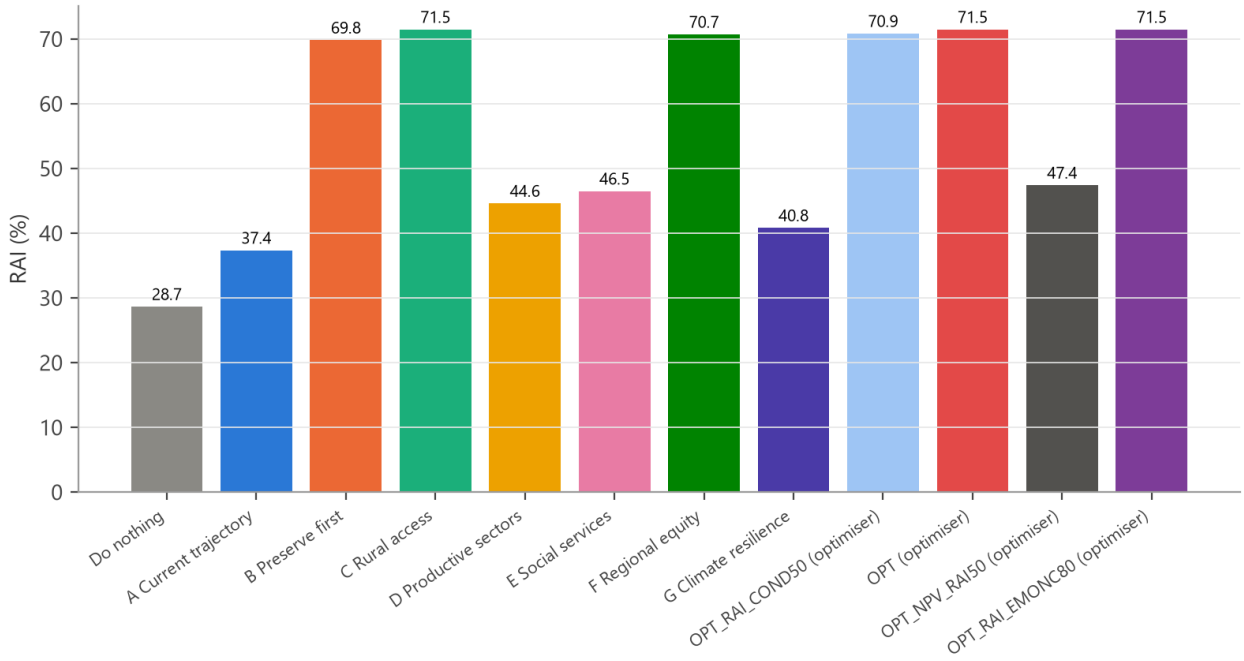

Figure 12.1: Rural Access Index at the horizon in every scenario (no scenario is recommended)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 12.1: Scenario comparison at the horizon (2030): each outcome in its own unit, compared, never scored

Scenario	RAI 2030	Poorest quintile	Richest quintile	Served, all-season (M)	Served, EmONC (M)	Served, market (M)	Agency poor km left	Interaction adj.
Do nothing	28.7%	10.0%	52.1%	21.04	12.59	10.38	107,837	-
A Current trajectory	37.4%	15.6%	61.7%	25.63	12.69	10.55	0	52.7%
B Preserve first	69.8%	54.1%	84.4%	42.35	12.73	10.67	11,204	86.9%
C Rural access	71.5%	55.8%	85.3%	43.34	12.78	10.71	107,837	49.4%
D Productive sectors	44.6%	28.1%	65.8%	28.85	12.74	10.86	22,976	66.2%
E Social services	46.5%	34.4%	65.9%	29.73	12.90	10.76	107,837	53.0%
F Regional equity	70.7%	55.2%	84.6%	42.95	12.76	10.70	107,837	51.2%
G Climate resilience	40.8%	21.7%	62.6%	27.05	12.69	10.67	21,657	81.0%
OPT_RAI_COND50 (optimiser)	70.9%	55.6%	84.7%	43.04	12.78	10.70	53,860	34.2%
OPT (optimiser)	71.5%	55.8%	85.3%	43.34	12.78	10.71	107,837	49.4%
OPT_NPV_RAI50 (optimiser)	47.4%	26.9%	68.4%	31.22	12.65	10.48	107,837	20.6%
OPT_RAI_EMONC80 (optimiser)	71.5%	55.8%	85.3%	43.34	12.78	10.71	107,837	49.4%

No scenario is recommended. Optimiser runs are labelled by objective and floors. Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_scenarios.csv).

Doing nothing leaves the RAI at 28.7 %. Every allocation raises it. The results fall into three groups:

- **About 70 % or more:** C (rural access) 71.5 %, OPT (maximising rural access) 71.5 %, OPT_RAI_EMONC80 71.5 %, OPT_RAI_COND50 (rural access with a 50 % condition floor) 70.9 %, F (regional equity) 70.7 % and B (preserve first) 69.8 %. These allocations raise the RAI by 41 to 43 points.
- **About 45 %:** OPT_NPV_RAI50 (maximising NPV with a 50 % access floor) 47.4 %, E (social services) 46.5 % and D (productive sectors) 44.6 %.
- **About 40 % or less:** G (climate resilience) 40.8 % and A (current trajectory, as implemented here) 37.4 %.

OPT, C and OPT_RAI_EMONC80 have identical outcomes because they fund the identical set of projects (Chapter 11). OPT_RAI_EMONC80's floor was not met within the solver's time limit, so it returned C's allocation.

In people, the do-nothing state has 21.04 M persons within 2 km of an all-season road at the horizon. C adds 22.31 M, F 21.91 M, B 21.31 M and OPT_RAI_COND50 22.00 M. A adds 4.59 M.

The spread between B and C is small (1.7 points), but their routes differ. C reaches its result with 11.3 tn UGX of capital and treats no poor km. B uses 17.1 tn, of which 5.8 tn goes to rehabilitation and periodic maintenance, and leaves 11,204 poor km. Each is a different answer to a different question. The table reports both without ranking them.

These RAI values are far above the 2015 benchmark of 53.1 % only because they start from the lower bound of 28.7 %. Under the upper-bound convention, many of the upgrades would add nothing, because the roads would already count as gravel. The gain in points is the meaningful quantity, and it is conditional on the surface convention.

12.4 Results: who gains

Figure 12.2 shows the RAI by wealth quintile in each scenario.

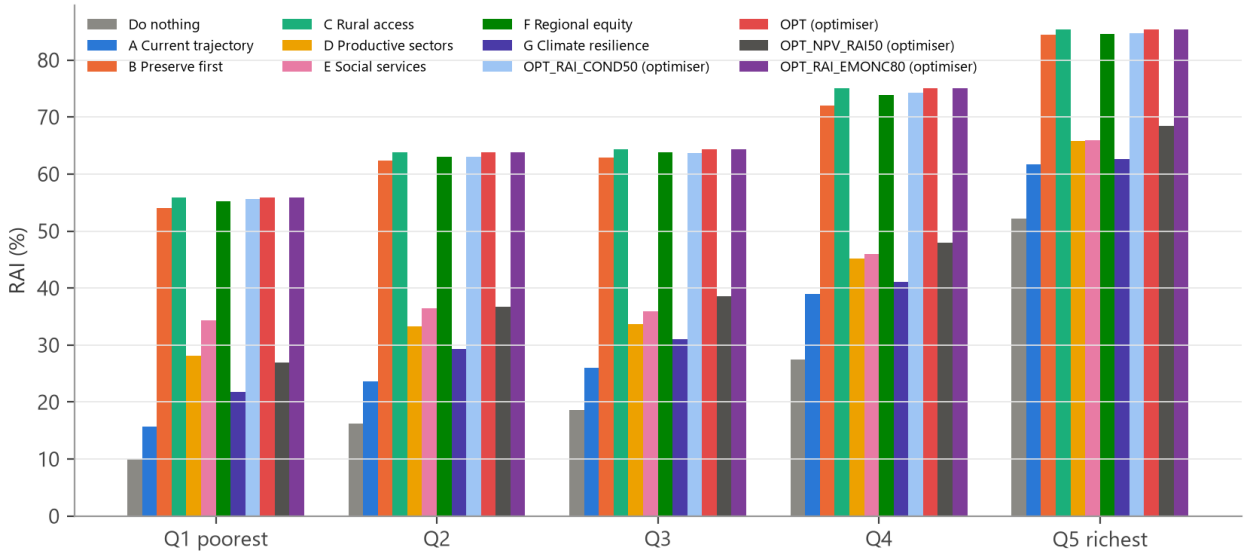

Figure 12.2: Who gains: rural access by household wealth quintile (IWI) in each scenario

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Under doing nothing, access rises steeply with wealth: 10.0 % for the poorest quintile, 16.2 % for the second, 18.5 % for the third, 27.5 % for the fourth and 52.1 % for the richest. The richest fifth of rural Ugandans is five times as likely as the poorest to live within 2 km of an all-season road.

Every allocation raises every quintile, but by very different amounts.

- **C and OPT** raise the poorest quintile from 10.0 % to 55.8 % (+45.9 points) and the richest from 52.1 % to 85.3 % (+33.2 points).
- **F** raises the poorest to 55.2 % and the richest to 84.6 %.
- **B** raises the poorest to 54.1 % and the richest to 84.4 %.
- **E** raises the poorest to 34.4 %; its second and third quintiles reach 36.4 % and 36.0 %, so the middle gains least.
- **D** raises the poorest to 28.1 %, **G** to 21.7 % and **A** to only 15.6 % (+5.7 points), while A raises the richest by 9.5 points to 61.7 %.

Figure 12.3 reduces this to one number per scenario: the gap between the richest and the poorest quintile.

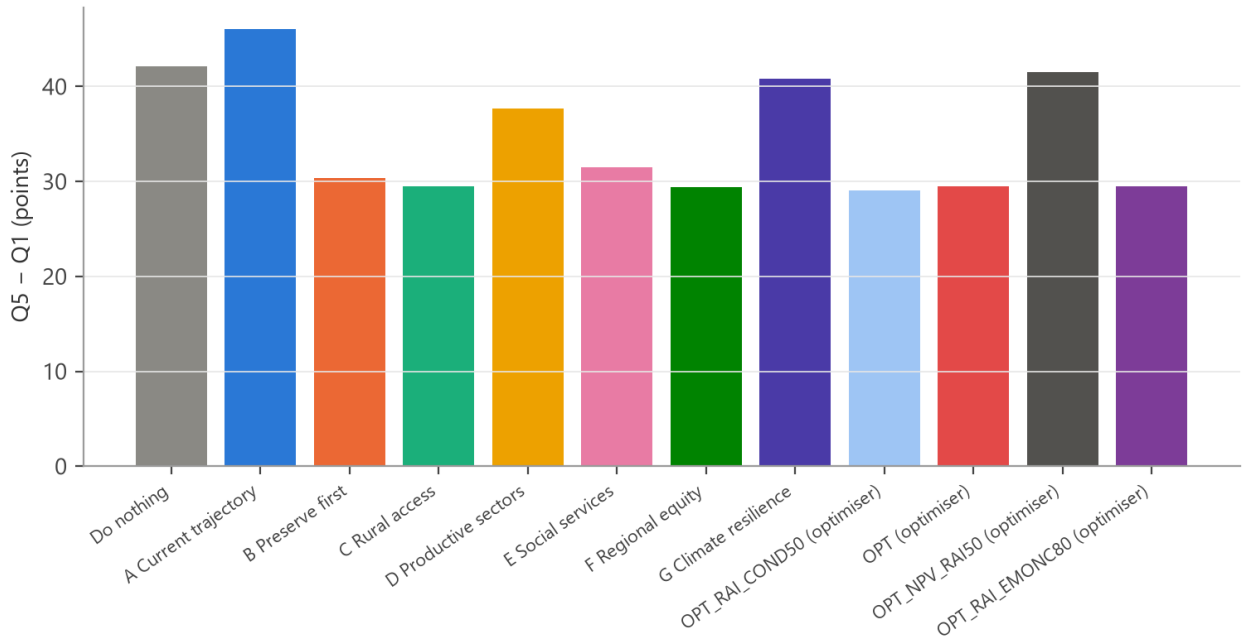

Figure 12.3: The wealth gap in access: richest minus poorest quintile, by scenario (points)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- Doing nothing leaves a gap of **42.2 points**.
- C, OPT and F narrow it to about 29.5 points (C 29.5, F 29.4), and OPT_RAI_COND50 to 29.1. B narrows it to 30.3 and E to 31.5.
- D (37.7), G (40.8) and OPT_NPV_RAI50 (41.5) narrow it only a little.
- **A widens it to 46.0 points.** It is the only allocation that leaves the poor further behind. Its fallback ordering funds national-class projects first; the gains it buys fall disproportionately to the richer quintiles.

The equity gap is a property of each allocation, not a weight on its benefits. It is reported beside the RAI, never multiplied into it.

12.5 Results: access to services and markets

Figure 12.4 shows the persons newly meeting each access standard against doing nothing, with each destination class in its own panel and its own scale.

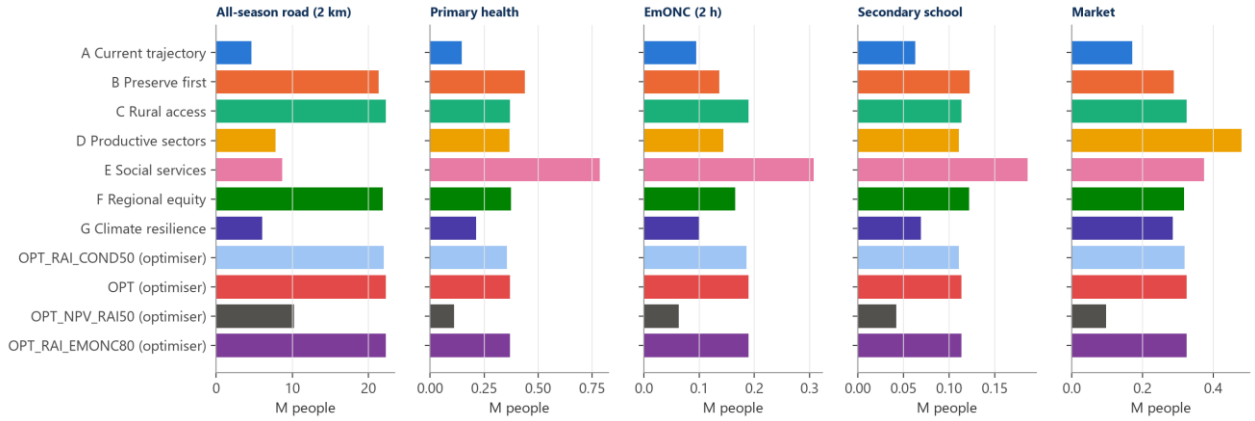

Figure 12.4: People newly meeting each access standard against doing nothing, by scenario (each class in its own panel)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 12.2: Persons newly meeting each access standard against doing nothing, 2030 (thousand persons)

Scenario	All-season road (2 km)	Primary health	EmONC	Secondary school	Market
A Current trajectory	4,590.9	146.4	94.9	63.2	170.4
B Preserve first	21,313.6	439.8	136.8	122.6	288.7
C Rural access	22,305.1	369.6	189.0	113.6	324.2
D Productive sectors	7,815.1	367.2	144.1	110.6	479.8
E Social services	8,698.5	786.1	307.5	186.1	374.4
F Regional equity	21,911.2	373.6	165.3	122.1	317.0
G Climate resilience	6,009.7	212.2	100.1	69.0	285.1
OPT_RAI_COND50	21,999.7	355.0	185.4	110.9	318.6
OPT	22,305.1	369.6	189.0	113.6	324.2
OPT_NPV_RAI50	10,186.1	111.2	63.4	41.9	97.1
OPT_RAI_EMONC80	22,305.1	369.6	189.0	113.6	324.2

The panels show that each scenario does what it was built to do, and that no scenario does everything:

- **E (social services)** adds the most people within the health and school standards: 786,100 to primary health, 307,500 to EmONC and 186,100 to secondary school. That is more than twice C's gain for primary health and 1.6 times C's for EmONC. But E adds only 8.70 M to rural access, against C's 22.31 M.
- **D (productive sectors)** adds the most people within the market standard: 479,800, against 324,200 for C. It adds 7.82 M to rural access.
- **B, C, F and OPT** dominate rural access (21.3 to 22.3 M) and give middling gains in the other classes.
- **A, G and OPT_NPV_RAI50** add the least in most classes.

The gains in the service classes are small relative to the population because base-year access to services is already high under the draft standards: in the wet season, 93.8 % of the requiring population meets the primary-health standard, 91.5 % the EmONC standard and 89.0 % the secondary-school standard (Chapter 7). Much of the remaining gap is caused by the absence of a facility, not of a road: Chapter 8 attributes 614,403 of the 1,035,094 persons without EmONC access to "no destination" within reach. A road cannot close that part. It is a siting question for the health sector, and the engine reports it to HealthModellingx as such.

The EmONC figure for OPT_RAI_EMONC80 here (189,000 persons, the same as C) differs from the optimiser's own measure in Chapter 11 (17,789). The two are computed differently: the optimiser counts persons in coverage groups whose enabling sets are complete, while TI-09 recomputes the network. Neither meets the 230,623 floor. The difference between the two measures has not yet been decomposed.

12.6 Results: where access improves

Figure 12.5 maps the RAI at the horizon by district for every scenario, on one shared colour scale.

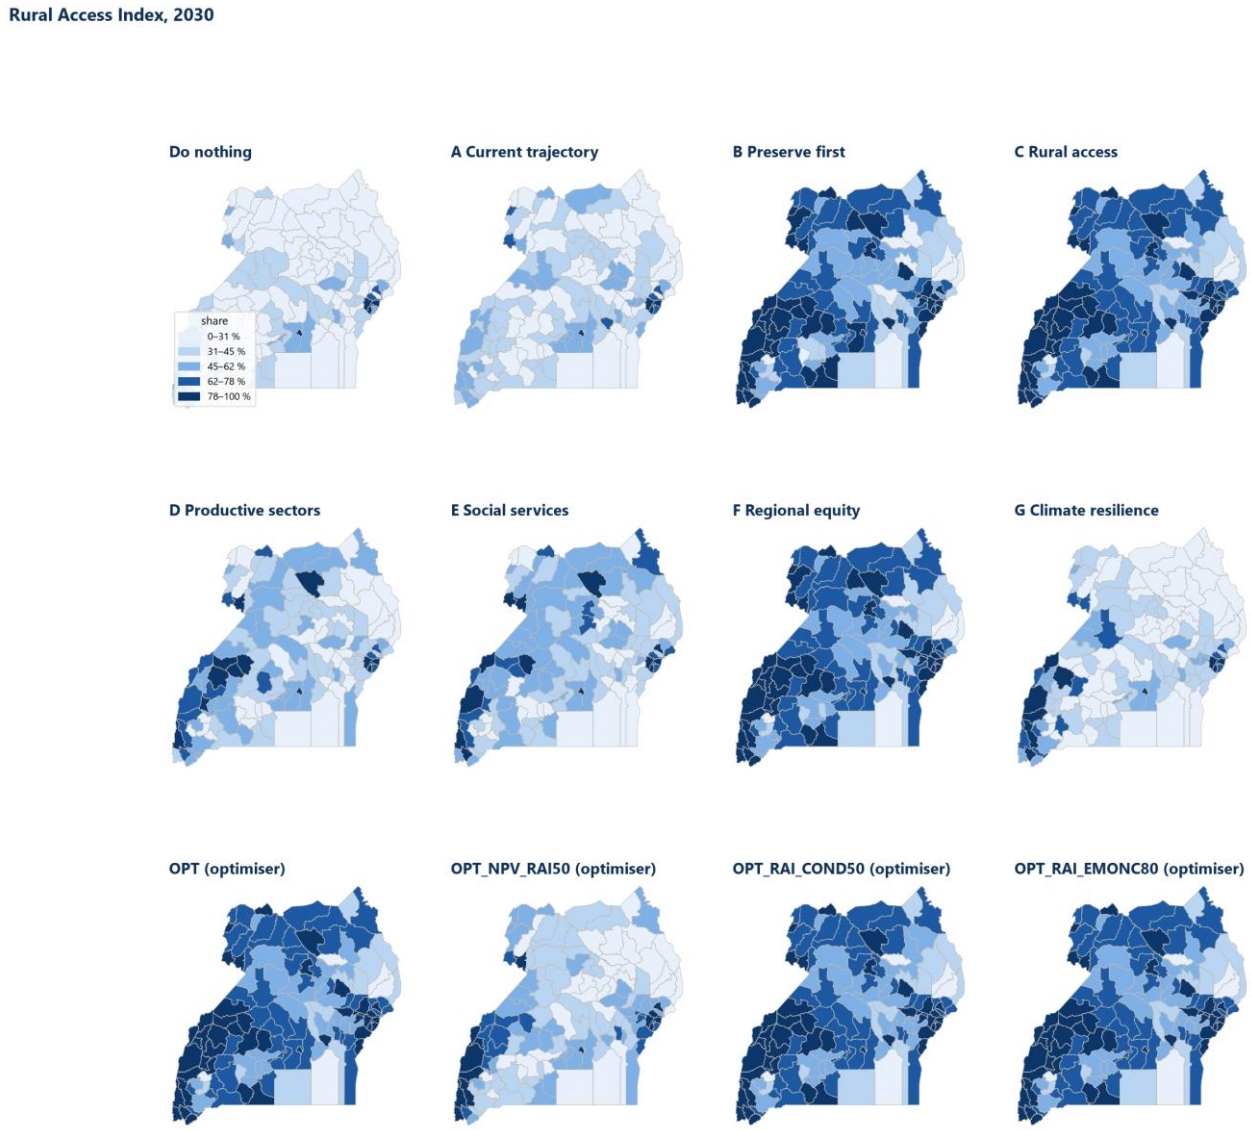

Figure 12.5: Rural Access Index at the horizon by district in every scenario — one shared colour scale

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **Doing nothing:** 80 of 135 districts have an RAI below 30 %, and only 14 are at 50 % or more. The median district is at 26.8 %. Buvuma and Kalangala, both island districts, are at 0 %; Manafwa (83.0 %), Kampala (82.6 %) and Mbale (81.5 %) are the highest.
- **C and OPT:** only 4 districts remain below 30 %, 114 reach 50 % or more and 26 reach 90 % or more. The median district reaches 69.4 %. The lowest are Buvuma (6.6 %), Buhweju (27.9 %), Nabilatuk (28.1 %) and Nakapiripirit (28.4 %).
- **F:** 6 districts below 30 % and 108 at 50 % or more.
- **B:** 10 districts below 30 %, 103 at 50 % or more and 30 at 90 % or more, the most of any scenario. The lowest are Buvuma (6.6 %), Nakapiripirit (11.5 %), Amudat (13.8 %) and Nabilatuk (14.4 %), all but Buvuma in Karamoja.
- **E:** 25 districts below 30 %; **D:** 32; **G:** 53; **A:** 49.

B reaches a higher district RAI than C in 52 districts, although C is higher nationally. B lifts more districts close to full coverage; C lifts more of the low districts above 30 %. Amudat illustrates the difference: 40.7 % under C, 13.8 % under B.

Figure 12.6 maps the gain against doing nothing.

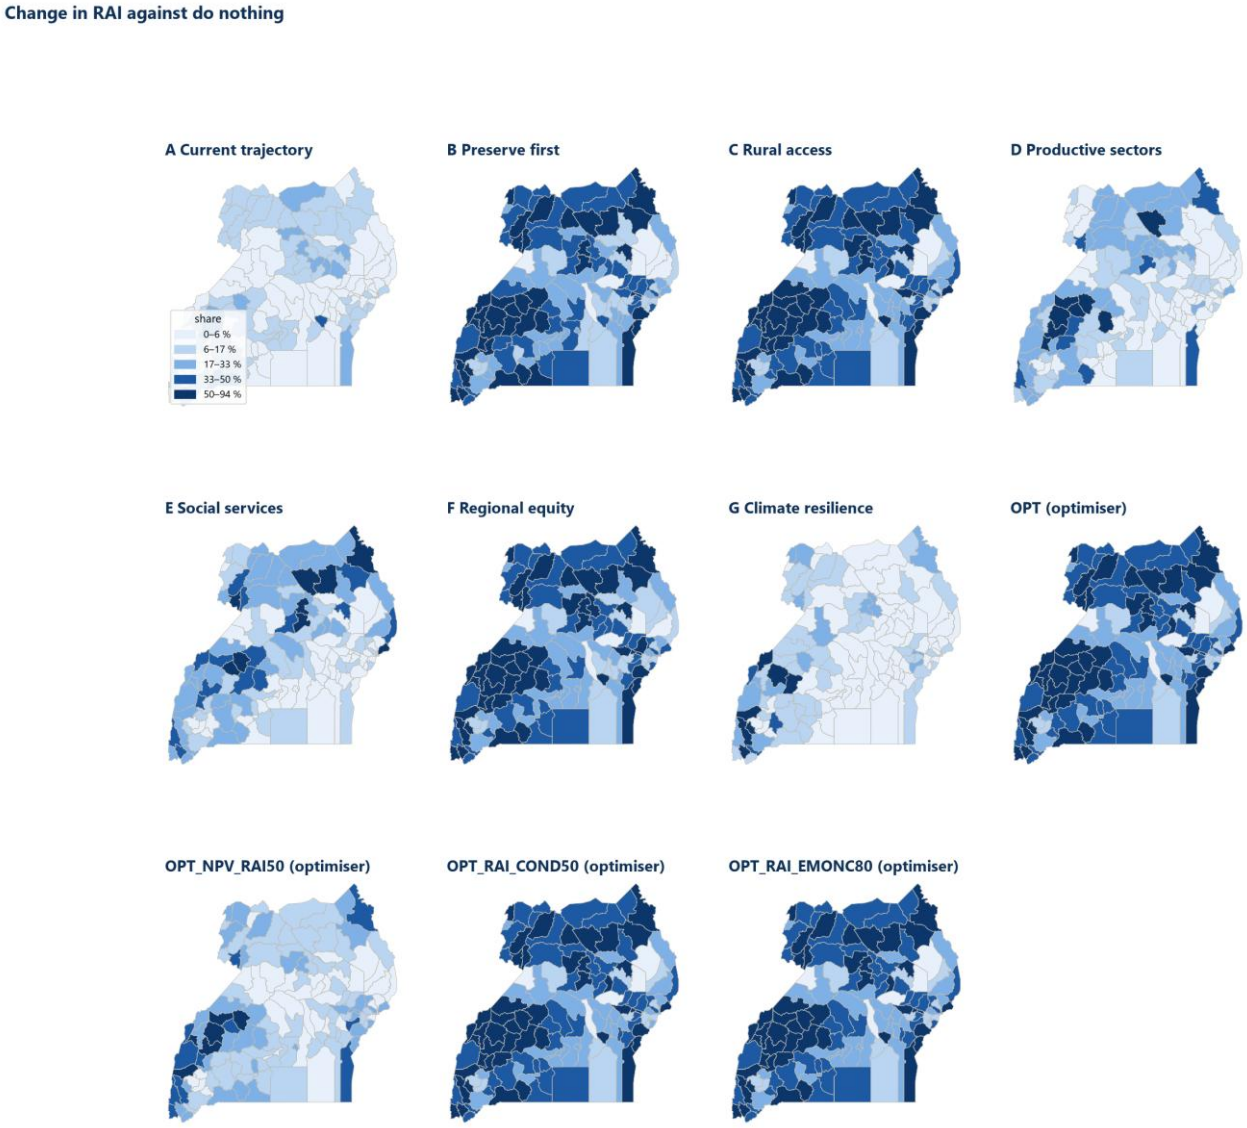

Figure 12.6: Gain in rural access against doing nothing, by district and scenario (shared scale)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **The largest gains are in the mid-west.** Under C, Kakumiro gains 93.0 points, Kibaale 89.6, Kole 88.6, Kagadi 81.8 and Kamwenge 79.5. B's largest gains are in the same districts: Kakumiro 93.6, Kibaale 90.9, Kole 88.8, Rakai 85.2 and Kamwenge 84.1. These districts start from low RAI values and large rural populations, so each upgrade reaches many people.
- **The smallest gains under C** are in Katakwi (3.9 points), Buliisa (4.2), Kayunga (5.1), Napak (5.3) and Serere (5.8), and in the islands of Buvuma (6.6), where few upgrading candidates reach many people.
- **The median district gains** 40.8 points under C, 39.6 under F, 38.7 under B, 13.2 under E, 8.4 under D, 6.8 under A and 5.5 under G.
- **Narrow scenarios leave many districts untouched.** D leaves 25 districts with no gain at all, E 18, G 14 and A 13.

The gain maps of B, C, F and the three access-led optimiser runs look much alike. The difference between them is in the money (Chapter 11) and in condition, not in where access improves.

12.7 Results: the network left in poor condition

Figure 12.7 shows the agency-reported poor km left untreated in 2030.

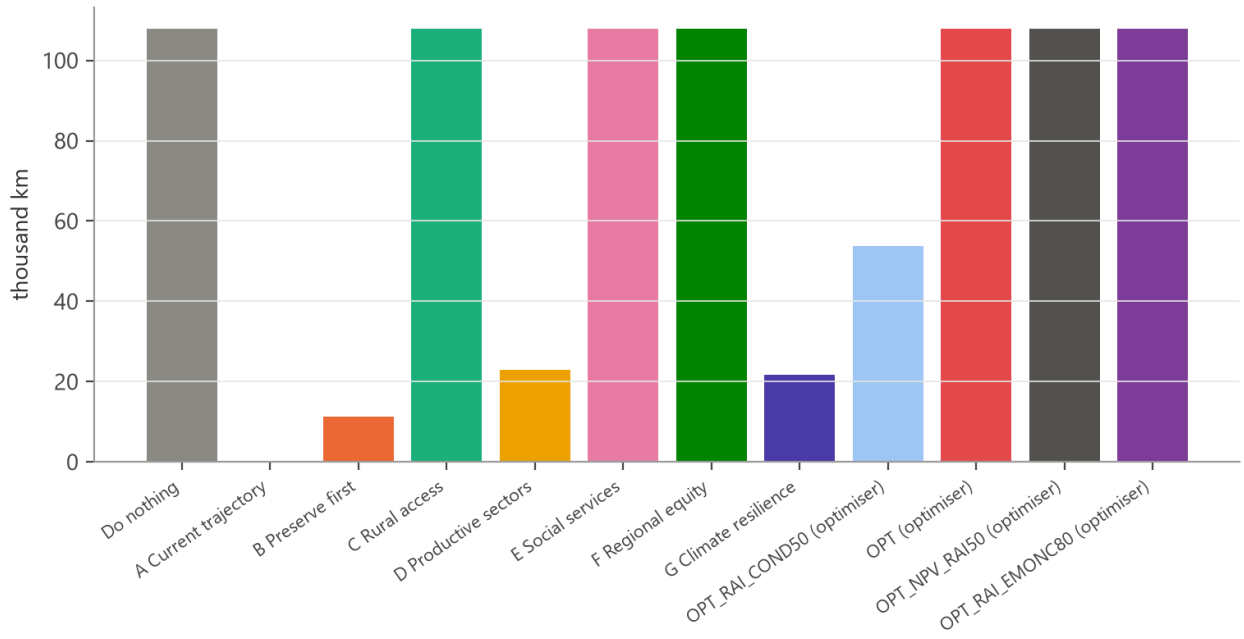

Figure 12.7: Agency-reported poor km left untreated at the horizon, by scenario

Source: AICTEM TIPRAS engine, reference run 014c9cd7

- **A treats all 107,837 poor km** outside urban areas, leaving 0. It funds every condition programme (1.07 tn of periodic maintenance and 6.18 tn of rehabilitation).
- **B leaves 11,204 km.** It funds all periodic maintenance and 4.74 tn of rehabilitation.
- **G leaves 21,657 km and D 22,976 km.** Both fund about 0.6 tn of periodic maintenance and 5.7 to 5.8 tn of rehabilitation.
- **OPT_RAI_COND50 leaves 53,860 km**, having treated 53,977 km to meet its 50 % floor.
- **C, E, F, OPT, OPT_RAI_EMONC80 and OPT_NPV_RAI50 leave all 107,837 km.** They fund no condition programme.

Rural access and condition pull against each other only at the margin within this envelope. B reaches 69.8 % RAI and treats 96,633 km; C reaches 71.5 % and treats none. OPT_RAI_COND50 shows a middle path: 70.9 % RAI and half the poor km treated. The frontier in Chapter 11 shows that the optimiser can treat about 81,000 km with no measurable loss of access, because the access-led allocations leave money unspent.

This indicator rests on agency statistics from FY2022/23. It treats a funded programme as fixing its km and does not model deterioration of other roads over the four years. The link-based measure of poor or bad km (`km_poor_or_bad` in the scenario table) is zero in every scenario because condition is unknown on every link. That zero means "not known", not "good".

12.8 Results: why project gains do not add up

Figure 12.8 shows the interaction adjustment for each allocation: the share by which the sum of the stand-alone access gains of the funded projects overstates the gain the network actually delivers when they are built together.

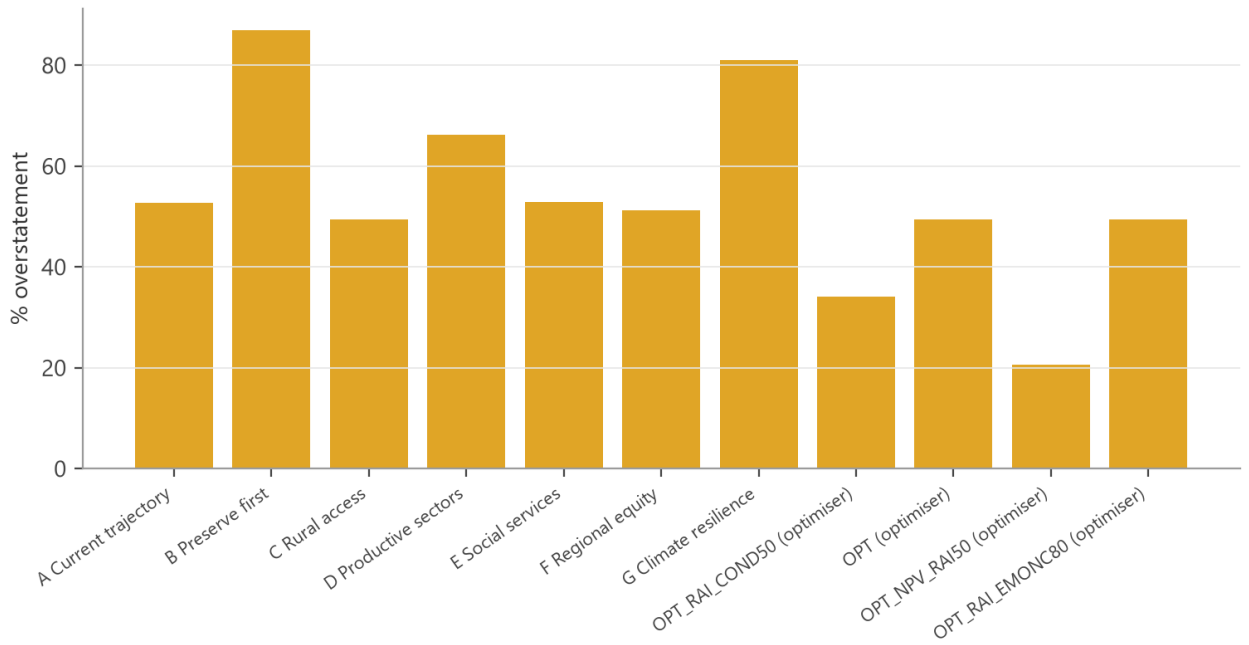

Figure 12.8: Interaction adjustment: how far stand-alone project gains overstate the network result

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Every allocation exceeds the 15 % threshold, and the run log flags each one:

- **B: 86.9 %** and **G: 81.0 %**. These scenarios fund many projects in the same areas, and their stand-alone gains overlap heavily.
- **D: 66.2 %**, **E: 53.0 %**, **A: 52.7 %**, **F: 51.2 %**.
- **C and OPT: 49.4 %.** Even the access-ranked allocation, which updates benefits for shared coverage sets as it goes, overstates by half when projects are summed.
- **OPT_RAI_COND50: 34.2 %** and **OPT_NPV_RAI50: 20.6 %**, the lowest. The NPV-led allocation funds the fewest projects (2,289), and their stand-alone gains overlap least.

The finding is methodological but matters for planning. A list of projects each claiming to connect a number of people will overstate what the list delivers by between a fifth and seven-eighths. Project-by-project claims cannot be summed into a programme target. The network must be recomputed, as TI-09 does.

12.9 Results: what each allocation buys and gives up

The outcomes above can be read across, one allocation at a time, to see what each buys and what it gives up. Table 12.3 does this without scoring: each column stays in its own unit, and the last column is a ratio of one outcome to capital, not an index.

Table 12.3: What each allocation buys and gives up, 2030 (each column in its own unit)

Allocation	Capital (tn UGX)	RAI gain (points)	Poorest-quintile gain (points)	Poor km left	Persons gaining all-season access per bn UGX
A Current trajectory	18.55	8.7	5.7	0	247
B Preserve first	17.08	41.2	44.1	11,204	1,248
C Rural access	11.28	42.8	45.9	107,837	1,978
D Productive sectors	17.58	15.9	18.1	22,976	445
E Social services	11.29	17.8	24.4	107,837	771
F Regional equity	11.29	42.1	45.2	107,837	1,941
G Climate resilience	17.66	12.2	11.7	21,657	340
OPT_RAI_COND50	13.48	42.2	45.6	53,860	1,632

Allocation	Capital (tn UGX)	RAI gain (points)	Poorest-quintile gain (points)	Poor km left	Persons gaining all-season access per bn UGX
OPT	11.28	42.8	45.9	107,837	1,978
OPT_NPV_RAI50	3.66	18.8	16.9	107,837	2,784

Read row by row:

- **A** treats all the reported poor km but buys the least access per shilling (247 persons per bn UGX) and the smallest gain for the poorest.
- **B** buys almost as much access as C, treats 90 % of the poor km and costs 5.8 tn more. Its access per shilling is about two-thirds of C's.
- **C and OPT** buy the most access for the poorest, at 1,978 persons per bn UGX, and treat no poor km.
- **D** buys the largest market-access gain (479,800 persons) and treats 79 % of the poor km, at modest access gains.
- **E** buys the largest health and school gains and little else.
- **F** is within a point of C on every access measure, with a flatter spread of money across districts (Chapter 11).
- **G** protects flood-exposed critical links and treats 80 % of the poor km, with the second-smallest access gain.
- **OPT_RAI_COND50** sits between B and C on both access and condition.
- **OPT_NPV_RAI50** spends least and buys the most access per shilling (2,784 persons per bn UGX), because it funds only projects with positive NPV plus the cheapest access projects needed for its floor. It leaves 18.3 tn unspent and treats no poor km.

The capital in each row carries induced routine maintenance once the assets are in service; Chapter 13 shows the steady-state requirement for scenario B. No row is marked as the one to choose. Choosing between them means deciding how much weight the government puts on access, on the poorest, on condition, on services and on cost. That is a political judgement, which the table informs but does not make.

12.10 Simulations: what changes the result

The scenarios A–G answer one question: *with this money, these prices and this evidence, what does each allocation rule buy?* A second question is just as important for a planner: *which of the inputs, if it were different, would change the answer?* To answer it, the engine was run seven more times. Each run changes one input of the reference configuration and keeps everything else fixed. All seven run the full chain TI-01 to TI-10 and the eight allocation rules; they do not run the optimiser, which adds an hour a run and does not change the rule-based results. Table 12.4 lists them.

Table 12.4: The simulation set: one change at a time from the evidence base case

Simulation	Norms	Prices	Unknown surface	What changes
S0 Base (evidence)	UGA-2026.09-draft-f	UCSN+EV-UGA-2026.09b	survey	Evidence prices (central), evidence envelope, unknown surface listed for survey (the reference configuration)
S1 Prices low	UGA-2026.09-draft-f	UCSN+EV-UGA-2026.09b-low	survey	Every evidence price at its low bound
S2 Prices high	UGA-2026.09-draft-f	UCSN+EV-UGA-2026.09b-high	survey	Every evidence price at its high bound
S3 Envelope −30 %	UGA-2026.09-draft-f-env70	UCSN+EV-UGA-2026.09b	survey	Roads share of the MTEF transport ceiling at 70 % of the evidence value (0.38)

Simulation	Norms	Prices	Unknown surface	What changes
S4 Envelope +30 %	UGA-2026.09-draft-f-env130	UCSN+EV-UGA-2026.09b	survey	Roads share of the MTEF transport ceiling at 130 % of the evidence value (0.71)
S5 Unknown surface = earth	UGA-2026.09-draft-f	UCSN+EV-UGA-2026.09b	assume_earth	Links with unknown surface treated as earth and given works (sensitivity)
S6 UCSN-first prices	UGA-2026.09-draft-e	UCSN+EV-UGA-2026.09b	survey	UCSN-first price map: illustrative UCSN codes where UCSN has one, evidence elsewhere

Source: AICTEM TIPRAS engine, simulation runs S0–S6 (outputs/simulations/simulations.csv) (report/build/data/t_sim_design.csv).

S0 reproduces the reference run exactly: the same 41,867 priced projects, the same 92.8 tn programme, the same 71.5 % RAI under C, the same 954.7 bn a year of routine maintenance. It is the yardstick for the other six. Two of the simulations bracket the **price evidence** (S1 sets every evidence price to the low end of its documented range, S2 to the high end), two bracket the **money** (S3 and S4 move the roads share of the MTEF transport ceiling 30 % down and 30 % up), one tests the **largest evidence gap** (S5 treats the 68 % of km whose surface is unknown as earth and lets the engine propose works on them) and one tests the **price map** (S6 uses the UCSN-first map, in which the seven illustrative UCSN codes are used wherever UCSN has a code).

12.10.1 What the simulations change

Table 12.5 and Table 12.6 give the results; Figure 12.9 shows the Rural Access Index at the horizon for every rule under every simulation, and Figure 12.10 the same numbers as a grid.

Table 12.5: Simulation results: programme cost and access outcomes at the horizon

Simulation	Programme capital (tn UGX)	Range	Induced RM (bn UGX/yr)	Projects priced	RAI do nothing	RAI A	RAI B	RAI C	Poorest quintile (C)	Funding ratio (B)
S0 Base (evidence)	92.76	56.64–125.26	205	41,867	28.7%	37.4%	69.8%	71.5%	55.8%	0.58
S1 Prices low	56.64	56.64–125.26	126	41,867	28.7%	44.5%	70.3%	72.6%	57.1%	1.98
S2 Prices high	125.26	56.64–125.26	298	41,867	28.7%	42.7%	69.5%	70.7%	54.6%	0.30
S3 Envelope −30 %	92.76	56.64–125.26	205	41,867	28.7%	38.3%	69.7%	70.7%	54.6%	0.41
S4 Envelope +30 %	92.76	56.64–125.26	205	41,867	28.7%	41.1%	70.1%	72.1%	56.6%	0.75
S5 Unknown surface = earth	153.73	94.61–209.71	488	133,379	28.7%	48.0%	97.1%	96.0%	93.2%	0.52
S6 UCSN-first prices	75.59	59.85–97.67	106	8,970	28.7%	36.8%	59.2%	64.5%	45.2%	0.29

Runs without the optimiser; scenarios A–G as in the reference run. Source: AICTEM TIPRAS engine, simulation runs S0–S6 (outputs/simulations/simulations.csv) (report/build/data/t_sim_results.csv).

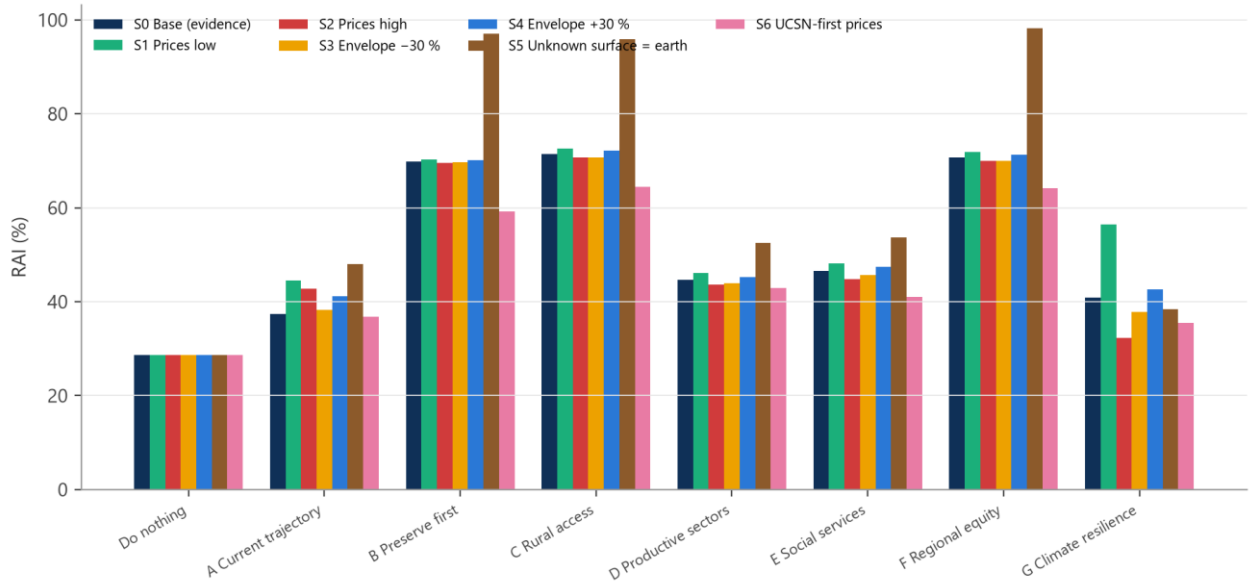

Figure 12.9: Simulations: Rural Access Index at the horizon for every scenario under each simulation

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Figure 12.10: Simulations × scenarios: RAI at the horizon (%) in every combination

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 12.6: Simulation results: what scenarios B and C fund

Simulation	Allocated, B (bn UGX)	Projects funded, B	Allocated, C (bn UGX)	Projects funded, C	Poor km left, B
S0 Base (evidence)	17,076	40,486	11,277	8,611	11,204
S1 Prices low	15,929	40,895	11,275	8,707	0
S2 Prices high	17,612	40,019	11,324	8,559	20,204
S3 Envelope −30 %	12,268	40,144	7,896	8,522	29,204
S4 Envelope +30 %	21,910	40,716	14,657	8,689	0
S5 Unknown surface = earth	17,081	127,127	11,346	19,815	11,204
S6 UCSN-first prices	12,918	6,876	11,294	4,045	70,656

Source: AICTEM TIPRAS engine, simulation runs S0–S6 (outputs/simulations/simulations.csv) (report/build/data/t_sim_allocation.csv).

Four findings stand out.

1. The access outcome is robust to prices and money; the cost is not. Moving every price to its low or high bound changes the programme from 56.6 tn to 125.3 tn — a factor of 2.2 — but moves the RAI under C only from 70.7 % to 72.6 % (Figure 12.11, Figure 12.13). Under C the money the rule can usefully spend is the binding quantity, not the price of each project: C spends about 11.3 tn in every price simulation and funds between 8,559 and 8,707 projects, because its candidates are short, cheap upgrades whose cost per person reached stays low even at the high bound. Moving the envelope 30 % either way changes the money C spends (7.9 tn and 14.7 tn) and the projects it funds (8,522 and 8,689) far more than the outcome: the RAI moves from 70.7 % to 72.1 %. The 3.4 tn added in S4 buys 0.6 points of RAI; the 3.4 tn removed in S3 costs 0.8 points. Most of the access gain the network can deliver within the evidence is already bought in the base case; additional money mainly buys condition and standard (B treats all the agency poor km in S4 and leaves 29,204 km in S3).

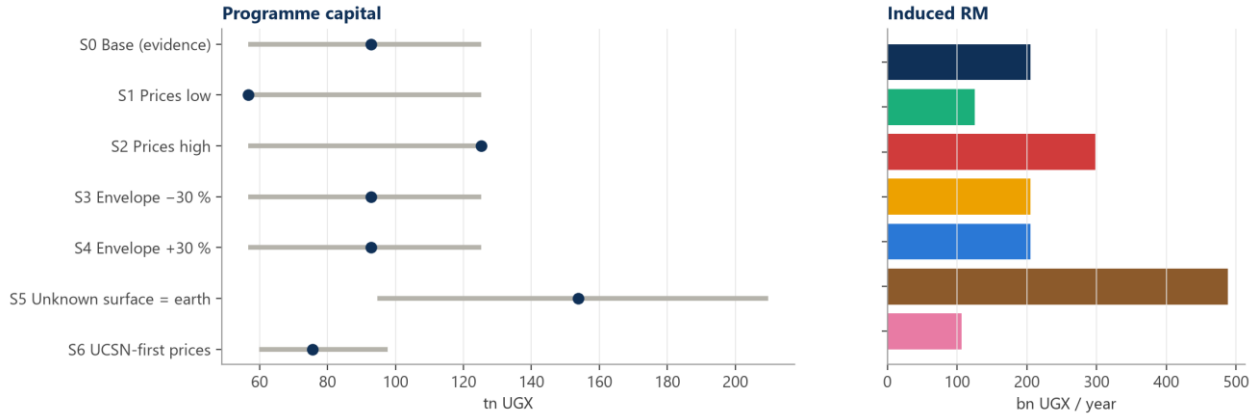

Figure 12.11: Simulations: the candidate programme's capital (low · central · high) and induced maintenance

Source: AICTEM TIPRAS engine, reference run 014c9cd7

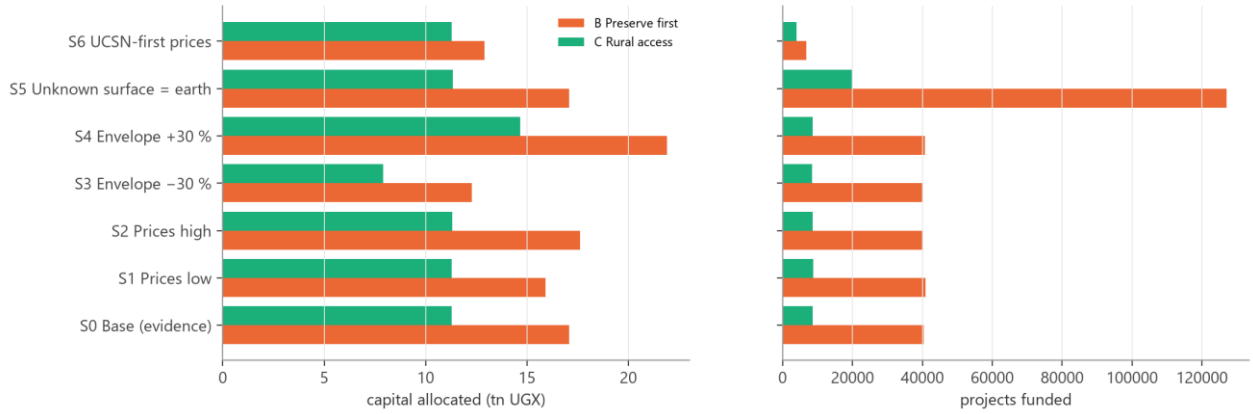

Figure 12.12: Simulations: capital allocated and projects funded under scenario B and scenario C

Source: AICTEM TIPRAS engine, reference run 014c9cd7

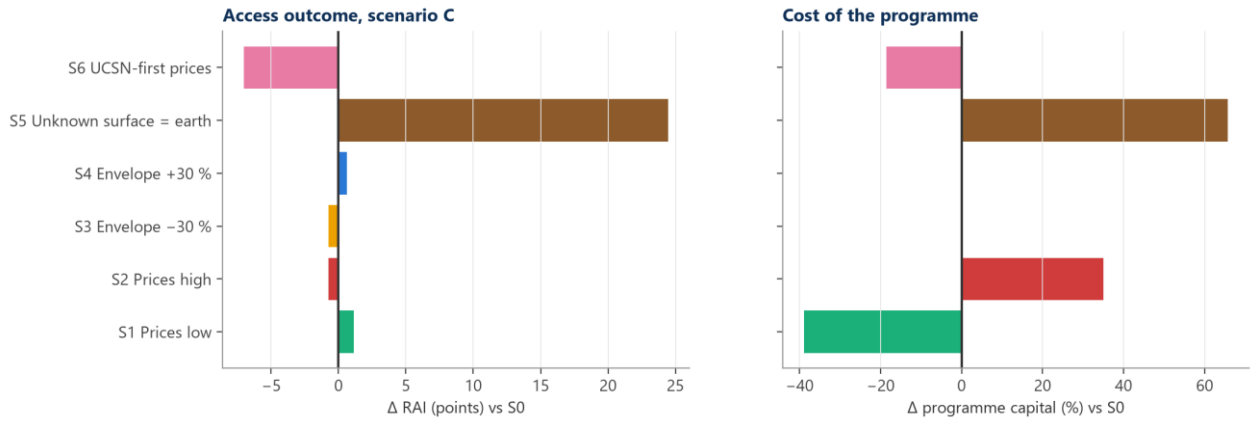

Figure 12.13: What moves the result: change in scenario C's RAI and in programme capital against the base simulation S0

Source: AICTEM TIPRAS engine, reference run 014c9cd7

2. Maintenance affordability is the most price-sensitive result. The steady-state maintenance funding ratio of scenario B is 0.58 in the base case, 1.98 with low prices and 0.30 with high prices. The reason is the routine-maintenance price for unpaved roads: its documented range runs from 1.0 to 8.9 M UGX per km a year, around a central 4.4 M, and the requirement is that price times about 70,000 km of maintainable unpaved road. A larger envelope helps (0.75 in S4) but does not close the gap; a smaller one widens it (0.41 in S3). Whether Uganda can maintain the network it has is therefore the question the next round of cost evidence must answer first: a Road Fund unit-rate series for routine maintenance by road class would narrow the range more than any other single dataset.

3. The largest uncertainty is not money or price but the unknown surface. In S5 the 169,343 km of gap links whose surface is unknown are treated as earth. The engine then proposes works on them: the candidate programme grows from 41,867 to 133,379 priced projects and from 92.8 tn to 153.7 tn (range 94.6–209.7 tn), with induced maintenance of 488 bn a year. The same 11.3 tn under C funds 19,815 projects and lifts the RAI to 96.0 %, and the poorest quintile to 93.2 % (Figure 12.14). This is not a forecast: it shows how much of the access gap in the base case is a *survey* gap rather than a works gap. If most of the unknown roads are in fact earth and short, cheap engineered-earth upgrades could close most of the access gap; if many are already gravel, the base case overstates the gap. Only an inventory survey can tell which, and Figure 12.15 shows that the answer differs most in the north and east.

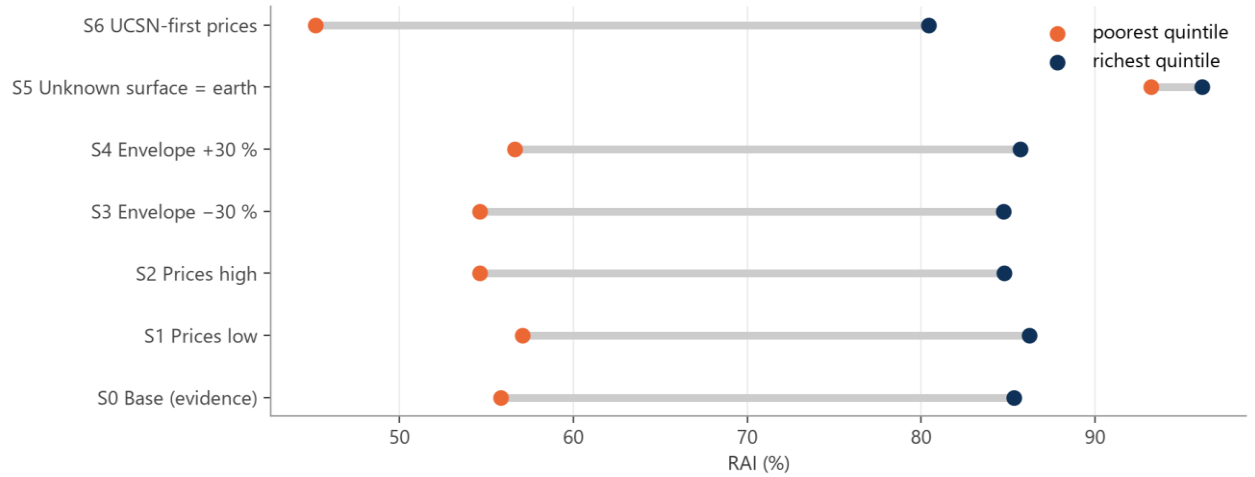

Figure 12.14: Simulations: rural access of the poorest and richest quintiles under scenario C

Source: AICTEM TIPRAS engine, reference run 014c9cd7

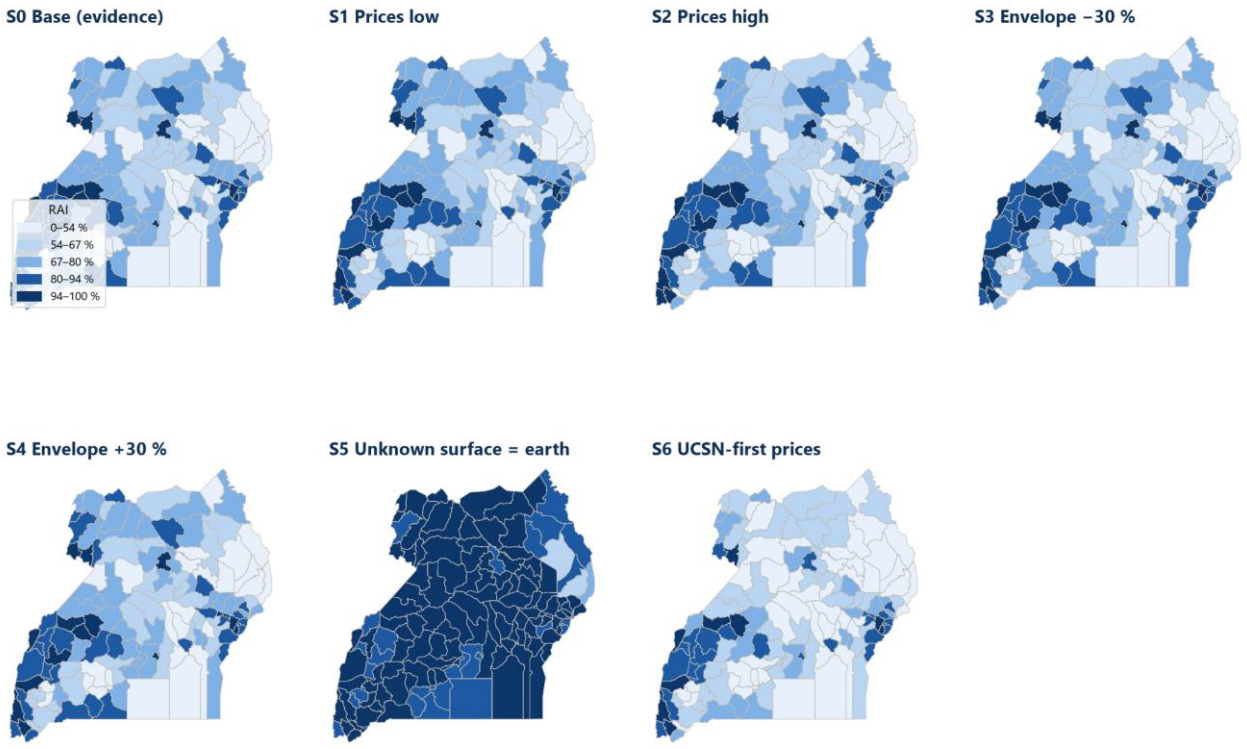

Figure 12.15: Simulations: Rural Access Index at the horizon under scenario C, by district — one shared scale

Source: AICTEM TIPRAS engine, reference run 014c9cd7

4. The price map matters more than the price level. S6 replaces the evidence codes with the UCSN-first map. UCSN publishes no code for engineered-earth upgrading or for rehabilitating earth roads, so those works become unpriced: only 8,970 projects are priced, the community condition programme (70,656 km) cannot be funded, and the programme's apparent cost falls to 75.6 tn — not because roads are cheaper but because most of the works disappear from the costed list. C then funds 4,045 projects and reaches 64.5 % instead of 71.5 %; B reaches 59.2 % instead of 69.8 %. The difference between S0 and S6 is the value, in access terms, of pricing the works that the UCSN catalogue does not yet cover. It is the reason the evidence prices were built, and the reason they have been sent to UCSN for review.

Where the envelope matters. Figure 12.16 maps how scenario C's allocation per person changes by district when the envelope moves 30 % down or up. The districts that gain most from more money are not those that lose most from less: with less money C drops the marginal, more expensive upgrades, which are spread across the west and north; with more money it adds them back and extends into districts where the cost per person reached is higher.

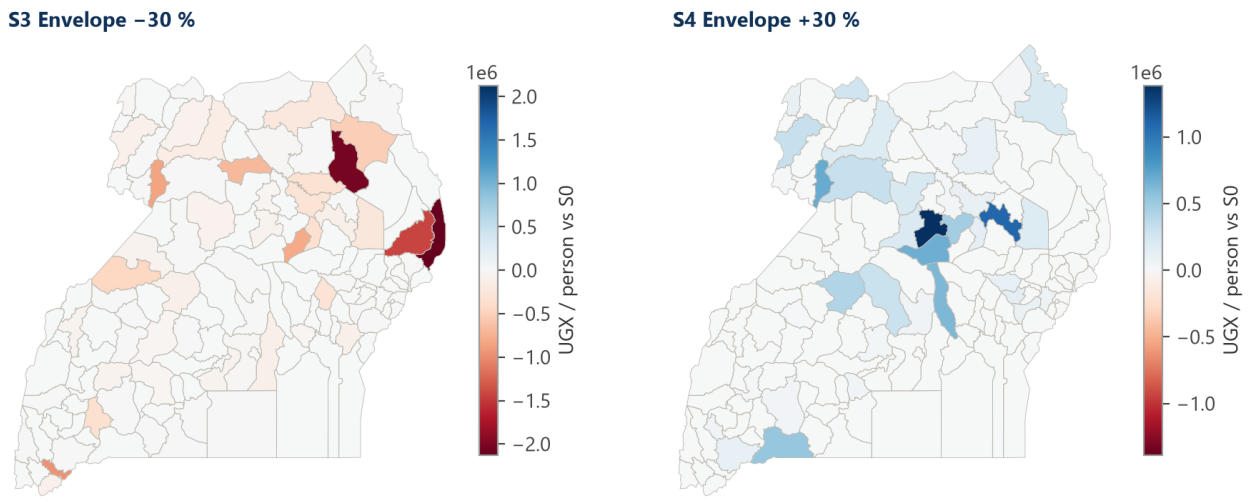

Figure 12.16: Simulations: where a 30 % larger or smaller envelope changes scenario C's allocation per capita

Source: AICTEM TIPRAS engine, reference run 014c9cd7

> **Reading the simulations.** The simulations do not rank the scenarios and do not make any of them more or less advisable. They show which inputs the conclusions depend on. Across all seven, the grouping of the rules by access outcome is stable — rural access (C), regional equity (F) and preserve-first (B) form a leading group within 1–5 points of each other in every simulation (C first in six, F first in S5), and A, D, E and G a group 20–60 points below — and the conclusions that do move are the cost of the programme (with prices and with the unknown surface) and the affordability of maintenance (with the routine-maintenance price).

12.10.2 What the simulations do not cover

They change one input at a time and do not combine changes; they do not vary the norms (speeds, access standards, value of time), which are drafts; they do not run the optimiser; and they inherit every limitation of the reference run listed below. The horizon could not be extended beyond 2030 because the MTEF ceilings stop in 2031. The registry of the runs is `outputs/simulations/simulations.csv`; each run's tables are kept under `outputs/<run_id>/`.

12.11 Data quality and limitations

- **Surface convention.** Every RAI and every access gain assumes that unknown surface is earth. Surface is known on only about 32 % of network km. Under the opposite convention the base RAI would be 93.5 % and most upgrading gains would vanish. The gains in this chapter are therefore conditional, and the base-year figure cannot be compared directly with the 53.1 % benchmark of 2015.
- **Horizon only.** TI-09 evaluates 2030, not each year. The year-by-year network path of the manual, which would show settlements served as works complete, is not built.
- **No deterioration.** Condition is unknown per link, so the do-nothing path does not deteriorate. A plan that holds condition would show no gain against doing nothing in this run, although it would against a deteriorating path.
- **Mobility, freight, safety, emissions and economics are not yet reported per scenario.** The manual's person-hours, vehicle operating cost, freight cost by sector, crashes, emissions and programme NPV are not in the scenario table of this run. The appraisal behind them uses uncalibrated vehicle operating cost coefficients, generic speeds marked "NOT Ugandan" and a placeholder value for wet-season closures.

- **No ranges on outcomes.** Cost uncertainty (low, central, high) is carried in costs but not propagated to outcomes. Access outcomes do not depend on prices, but which projects each scenario funds does.
- **Wealth quintiles.** IWI is a survey-based wealth index attached to the nearest point within 5 km. It is an equity weight, not a poverty statistic, and cells without a value are excluded.
- **Population.** UBOS 2022 district totals carried forward at an assumed 3 % a year.
- **Service standards are draft norms.** The thresholds (2 km, two hours to EmONC and so on) are draft, mostly assumed values. A different threshold changes every service-class count.
- **Scenario A is a fallback.** Without historical allocation data, A funds national-class projects first. Its outcomes do not represent the ministry's actual practice.

12.12 Implications for planning

- **An access-led allocation of about 11.3 tn UGX could raise the RAI from 28.7 % to about 71 % by 2030 under the engine's surface convention.** That conclusion is only as good as the surface data. Recording the surface of the 68 % of km where it is unknown would settle how much of the gain is real.
- **The poorest gain most from access-led and equity-led allocations.** C, F and B narrow the wealth gap from 42.2 to about 29–30 points. The fallback current-trajectory allocation widens it to 46.0. The distribution of gains, not only the total, should be part of every comparison.
- **Health and school access gains come mainly from a social-services rule, and are limited by facility siting.** E doubles the primary-health gain of C, but most of the remaining EmONC gap is "no destination". Joint planning with HealthModellingx on facility siting will do more there than roads alone.
- **District results differ by rule even when national totals are close.** B and C differ by 1.7 points nationally but by 27 points in Amudat. District profiles should be shown for each allocation.
- **Project-level access claims must not be summed.** Interaction adjustments of 21–87 % mean that programme targets need network recomputation.
- **The condition outcome needs link-level data.** Poor km treated is an accounting of agency statistics. A condition survey would turn it into a projected network state.

CHAPTER 13

13 Budget and the medium-term expenditure framework

An allocation is not yet a budget. A finance ministry needs to see the plan in its own chart of accounts, by fiscal year, by programme and by economic class, and to know what the plan commits future budgets to once the new roads have to be maintained. This chapter answers two questions. First, how does a scenario look when it is written as a programme budget and a medium-term expenditure framework (MTEF)? Second, can Uganda afford to maintain the network it has, and the one the plan would add? The TI-10 module answers them. It also writes the files that carry the transport requirement to the fiscal models. The chapter implements design chapter 12 (`docs/12_TI10_BUDGET_MTEF.md`). It begins with the evidence on which the envelope rests: the approved roads budgets for FY2024/25 to FY2026/27, classified line by line in `data/reference/budget_roads_evidence_UGA.md`. The budget is shown for **scenario B (preserve first)**, the default scenario TI-10 is run on. That choice is a convention of the engine for this report, not a recommendation: any allocation can be passed to TI-10, and every allocation would be submitted in the same form. All figures are from the reference run `014c9cd7`, in constant 2026 UGX, and are **exploratory**.

13.1 What the module does

TI-10 takes the allocation of one scenario from TI-08 and the routine maintenance requirement from TI-06, and produces:

- **three budget views** of the same money, which must reconcile exactly:
  - a **programme view** by programme code (the budget submission);
  - an **economic view** by Government Finance Statistics (GFS) class (for the fiscal model and the IMF programme);
  - an **administrative view** by vote and cost centre (for appropriation and district grants);
- a **fiscal sustainability statement**: the routine maintenance requirement of the existing network, plus the maintenance induced by the plan's assets, against the resources that may legally fund maintenance;
- **district layers** of the capital budget and of the maintenance requirement, mapped separately and never summed;
- **outbound files** for MTEFModelling and FINEX, and access files for the health, education and agriculture models (Chapter 14).

The reconciliation of the three views is asserted in code: if they do not agree, the run stops.

13.2 The evidence: Uganda's approved roads budget

The envelope used in Chapter 11 rests on two numbers taken from the approved budget estimates: the roads share of transport and the split of roads money between pools. Both were derived from the item-level approved estimates for three years, held in `D:\EDIP\Data\Budgetdata`.

13.2.1 How the roads budget was identified

- **The transport total** is programme 09, Integrated Transport Infrastructure and Services (ITIS), summed across all votes and funding columns. Its votes are 016 (Ministry of Works and Transport), 113 (UNRA) and 118 (URF), which are separate votes in FY2024/25 only and are mainstreamed into vote 016 from FY2025/26, 122 (KCCA) and 609 (the central line for local-government grants).

- **Roads lines** are all of votes 113, 118, 122 and 609, the vote 016 sub-subprogrammes for national roads, national and district road maintenance, and district, urban and community access roads, road equipment, and road construction-standards outputs.
- **Not roads:** rail (including the standard gauge railway), air, inland water, ferries and landing sites.
- **Unclear, excluded:** policy and planning support (pensions and arrears), multimodal regulation and non-road construction standards. Together these are 1.2 to 2.7 % of transport.

The check against the MTEF crosswalk is close. MTEFModelling maps transport to programme 09 plus votes 016 and 609, with an FY2026 baseline of 8,830.16 bn UGX. The ITIS total here is 8,827.18 bn. The 2.98 bn difference is vote 016 lines under the Sustainable Urbanisation programme, excluded as not transport.

Table 13.1: Transport and roads in the approved budget estimates (bn UGX)

	FY2024/25	FY2025/26	FY2026/27
Transport (ITIS) total	4,968.2	6,602.1	8,827.2
of which external financing	2,985.1	4,335.2	4,615.8
Roads	3,005.7	3,572.1	4,813.6
Not roads (rail, air, water)	1,900.5	2,849.6	3,814.1
Roads share of transport	0.605	0.541	0.545

Transport spending grew by 78 % in two years, from 4,968.2 to 8,827.2 bn UGX. Roads grew more slowly, by 60 %, from 3,005.7 to 4,813.6 bn. The non-road share rose because of external financing for rail. The roads share of transport fell from 0.605 to 0.545. The engine uses the FY2026/27 figure, 0.545: it is the latest approved year, it is the MTEF base year, it reflects the post-mainstreaming vote structure, and FY2025/26 (0.541) agrees with it.

13.2.2 The roads budget by pool

Figure 13.1 and Table 13.2 split the roads budget into the pools TIPRAS uses.

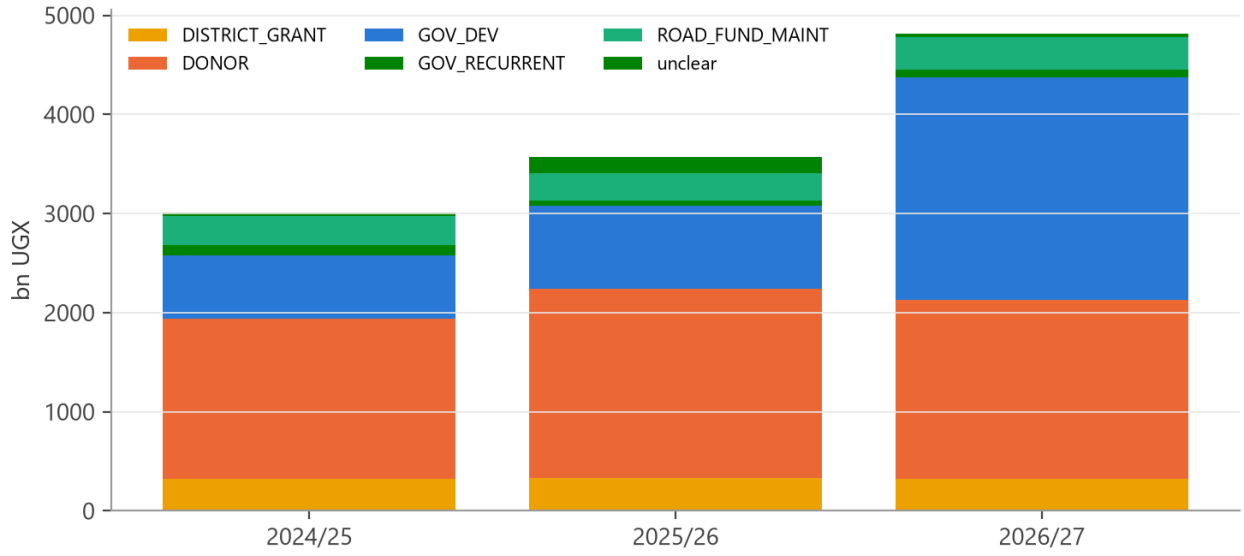

Figure 13.1: The approved roads budget by funding pool, FY2024/25–FY2026/27

Source: Approved Budget Estimates (Budgetdata), classified in data/reference/budget_roads_evidence_UGA.csv

Table 13.2: Approved roads budget by funding pool, FY2024/25–FY2026/27 (bn UGX)

Fiscal year	DISTRICT_GRANT	DONOR	GOV_DEV	GOV_RECURRENT	ROAD_FUND_MAINT	unclear	Total (bn UGX)
2024/25	323.4	1,619.5	631.3	105.7	294.7	31.1	3,005.7
2025/26	330.2	1,911.3	840.6	51.8	273.2	165.1	3,572.1
2026/27	327.5	1,804.1	2,239.9	79.6	327.9	34.5	4,813.6

Classified from the approved budget estimates (data/reference/budget_roads_evidence_UGA.md). Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_budget_pools.csv).

The pool rules are:

- **DONOR:** external financing (ExtFin).
- **DISTRICT_GRANT:** all of vote 609 (the 176.0 bn maintenance transfer to district, urban and community access roads, and 44.3 to 50.1 bn of development grants), plus the Road Fund's non-wage transfers to district road agencies (about 100.8 bn a year).
- **ROAD_FUND_MAINT:** in FY2024/25, the Road Fund transfer for national road maintenance (294.7 bn). From FY2025/26, the Ministry's non-wage spending on the output "National Road Maintenance", its successor after the Road Fund was mainstreamed. This is an **inference**: the fuel-levy earmark status of those lines is not shown in the files.
- **GOV_DEV:** the domestic development budget of votes 016, 113 and 122.
- **GOV_RECURRENT:** wages and all other road non-wage spending.
- **unclear:** arrears (31.1, 165.1 and 34.5 bn), which cannot fund new works.

The figure shows three features.

- **Donor money is the largest pool every year**: 1,619.5 bn in FY2024/25, 1,911.3 bn in FY2025/26 and 1,804.1 bn in FY2026/27. As a share of roads spending without arrears it falls from 0.544 to 0.377 in FY2026/27, because domestic development grows.
- **Domestic development jumps in FY2026/27**, from 840.6 to 2,239.9 bn. Much of the jump is one item: Land Acquisition Project II, at 689.5 bn against 150.0 and 113.1 bn in the two earlier years. Without it, the FY2026/27 development share would be 0.379, not 0.469.
- **The maintenance and grant pools are almost flat in shillings.** The road maintenance pool is 294.7, 273.2 and 327.9 bn; district grants 323.4, 330.2 and 327.5 bn; recurrent 105.7, 51.8 and 79.6 bn. While capital spending grows, the money for keeping roads open does not.

Because the split is dominated by lumpy items (the timing of loan disbursements, one-off land acquisition), the engine uses a three-year amount-weighted split, over 11,160.7 bn of roads spending without arrears:

Table 13.3: Pool shares used for the envelope (three-year pooled, quality normed)

ROAD_FUND_MAINT	GOV_DEV	GOV_RECURRENT	DISTRICT_GRANT	DONOR
0.080	0.333	0.021	0.088	0.478

The evidence note records the alternatives. The FY2026/27 split alone would be 0.069 / 0.469 / 0.017 / 0.069 / 0.377. If the Road Fund's transfers to district agencies were counted as road-fund maintenance rather than district grants, the road maintenance share would be 0.107 and district grants 0.061.

13.2.3 What the evidence does not cover

The budget evidence has seven stated caveats. Four matter most for this chapter.

- **Approved estimates, not releases or outturn.** External financing typically under-executes, so the donor pool overstates the cash that will arrive.
- **Only on-budget external finance.** Off-budget donor and PPP finance is not captured.
- **Local-government coverage is partial.** Only vote 609 transfers are captured, not local own-source revenue or other grants spent on roads.
- **Possible double counting.** It is not verified whether vote 609's 176.0 bn transfer overlaps with the Road Fund's 100.8 bn transfer to district agencies. If it does, the district grant share would fall by up to about 0.03.

Road-related lines outside ITIS, such as the urban development programmes (USMID II and UCMID, 88.0 / 86.5 / 22.8 bn of external finance; the Greater Kampala urban development project, 444.4 / 1,177.8 / 310.7 bn) and farm access roads under the agriculture ministry (6.2 / 30.2 / 32.3 bn), are recorded but not counted.

13.3 Method

13.3.1 From plan years to fiscal years

Uganda's fiscal year runs from July to June. The design maps plan-year amounts to fiscal years by monthly phasing weights:

$$y_{\phi}^{FY} = \sum_{t} y_t \sum_{m \in t \cap \phi} w_{t,m}, \qquad \sum_{\phi} y_{\phi}^{FY} = \sum_{t} y_t \tag{13.1}$$

where y_t is the amount in plan year t, ϕ a fiscal year and $w_{t,m}$ the share of year t's spending in month m. The current implementation uses the simplest case: plan year t is fiscal year $t/t{+}1$ in full (2027 is FY2027/28), so totals are preserved exactly but no split by month is made.

13.3.2 Three views that reconcile

Every cost line is mapped to a programme code, a GFS economic class and an administrative code. The mappings are data in the design; in the current code they are small tables:

- works RM and PM map to programme `ITIS.roads_maintenance`, all other works to `ITIS.roads_development`;
- RM and PM map to GFS class 22 (use of goods and services), capital works (RH, UP, SI, ST, NC) to class 31 (non-financial assets);
- the administrative code is a placeholder cost centre per district (`vote.roads/<district>`), with routine maintenance on a national line.

The reconciliation identity is asserted for every fiscal year ϕ and pool π:

$$\sum_{o} Prog_{o,\phi,\pi} = \sum_{e} Econ_{e,\phi,\pi} = \sum_{a} Admin_{a,\phi,\pi} \tag{13.2}$$

13.3.3 The MTEF outer years and the fiscal sustainability statement

In the design, each MTEF year carries the cash of funded projects, routine maintenance of the whole network, the induced maintenance of every asset completed by that year, and debt service or availability payments supplied by FINEX:

$$MTEF_{\phi} = \sum_{j} cash_{j,\phi} + M_{\phi} + \sum_{j\,:\,t_c(j) < \phi} IM_{j,\phi} + DS_{\phi} + AP_{\phi} \tag{13.3}$$

The fiscal sustainability statement is the steady-state version: the annual maintenance commitment once every plan asset is in service, against the resources that may legally fund it:

$$M^{ss} = M^{exist} + \sum_{j} IM_j, \qquad mfr^{ss} = \frac{R^{maint}}{M^{ss}} \tag{13.4}$$

where M^{exist} is the routine maintenance requirement of the existing maintainable network, IM_j the induced routine maintenance of funded project j and R^{maint} the road-fund and recurrent resources in the final plan year. A ratio below 1 means the network, as planned, cannot be maintained from maintenance money.

13.4 Results: scenario B as a budget

Figure 13.2 and Table 13.4 show scenario B in the national chart of accounts, by programme code and fiscal year.

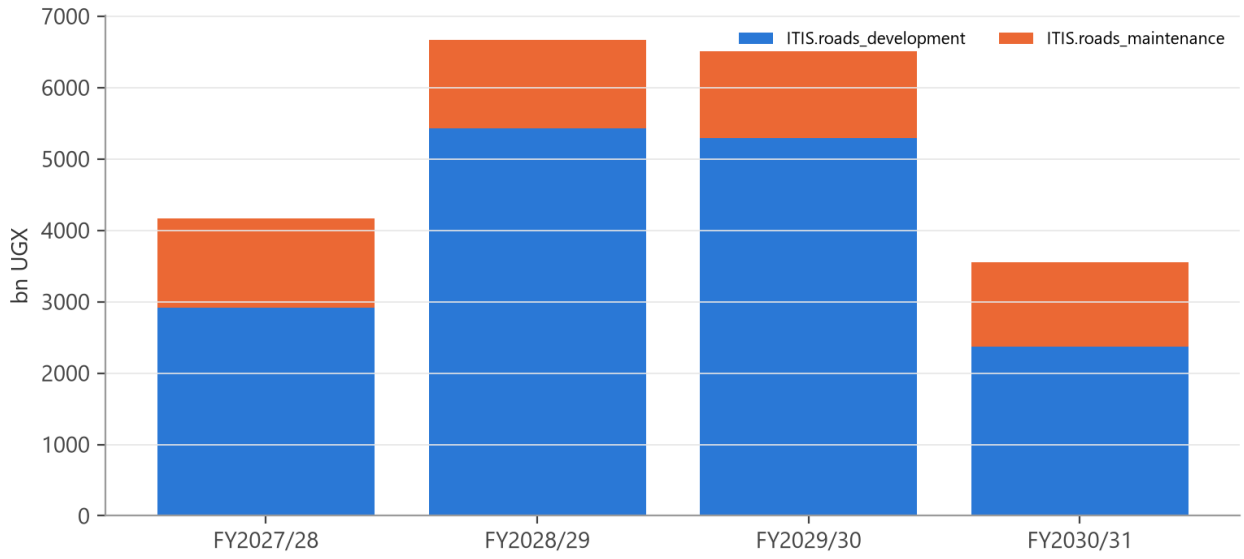

Figure 13.2: Scenario B in the national chart of accounts: capital by programme code and fiscal year

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 13.4: Scenario B capital by programme code and fiscal year (bn UGX)

Programme code	FY2027/28	FY2028/29	FY2029/30	FY2030/31
ITIS.roads_development	2,921.7	5,426.1	5,289.4	2,367.7
ITIS.roads_maintenance	1,243.7	1,244.4	1,215.4	1,186.4

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_budget_programme.csv).

Table 13.5: Scenario B budget by fiscal year against the envelope (bn UGX)

Fiscal year	Roads development	Roads maintenance	Total	Envelope	Share used
FY2027/28	2,921.7	1,243.7	4,165.4	6,684.5	62 %
FY2028/29	5,426.1	1,244.4	6,670.5	6,693.5	100 %
FY2029/30	5,289.4	1,215.4	6,504.8	6,528.3	100 %
FY2030/31	2,367.7	1,186.4	3,554.1	5,861.9	61 %
Total	**16,004.9**	**4,889.9**	**20,894.8**	**25,768.2**	**81 %**

- **Development peaks in the two middle years**, at 5,426.1 bn in FY2028/29 and 5,289.4 bn in FY2029/30, and is about half that in the first and last years. This is the S-curve phasing of three-year works that all start in the first plan year (Chapter 11).
- **Maintenance is steady at 1.19 to 1.24 tn a year.** It has two parts. Routine maintenance on the existing network is 954.7 bn every year, labelled `ROAD_FUND_MAINT+` because it is drawn from the road-fund, recurrent and district-grant pools together. Periodic maintenance programmes add 289.0, 289.7, 260.7 and 153.2 bn from district grants, and 78.5 bn from domestic development in FY2030/31.
- **In the two peak years B uses all but 0.3 % of the envelope.** In the first and last years it uses about 60 %. The leftover is donor money that the S-curve cannot draw in time (Chapter 11).
- **The peak is large against current budgets.** Scenario B's FY2028/29 total, 6,670.5 bn, is 1.39 times the FY2026/27 approved roads budget of 4,813.6 bn. It fits the envelope only because the MTEFModelling baseline ceiling for transport in FY2028/29 (12.28 tn) is itself 1.39 times the FY2026/27 ITIS budget.

Every capital line in this table carries its induced maintenance, which appears in the fiscal sustainability statement below rather than in the four MTEF years. In the MTEF period itself, induced maintenance is small because most assets come into service only in the last years.

The budget is in constant 2026 UGX. The conversion to nominal prices with FINEX deflators (design equation 10.4) is not yet applied. Debt service and availability payments are not computed in TIPRAS; they are FINEX's to supply, and no FINEX return exists yet.

13.5 Results: three views, one total

Figure 13.3 and Table 13.6 show that the three views of scenario B sum to the same total.

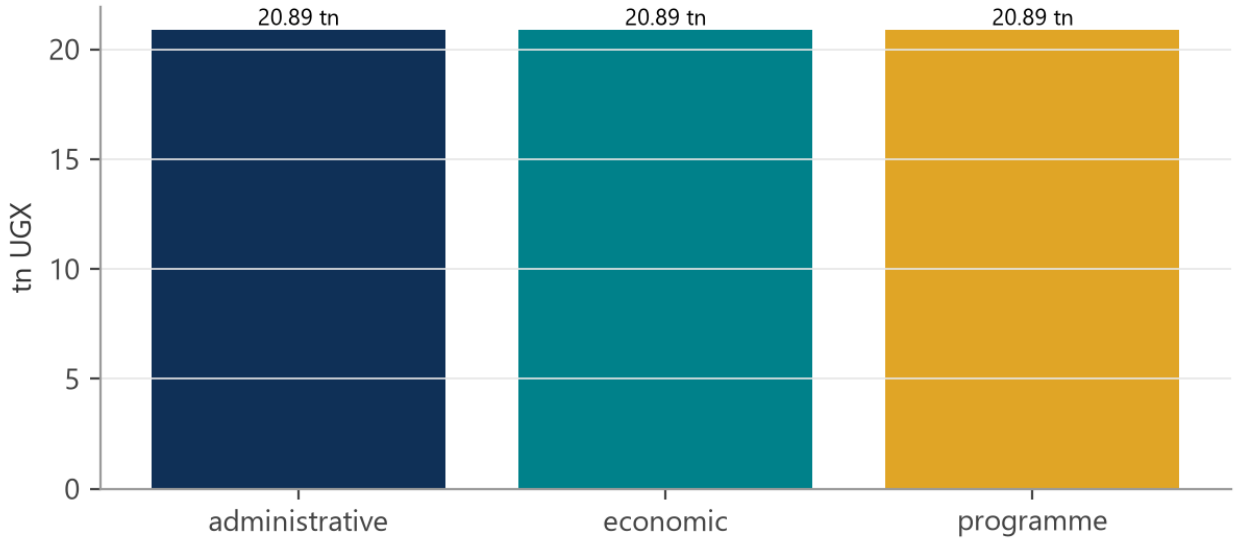

Figure 13.3: The three budget views of scenario B reconcile to the same total

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 13.6: Scenario B budget: the three views reconcile

View	Total (bn UGX)
administrative	20,894.9
economic	20,894.9
programme	20,894.9

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_budget_views.csv).

Each view totals **20,894.9 bn UGX** over FY2027/28 to FY2030/31. The total is B's capital (17,076.1 bn) plus four years of routine maintenance (4 × 954.7 = 3,818.8 bn). The reconciliation holds per fiscal year and per pool, not only in total; TI-10 asserts it and the run log records "budget views reconcile".

In the economic view, 16,004.9 bn falls in class 31 (non-financial assets) and 4,889.9 bn in class 22 (use of goods and services). The 22 % share of goods and services is the maintenance share of the plan. For FINEX and the IMF programme, that split matters: class 31 is investment, class 22 is consumption, and the two have different import content and different macroeconomic effects (Chapter 14).

The reconciliation is a necessary check, not a sufficient one. Two parts of the mapping are still placeholders:

- **Programme codes** `ITIS.roads_development` and `ITIS.roads_maintenance` are marked in the code "to confirm against the national crosswalk". The design requires every submission code to exist in the MTEFModelling crosswalk `uganda_crosswalk_FY2026_27.csv`. That test is not yet run.

- **Administrative codes** are one cost centre per district, not the real votes (UNRA stations, the Ministry's departments, the local governments). A real administrative view needs the map from works, network class and agency to vote.

The design also requires a physical output on every budget line, such as "km of national paved road periodically maintained" or "rural population newly within 2 km of an all-season road". The output columns exist in the budget table but are empty in this run. Attaching TI-09's outcomes to each line is the next step for a submission-ready budget.

13.6 Results: can the network be maintained?

Figure 13.4 and Table 13.7 give the fiscal sustainability statement for scenario B.

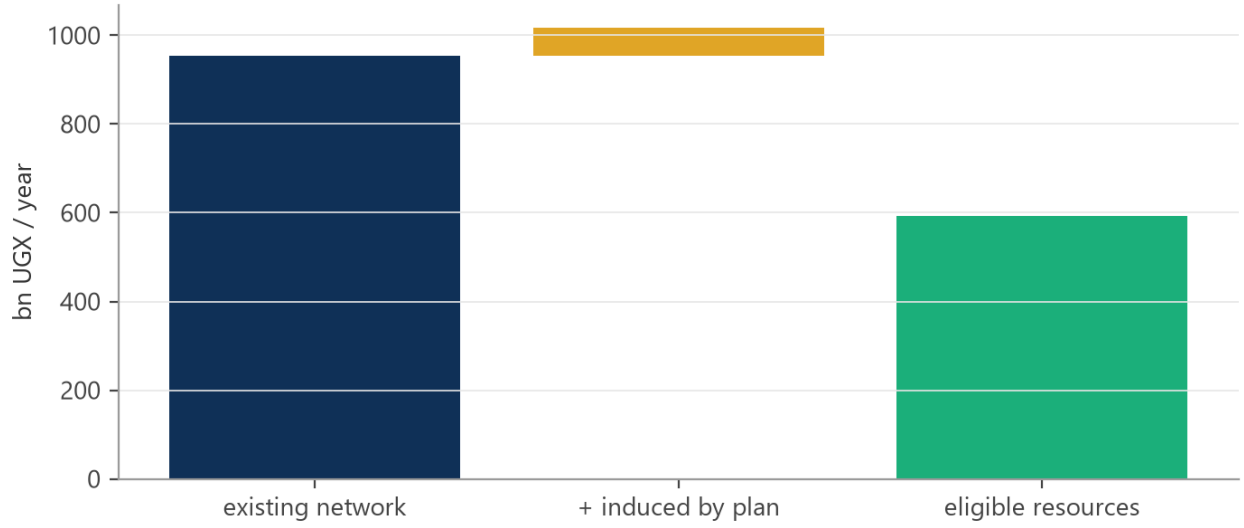

Figure 13.4: Can the network be maintained? Routine maintenance requirement against maintenance-eligible resources (steady state)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 13.7: Fiscal sustainability statement, scenario B (steady state)

Item	Value
maintenance requirement, existing network (RM, LCU/yr)	0.95 tn
plus induced RM of this plan's assets (LCU/yr, steady state)	0.06 tn
= total routine maintenance requirement (steady state)	1.02 tn
maintenance-eligible resources, final plan year	0.59 tn
maintenance funding ratio (steady state)	0.58

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_fiscal.csv).

- **The existing network needs 954.7 bn UGX a year** of routine maintenance at the evidence prices (routine maintenance of unpaved roads at 4.4 M UGX per km a year and of paved roads at 16.4 M UGX per km a year, both tier E5).
- **Scenario B's assets add 62.7 bn a year** once all are in service. The total steady-state requirement is **1,017.4 bn a year**.
- **The resources that may fund it are 592.1 bn** in the final plan year: the road-fund maintenance pool (469.0 bn in 2030) plus the recurrent pool (123.1 bn).
- **The maintenance funding ratio is 0.58.** Each year about 425 bn UGX of routine maintenance would go unfunded from maintenance money.

The allocation step in Chapter 11 does fund 954.7 bn a year of routine maintenance, but only by also drawing on district grants, which are eligible for routine maintenance. In 2030 the three maintenance-eligible pools together

hold 1,107.9 bn, which would cover the 1,017.4 bn requirement. But that would take 425.3 bn of the 515.8 bn district-grant pool, leaving 90.5 bn for the district and community roads that the grant exists to improve. The statement therefore counts only the road-fund and recurrent pools as maintenance resources. Whether district grants should be counted is a policy question; the engine shows both readings.

The shortfall is not new. Chapter 9 shows a condition backlog of 7.25 tn UGX at central prices on the roads reported in poor condition, most of it on community roads (70,656 poor km, 5.61 tn) and district roads (35,320 poor km, 1.02 tn). A network maintained at 58 % of its requirement adds to that backlog every year. The design asks for this shortfall to be shown as a condition trajectory: the km that fall to poor and the asset value lost. That cannot be done in this run, because deterioration is not modelled without link condition.

Two points help to read the ratio. First, the approved road maintenance pool was 327.9 bn in FY2026/27, against 534.8 bn in the 2027 envelope: the envelope already assumes that maintenance money grows with the transport ceiling. Second, an earlier run with the illustrative UCSN routine-maintenance coefficients (10 M UGX per km a year on unpaved roads, tier E7) gave a requirement of 2,049.6 bn a year and a ratio of 0.29. The move to evidence prices roughly halved the requirement, but it did not close the gap.

The steady-state row is the one that should determine approval. A plan whose capital fits the envelope but whose maintenance does not is not affordable: it adds roads that will fall back into the backlog.

13.7 Results: where the budget and the maintenance fall

Figure 13.5 maps, side by side, scenario B's capital budget by district and the routine maintenance requirement by district. They are two maps, never summed into one.

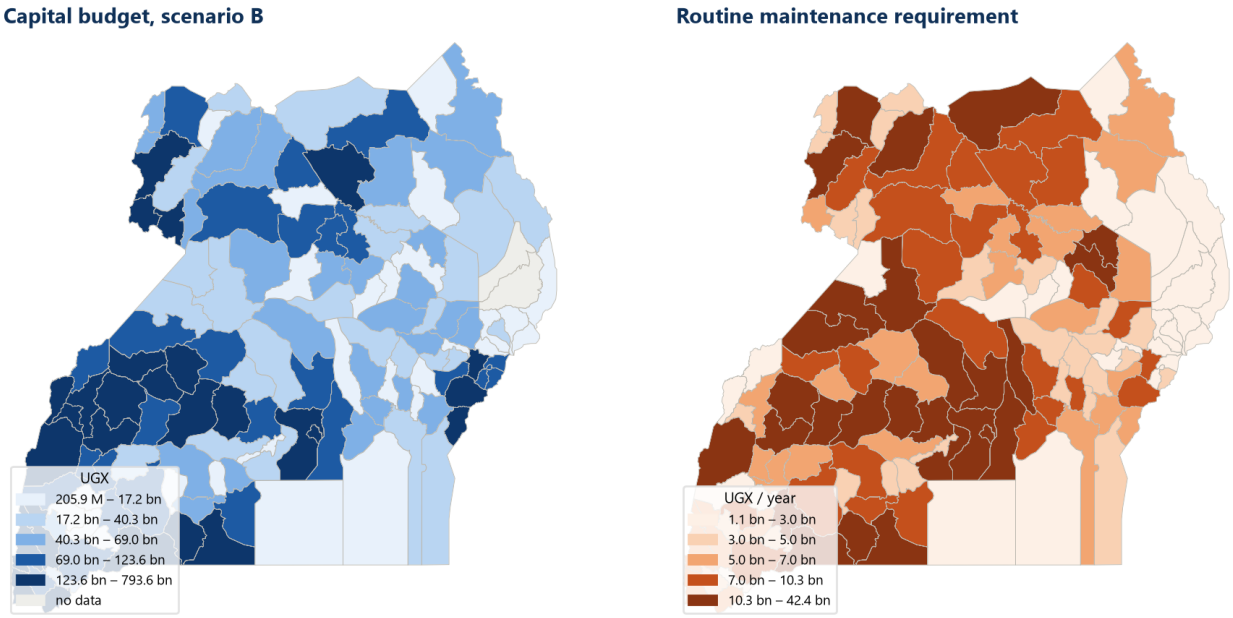

Figure 13.5: Scenario B capital budget by district (left) and the routine maintenance requirement of the existing network (right)

Source: Simulation S0 (reference configuration: draft-f norms, evidence prices)

A note on the source. The two panels are drawn from simulation S0, which has exactly the reference run's configuration (norms `draft-f`, evidence prices). Its district values reproduce the reference run's totals: scenario B's capital budget of 17,076.1 bn (133 districts plus 5,806.5 bn of national programmes without a district) and the routine maintenance requirement of 954.7 bn a year.

The panels show:

- **Capital is concentrated in a few districts.** The largest district budgets are Wakiso (793.6 bn), Kasese (566.4 bn), Mbale (331.6 bn), Kyenjojo (296.0 bn), Kampala (276.3 bn); the median district receives 59.1 bn, and the smallest are Buhweju (1.7 bn), Rwampara (1.1 bn), Buvuma (0.2 bn).
- **The maintenance requirement follows network length.** The largest requirements are Wakiso (42.4 bn), Mukono (23.5 bn), Kyenjojo (20.2 bn), Yumbe (17.5 bn), Arua (17.2 bn); the smallest are Kalangala (1.3 bn), Nabilatuk (1.2 bn), Bududa (1.1 bn).
- **The two patterns overlap.** Wakiso and Kyenjojo are high on both maps: they would receive large capital programmes and already carry large maintenance claims, which grow when the plan's assets come into service.

The district maintenance map is the basis for the requirement-based district grant the design calls for. Today's district maintenance transfer (vote 609) is allocated by a formula outside this engine; the map shows what a requirement-based allocation would look like.

13.8 Results: who finances scenario B, and who maintains it

Scenario B's capital of 17,076.1 bn is financed 8,580.8 bn (50.3 %) from domestic development money, 7,484.2 bn (43.8 %) from donor project finance and 1,011.1 bn (5.9 %) from district grants (Chapter 11). Its routine maintenance is financed entirely from domestic pools. The design (§10.6) asks TI-10 to track project loans by disbursement profile and to show the domestic maintenance requirement rising as each donor-financed asset is completed, because a corridor built with a loan becomes a domestic maintenance liability on the day it opens.

That tracking is not yet implemented. What the run does show is the end point: the 62.7 bn a year of induced maintenance in the fiscal statement is the domestic claim of all B's new assets, whoever financed them. Almost half of the capital that creates this claim is donor money, and none of the money that must meet it is.

Two further fiscal items of the design are not computed here, by intent. Debt service on the project loans and availability payments on any PPP road belong to FINEX, which receives the requirement and returns them. TIPRAS does not compute a macroeconomic effect of the programme either; that is CDCGE's answer (Chapter 14).

13.9 Outbound to the fiscal models

TI-10 writes two files for the fiscal models. Each carries the shared exchange header (country, run id, scenario id, schema version `tipras-1`, producer and exploratory flag):

- `data/exchange/outbound/mtef/transport_mtef_UGA_<date>.csv`: all three views by code, fiscal year and pool: 749 lines, of which 17 are programme lines, 17 economic lines and 715 administrative lines across 134 cost centres (133 districts and the national routine-maintenance line);
- `data/exchange/outbound/finex/transport_requirement_UGA_<date>.csv`: the economic view by fiscal year, GFS class and pool, 17 lines.

Both carry `exploratory = true`. The FINEX file does not yet carry the separate maintenance channel, the import-content share or the loan-disbursement and PPP columns of the design contract (Chapter 14). No ceiling has come back from MTEFModelling in response, and the bounded fiscal loop of the design, between TI-08, TI-10 and FINEX, is not run.

13.10 Data quality and limitations

- **The pool shares are evidence; the envelope is a model.** The shares come from approved estimates (quality normed). The ceiling comes from MTEFModelling's baseline and is 39 % above the FY2026/27 approved roads budget in the first plan year.

- **Approved, not executed.** External finance usually under-executes, so the donor pool overstates cash. The road maintenance pool from FY2025/26 is an inferred successor of the Road Fund transfer.
- **Fiscal years.** Plan years are mapped whole to fiscal years, with no monthly phasing.
- **Constant prices.** There is no conversion to nominal prices yet.
- **Placeholder codes.** Programme codes are still to be confirmed against the national crosswalk, and administrative codes are per-district placeholders, not votes.
- **No physical outputs** are attached to budget lines in this run.
- **No ceiling response.** The design re-runs TI-08 when an MTEF ceiling binds and shows what is given up. Because the envelope is derived from the ceiling, the ceiling cannot bind in this run, and the response is not exercised.
- **Maintenance prices.** The routine maintenance requirement rests on tier E5 evidence prices. Changing them moves the funding ratio directly: the illustrative coefficients gave 0.29.
- **The district maps** come from simulation S0, which reproduces the reference run's totals exactly.

13.11 Implications for planning

- **Maintenance, not capital, is the binding fiscal constraint.** The steady-state routine maintenance requirement is 1.02 tn UGX a year against 0.59 tn of maintenance resources, a ratio of 0.58. Every capital decision should be presented with the maintenance line it creates.
- **The maintenance pools are flat while capital grows.** Road maintenance, recurrent and district-grant money stayed roughly flat over three approved years, while development money grew. A plan that grows the network needs a decision on maintenance financing, such as the share of fuel-levy revenue that reaches road maintenance.
- **Counting district grants as maintenance money closes the gap on paper but empties the grant.** That trade-off should be made explicitly, not by accounting.
- **The budget can already be produced in the ministry's format**, with three reconciling views. The codes must be checked against the national crosswalk, and physical outputs attached, before it is used in a submission.
- **The peak years ask for 1.39 times the current roads budget.** The plan is affordable only if the MTEF ceiling grows as MTEFModelling's baseline assumes and if donor finance disburses as approved.
- **The budget evidence should be refreshed with releases and outturn**, not only approved estimates, and the vote 609 and Road Fund district transfers should be reconciled.

CHAPTER 14

14 Integration with the sector and fiscal engines

TIPRAS is one engine in a family. It takes geography, inventory, prices, sector facilities and a fiscal envelope from other systems in the estate. It gives back an expenditure requirement that changes the fiscal and macroeconomic picture, and changes in access and transport cost that the health, education, agriculture, mining, industry and trade models depend on. This chapter describes, for the Uganda build, which of those links exist, which carry real data, which run on a flagged fallback and which are still only designs. It implements design chapter 13 (`docs/13_INTEGRATION.md`); the sector-by-sector linkages are in Chapter 22. The central rule is simple: TIPRAS **consumes** the results of other models and **emits** physical changes back to them. It never re-models a sector, and it never presents a number produced against a stub as if it came from the counterpart model.

14.1 The rules every exchange follows

The design sets six rules, shared with the sibling health (HPRAS) and education (EPRAS) planning engines:

- **Files, not calls.** Engines exchange files or tables, never in-process function calls, because they run on different schedules and machines.
- **A common header.** Every exchange is versioned and dated and carries the same leading columns: country, run id, scenario id, schema version, producer and an exploratory flag. This is the `ExchangeHeader` dialect also used by TransportModelling.
- **A schema contract.** Every exchange has a contract file in `config/interfaces/`, because silent column changes are the commonest failure.
- **Flagged fallbacks.** A missing counterpart produces a documented, flagged fallback, never a silent zero.
- **Idempotence.** A re-run produces the same files.
- **Snapshot markers.** A `SNAPSHOT` file in an inbound folder marks a vendored copy as not live. A real export dropped in without the marker replaces it.

In the current build, the first two and the fourth rule are implemented. Every outbound file starts with the header columns and the schema version `tipras-1`, and every fallback is printed in the run log (the reference run lists eight). The third is not: `config/interfaces/` is empty, so no schema contract is yet checked on either side. The inbound folder `data/exchange/inbound/` is also empty. Inbound data reach TIPRAS today by being read from their source locations in the estate, not through versioned exchange files.

14.2 What comes in

Table 14.1 lists every inbound link the design defines and its state in the reference run `014c9cd7`.

Table 14.1: Inbound links and their state in the reference run

Counterpart	What TIPRAS needs	State in the reference run
SDI (spatial data infrastructure)	boundaries, population grid, noded road network, terrain, hazards	**Used.** The network was exported read-only from SDI's database into the geography snapshot `UGA-2026.09-c` (817,354 links). The `tipras_export` views the design asks SDI to create do not yet exist

Counterpart	What TIPRAS needs	State in the reference run
ATDGR (transport data repository)	canonical inventory: roads, structures, condition, traffic	**Partly used.** OSM features (ferries, fords, culverts, road numbers) were extracted from the estate's OpenStreetMap file; condition comes from agency statistics. No ATDGR export manifest is read, and no findings are written back
TransportModelling	origin–destination matrices, speeds, deterioration, alignments, mode capacities	**Not available.** TransportModelling writes no exchange files. Every item runs on its norms-tier fallback: gravity distribution for freight, generic speeds marked "NOT Ugandan", uncalibrated vehicle operating cost coefficients, no candidate alignments
UCSN (unit cost standards network)	unit prices with low, central and high values and an evidence tier	**Used, partly replaced.** The 7 UCSN transport coefficients are ILLUSTRATIVE (tier E7). The reference run prices works from 14 TIPRAS evidence codes (tier E5)
MTEFModelling	transport ceilings by fiscal year	**Used.** The final-round transport ceiling of the baseline scenario sets the envelope (Chapter 11)
FINEX (fiscal model)	the envelope by pool, with earmarks	**Not available.** No FINEX envelope reader exists in the estate. The pool split comes from approved budgets instead (Chapter 13)
HealthModellingx	facility register and planned facilities	**Used for nodes.** 3,169 geocoded facilities from the national register, as health nodes. The health engine keys facilities to 122 districts, not 135
EducationModellingx	school register	**Used for nodes.** 20,474 school nodes
AgricultureModelling	production surface by crop	**Not available as a surface.** The model publishes production for only four regions. TIPRAS builds 10 marketed-surplus zones of at least 5,000 t (9.76 Mt in total) itself, flagged as modelled
MineralModelling	mineral output and freight flows	**Not used.** The exchange is labelled "DEMONSTRATION DATA — synthetic". Mines enter as 309 nodes without magnitudes
TradeModelling	gateway flows and growth	**Not used.** The exchange is synthetic and for The Gambia (`geo_id` GMB), not Uganda. 28 gateways enter as nodes without throughput
MEL (monitoring)	km delivered, contract costs, execution	**Not available**

Two of these deserve comment.

The synthetic sector feeds are refused, not adapted. MineralModelling's `minerals_to_transport_flows.csv` (115 rows) declares in its manifest that it holds "synthetic values for testing and training only". TradeModelling's transport package is marked `synthetic: true` and describes another country. Loading either would give scenario D (productive sectors) freight magnitudes that look Ugandan but are not. The design allows a synthetic feed to be loaded only with `basis = assumed` and a demonstration label on every page that uses it; the reference run does not load them at all. The consequence is visible in the results: the run log records that the sector connection test finds 21 of 23 industrial nodes and 16 of 28 gateways lacking their required connection, but with "at stake 0.0", because no tonnage is attached to them. Only agriculture, with its own surplus zones, has a magnitude (9.76 Mt).

TransportModelling is the largest missing input. It owns speed calibration, deterioration, vehicle operating costs, traffic assignment and freight origin–destination matrices. None of these reaches TIPRAS as an exchange file. Its speed calibration was rejected (multipliers pinned at their bounds), and its deterioration ran on 96.9 % assumed

initial condition. TIPRAS does not re-implement these models; it runs on norms-tier fallbacks and says so. Adding exchange writers to TransportModelling, per the contract in the design, is a request to that project, not work for TIPRAS. Chapter 23 describes the small set of pure functions TIPRAS has ported from it.

14.3 What goes out

TI-10 writes the outbound files at the end of every run, under `data/exchange/outbound/`. The current set is dated 5 October 2026 and carries the run id `0bad432d`, a later run with the same geography, norms and prices as the reference run; its scenario B budget is identical to the reference run's. Every file of a UGA run carries `exploratory = true`.

Table 14.2: Outbound files written by TI-10

Folder	File	Content	Rows
`mtef/`	`transport_mtef_UGA_<date>.csv`	scenario B in three views (programme, economic, administrative) by code, fiscal year and pool	749
`finex/`	`transport_requirement_UGA_<date>.csv`	scenario B, economic view, by fiscal year, GFS class and pool	17
`healthmodellingx/`	`transport_access_UGA_<date>.csv`	population within 30, 60 and 120 minutes of a primary health facility, by district	135
`educationmodellingx/`	`transport_access_SCHOOL_PRIMARY_UGA_<date>.csv`	wet-season distance to the nearest primary school, per settlement cell, with whether it meets the standard	175,367
`educationmodellingx/`	`transport_access_SCHOOL_SECONDARY_UGA_<date>.csv`	wet-season travel time to the nearest secondary school, per settlement cell	175,367
`agriculture/`	`transport_cost_UGA_<date>.csv`	per surplus zone: surplus tonnes, crop mix, travel time to market on the required-standard network, whether connected at standard	10
`ucsn/`	`tipras_evidence_for_ucsn_<date>.csv`, `tipras_evidence_vs_ucsn_<date>.csv`	the 14 evidence prices in UCSN's format; the comparison with UCSN's own codes	14 and 7

14.3.1 The fiscal files

The MTEF and FINEX files are described in Chapter 13. They carry scenario B, in constant 2026 UGX. They lack several columns of the design contract for FINEX: a separate maintenance channel (road fund statutes ring-fence maintenance, so the finance ministry treats it as a separate line), the import content share (bitumen, steel, equipment and fuel are largely imported), and loan-disbursement and PPP availability columns for debt accounting. They are emitted without a schema contract, and neither counterpart has read them back. No file is written for CDCGE, the general-equilibrium model, although the design lists six channels for it: government investment, government consumption, construction demand, trade margins, productivity and market access.

14.3.2 The health file

The health file is the most mature sector exchange. It follows the schema HPRAS reads from its SDI health-accessibility snapshot, so that, dropped into HPRAS's inbound folder without a snapshot marker, it would replace that snapshot with no code change. In the current file, 52.01 M persons are covered across 135 districts; 79.7 % are within 30 minutes of a primary health facility, 94.0 % within 60 minutes and 96.3 % within 120 minutes, by mixed modes.

Two limits must be stated plainly. First, the values are the **base-year** network (year 2026), although the scenario id column says B. The code computes them from current access, not from the scenario's horizon network. HPRAS

cannot yet test its facility plan against the transport plan's future network, which is the purpose of the link. Second, the health engine keys its facilities to 122 districts, TIPRAS to 135, so a district crosswalk is needed on the receiving side.

14.3.3 The education files

The two education files give, for every one of the 175,367 settlement cells, the wet-season access to the nearest primary school (in km) and secondary school (in minutes), and whether the cell meets its standard. Counted over cells, not population, 48.3 % of cells meet the primary-school standard and 74.1 % the secondary-school standard. In 9,258 cells (primary) and 11,728 cells (secondary) no school is reachable, and the value is left blank rather than set to a large number. Like the health file, these are base-year values.

14.3.4 The agriculture file

The agriculture file lists the 10 surplus zones, which together market 9.76 Mt. One zone, in district UG2039, holds 8.47 Mt of it. The travel time from each zone to its nearest market on the required-standard network ranges from 2.5 to 114.9 minutes (median 54.0). None of the 10 zones is connected at its required standard today. The design asks for a farm-gate transport margin per tonne, which AgricultureModelling would use for its price and post-harvest-loss channels; the file carries travel time instead, because the freight cost model runs on uncalibrated coefficients. The yield and price response remains AgricultureModelling's to compute.

14.3.5 Files not yet written

No files are written for MineralModelling, IndustryModelling or TradeModelling (route capacity, logistics cost, gateway cost), for MEL (physical targets by district and year), for TransportModelling (the funded programme for detailed appraisal), for ATDGR (data findings) or for SDI (published layers). Each is designed; none is built.

14.4 The unit-cost link: evidence against UCSN

The UCSN link runs both ways. Inbound, the UCSN file `illustrative_coefficients_UGA.csv` supplies seven transport coefficients, all ILLUSTRATIVE (tier E7). Outbound, TIPRAS sends UCSN its own evidence prices in UCSN's format (14 codes, with the source basis of each) and a comparison of each evidence price with the UCSN code it would replace.

Table 14.3: TIPRAS evidence prices against the UCSN illustrative coefficients (UGX 2026, central values)

UCSN code	UCSN central	Evidence code	Evidence central	Evidence ÷ UCSN
TRN-OPS-001-ROUTINE-GRAVEL (per km a year)	10.0 M	TRN-OPS-003-RM-UNPAVED-EV	4.36 M	0.44
TRN-OPS-001-ROUTINE-PAVED (per km a year)	18.0 M	TRN-OPS-004-RM-PAVED-EV	16.37 M	0.91
TRN-CAP-001-GRAVEL (per km)	500 M	TRN-CAP-008-GRAVEL-EV	75.1 M	0.15
TRN-CAP-001-PAVED2L (per km, flat)	3,100 M	TRN-CAP-009-PAVED-EV	4,269 M	1.38
TRN-CAP-001-PAVEDROLL (per km, rolling)	3,700 M	TRN-CAP-009-PAVED-EV	4,269 M	1.15
TRN-CAP-001-PAVEDMTN (per km, mountainous)	5,200 M	TRN-CAP-009-PAVED-EV	4,269 M	0.82
TRN-CAP-004-BRIDGE (per m of span)	180 M	TRN-CAP-004-BRIDGE-EV	86.2 M	0.48

The comparison matters well beyond transport:

- **Gravel construction is the largest difference.** The evidence price, 75.1 M UGX per km, is 15 % of UCSN's 500 M. It rests on National Development Plan III actuals, two rural roads projects and 12 community access road contracts. Gravel works dominate the Uganda programme by km, so this one coefficient changes the whole cost picture.

- **Routine maintenance of unpaved roads** is 44 % of the UCSN value. That difference alone roughly halved the national routine maintenance requirement, from 2,049.6 bn to 954.7 bn a year between the draft-e and draft-f norm sets (Chapter 13).
- **Paved upgrading** is 15 to 38 % above UCSN on flat and rolling terrain and 18 % below on mountainous terrain. No source splits paved upgrading by terrain, so a single evidence code covers all three.
- **Bridges** are 48 % of UCSN, but the evidence is one Bailey bridge. It is not a price to rely on.

The design says outbound quantity weights should also tell UCSN which coefficients matter most. In Uganda the answer from this run is clear: gravel and engineered-earth upgrading, and routine maintenance of unpaved roads. These files have been written but there is no record that UCSN has reviewed them. The evidence prices remain tier E5 candidates until it does.

14.5 What integration means for the results

Every result in Chapters 4 to 13 inherits the state of these links:

- the envelope rests on MTEFModelling's baseline ceiling and approved budgets, not on a FINEX envelope with earmarks;
- the freight, speed and economic appraisal rest on norms-tier fallbacks because TransportModelling sends nothing;
- scenario D (productive sectors) sees agricultural surplus but no mineral, industrial or trade tonnage, because the feeds that exist are synthetic;
- prices rest on TIPRAS's own evidence, not on a reviewed UCSN vintage;
- the health and education models receive base-year access, not the access each scenario would deliver.

None of this is hidden: the run log and the run panel of the interface list the fallbacks in use, and this report repeats them.

14.6 Implications for planning

- **The fiscal link is the most important one to close.** A FINEX envelope reader with pools and earmarks, and a FINEX return of debt service and availability payments, would replace the assumed envelope and complete the fiscal sustainability statement.
- **TransportModelling's exchange writers would lift the economic results from fallback to model.** Speeds, vehicle operating costs and freight matrices are requests to that project; their absence is why NPV-based comparisons remain weak.
- **The health and education outbound files should carry scenario outcomes, not base-year access.** The change is small in code and is what makes joint health–transport and education–transport planning possible.
- **Real sector magnitudes are needed for the productive-sector scenario.** Until MineralModelling and TradeModelling publish Uganda data, scenario D is an agriculture scenario in practice.
- **Schema contracts should be written before more files are exchanged.** The `config/interfaces/` folder is empty, and silent column changes are the main failure mode the design warns against.
- **UCSN should review the evidence prices**, starting with gravel construction and unpaved routine maintenance, which drive most of the programme's cost and maintenance burden.

CHAPTER 15

15 The planning interface

A planning engine is only useful if planners can see what it says, question it and run it again with their own choices. This chapter describes the web interface built for TIPRAS: what it shows, how it is organised around the five planning questions, how it enforces the engine's rules on screen, and what is still missing. It implements design chapter 14 (`docs/14_INTERFACE.md`). The interface was rebuilt on 24 September 2026 on the pattern of the energy planning engine (`../EnergyModellingx`), kept on the Genie web framework as the design requires. It is served by `app.jl` at `http://127.0.0.1:8050` and started with `run.ps1 web` from the repository root. Everything it shows is read from a finished run: the interface recomputes nothing.

15.1 The principle: five questions, every screen a map and a table

The design fixes one principle. The interface is organised around the five planning questions, permanently. It does not open on a model, a network editor or a map tool; it opens on the questions.

- **Q1 Needs:** who must be connected, and to what standard?
- **Q2 Have:** what network exists, and what access does it give today?
- **Q3 Gaps:** where are the gaps, one dimension at a time?
- **Q4 Cost:** what would closing them cost, with maintenance?
- **Q5 Where:** how should a fixed envelope be allocated?

Every screen pairs a map with a table of the same thing. Selecting a table row zooms to its district; clicking a district highlights its row.

Figure 15.1 shows the overview page of the reference run.

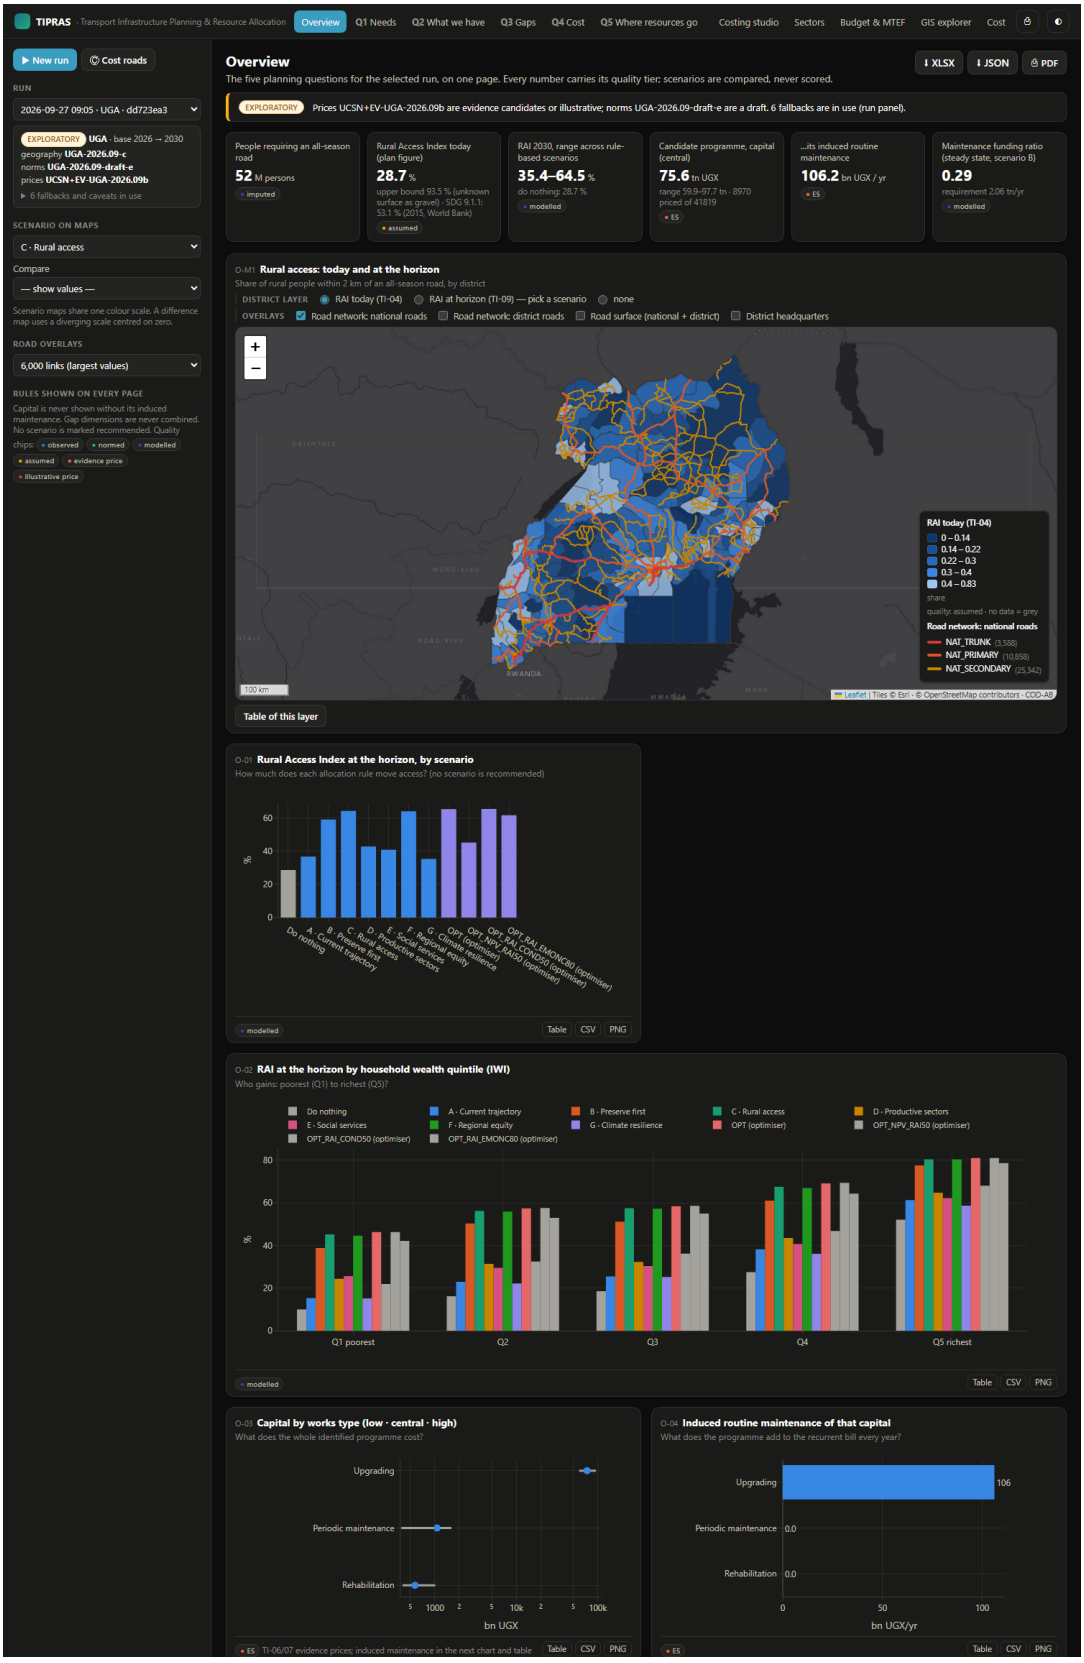

Figure 15.1: The planning interface: the Overview page (KPIs with quality chips, access map, scenario charts)

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

The overview gathers the headline figures as tiles, each with a quality chip that says how far to trust it (observed, normed, modelled, assumed, or a price tier such as E5). In the reference run they include the 52.0 M people requiring an all-season road, the base-year Rural Access Index (28.7 %, chip *assumed* because it depends on the surface convention), the range of RAI at the horizon across the rule-based scenarios, and the candidate

programme's capital at central prices (92.8 tn UGX) **with its induced routine maintenance beside it** (205.2 bn UGX a year). A side panel states the run: its exploratory status, the fallbacks in use (eight in the reference run), and the norm version, price vintage and geography it was built from.

15.2 The pages

The interface has twelve pages, each reachable by a deep link.

Table 15.1: Pages of the planning interface

Link	Page	What it shows
`#overview`	Overview	headline tiles with quality chips; validation checks; run panel
`#needs`	Q1 Needs	population requiring each destination class; required standards and the binding test; sector nodes
`#have`	Q2 Have	surface, wet-season closures, RAI by district, access to each class, condition by network
`#gaps`	Q3 Gaps	the seven gap dimensions, one at a time; the access cause decomposition; the district gap matrix
`#cost`	Q4 Cost	capital by works type (low, central, high) beside induced maintenance; backlog; candidate programme
`#where`	Q5 Where	all scenarios side by side; optimiser runs with objectives and floors; frontier; fiscal statement
`#sectors`	Sectors	what each scenario buys for health, education, agriculture, minerals, industry and trade
`#budget`	Budget	scenario B in three reconciling views; the fiscal sustainability statement
`#gis`	GIS explorer	every district layer; Moran's I; LISA clusters; histogram; Moran scatter
`#evidence`	Evidence	the evidence prices, their rows and sources
`#explain`	Explain	the factor decomposition of any cell
`#data`	Data	the output tables of the run, for download

Each page is a payload built by `web/PageData.jl` from the run's stored layers and tables, and drawn by one generic renderer, `web/static/app.js`. The design system (palette, Plotly chart factories, map and tile specifications) is in `web/Pages.jl`. Charts come with buttons to show the underlying table, download it as CSV or save the chart as PNG. Every page can be exported to XLSX (one sheet per chart table) or JSON, and printed to PDF. The manual itself can be read inside the app. The interface has light and dark themes.

15.2.1 Q5: where resources go

Figure 15.2 shows the Q5 page, which carries the material of Chapters 11 and 12.

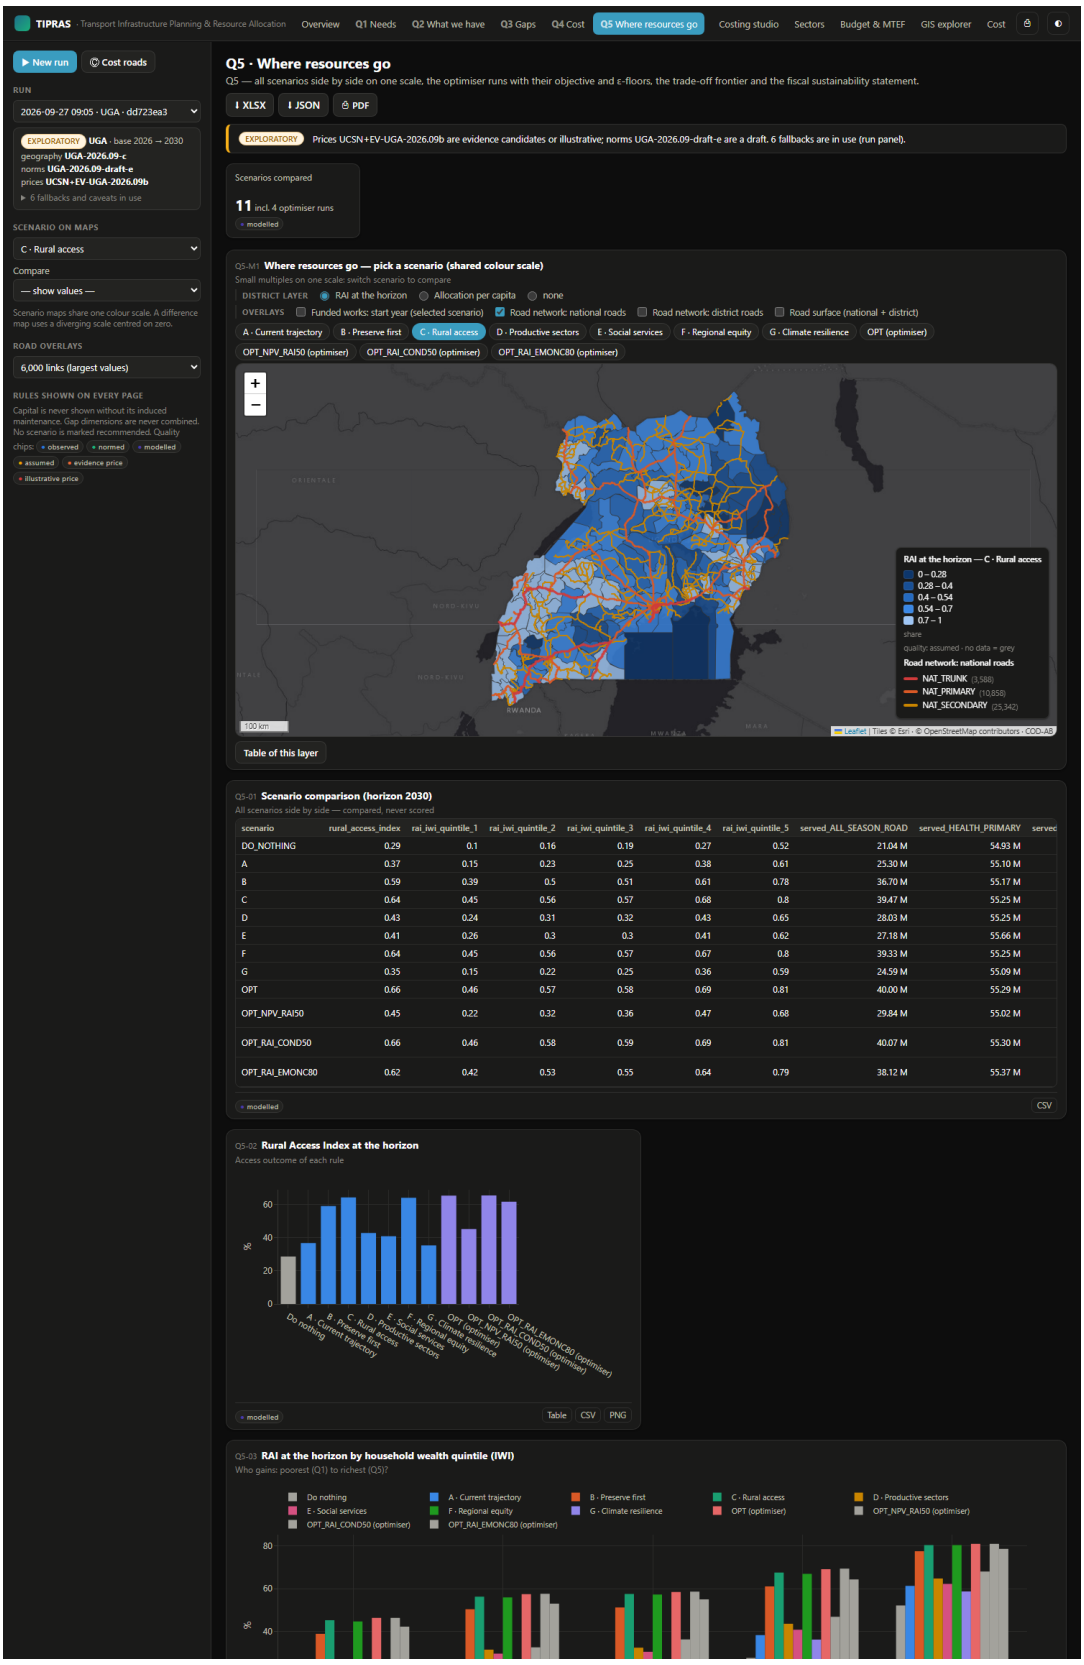

Figure 15.2: Q5 Where resources go: scenarios side by side on one scale

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

The Q5 page holds:

- a tile counting the scenarios compared (eleven in the reference run, including four optimiser runs), with the note "none is marked recommended";
- the side-by-side scenario comparison table, with no ranking or score column;

- the RAI at the horizon by scenario, with optimiser runs drawn in a separate colour from the rule-based scenarios;
- the equity chart by wealth quintile;
- persons newly meeting each access standard, class by class, in persons, never summed;
- the interaction adjustment of each scenario;
- capital by pool, by works type and by year;
- the trade-off frontier between rural access and poor km treated. Points where the solver hit its time limit and the floor was not met are **left off the curve** and listed in a note below it; in the reference run that is the 100 % point;
- the optimiser runs table, with the note that a positive floor shortfall is a finding, not an error;
- the fiscal sustainability statement for scenario B.

Its map lets the user pick one scenario for all layers: RAI at the horizon, allocation per capita, and the funded works coloured by start year. All scenario maps share one colour scale, and a **difference-from** mode redraws any scenario as its difference from another on a diverging scale, the interactive form of the C − B map in Chapter 11.

15.2.2 The GIS explorer

Figure 15.3 shows the GIS explorer.

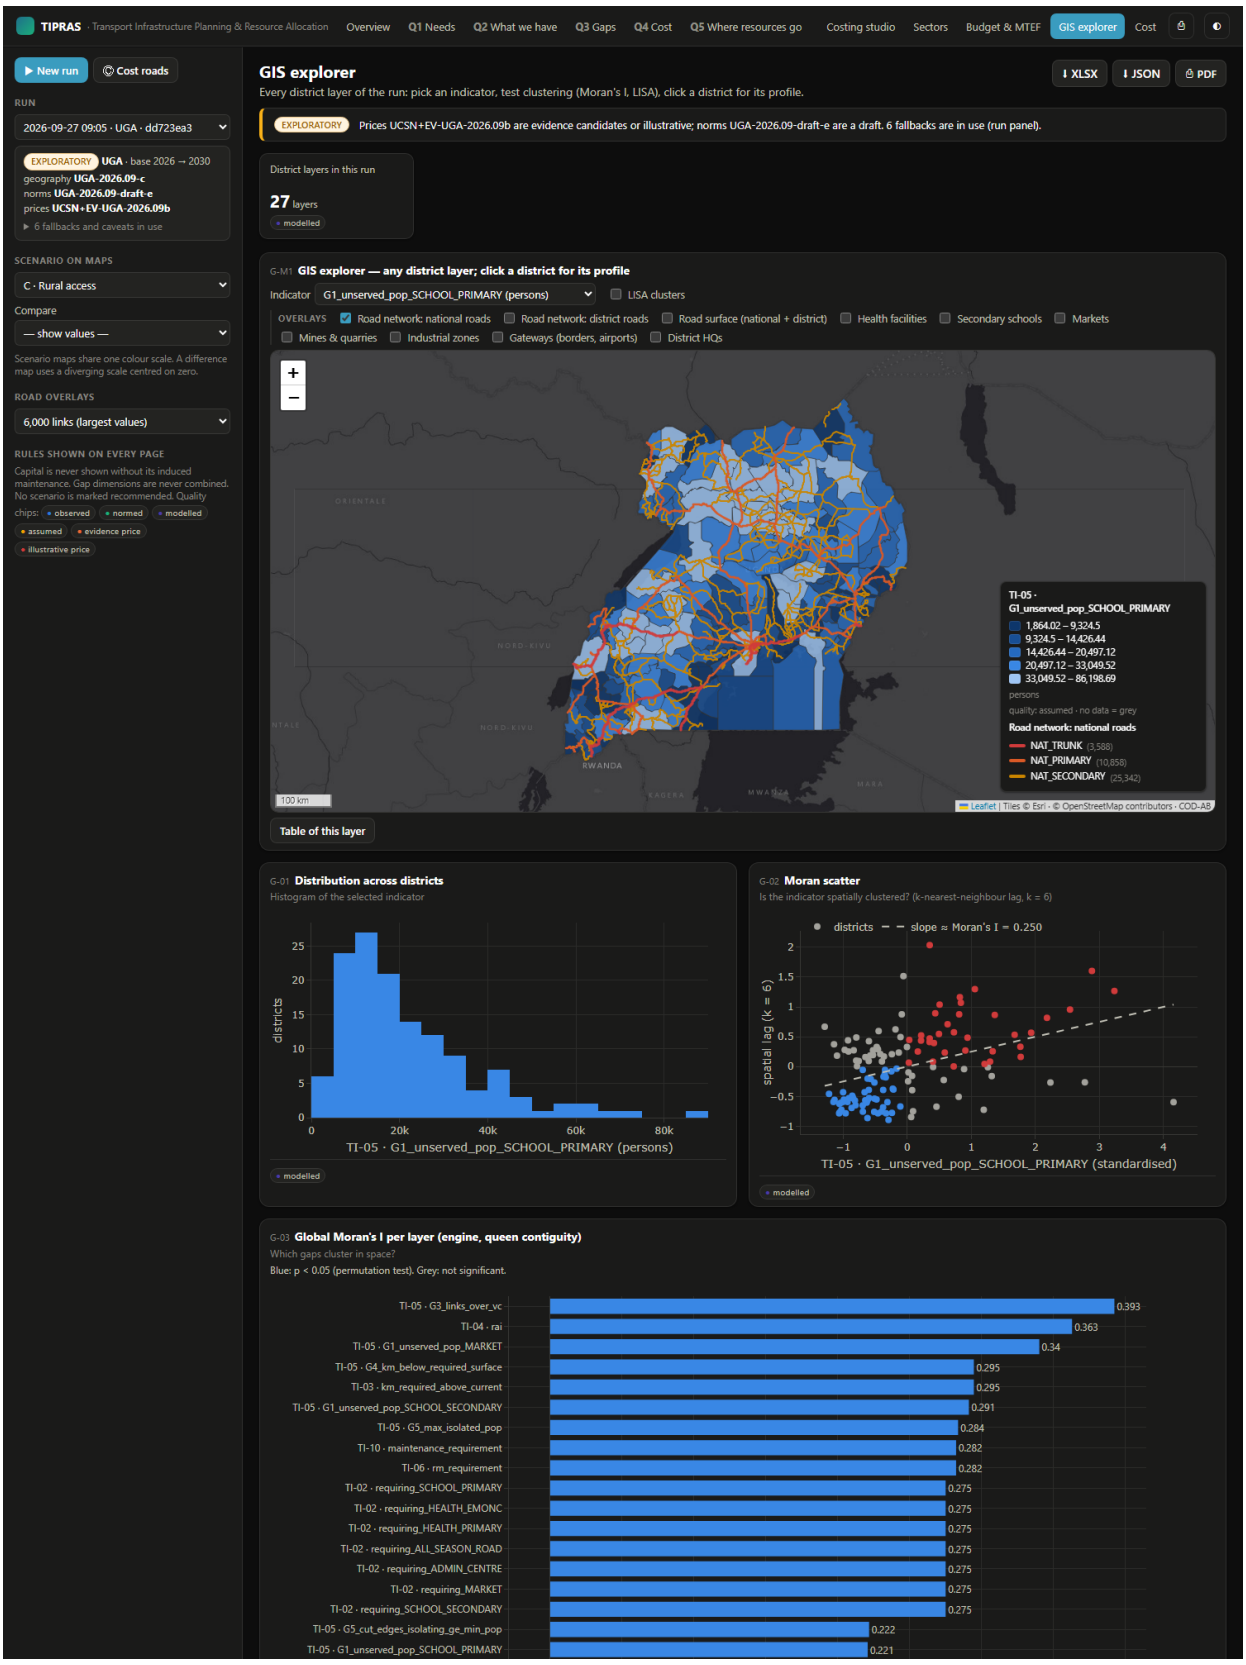

Figure 15.3: The GIS explorer: any district layer, histogram, Moran scatter and Moran's I per layer

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

The explorer holds every district layer the engine writes: each gap dimension, each access class, RAI, allocations and budgets. For the selected layer it shows the choropleth, a histogram, a Moran scatter (each district's value against the average of its nearest neighbours), the engine's Moran's I and the engine's LISA clusters (high–high,

low–low and the outliers). Chapter 21 reports these statistics; of the 25 layers summarised there, 23 show significant spatial clustering, the strongest being links over capacity (Moran's I 0.39) and RAI (0.36).

Only one district layer is displayed at a time, and no layer blends two gap dimensions into one colour, because the seven gap dimensions are never combined.

15.2.3 Q4: cost

Figure 15.4 shows the Q4 page.

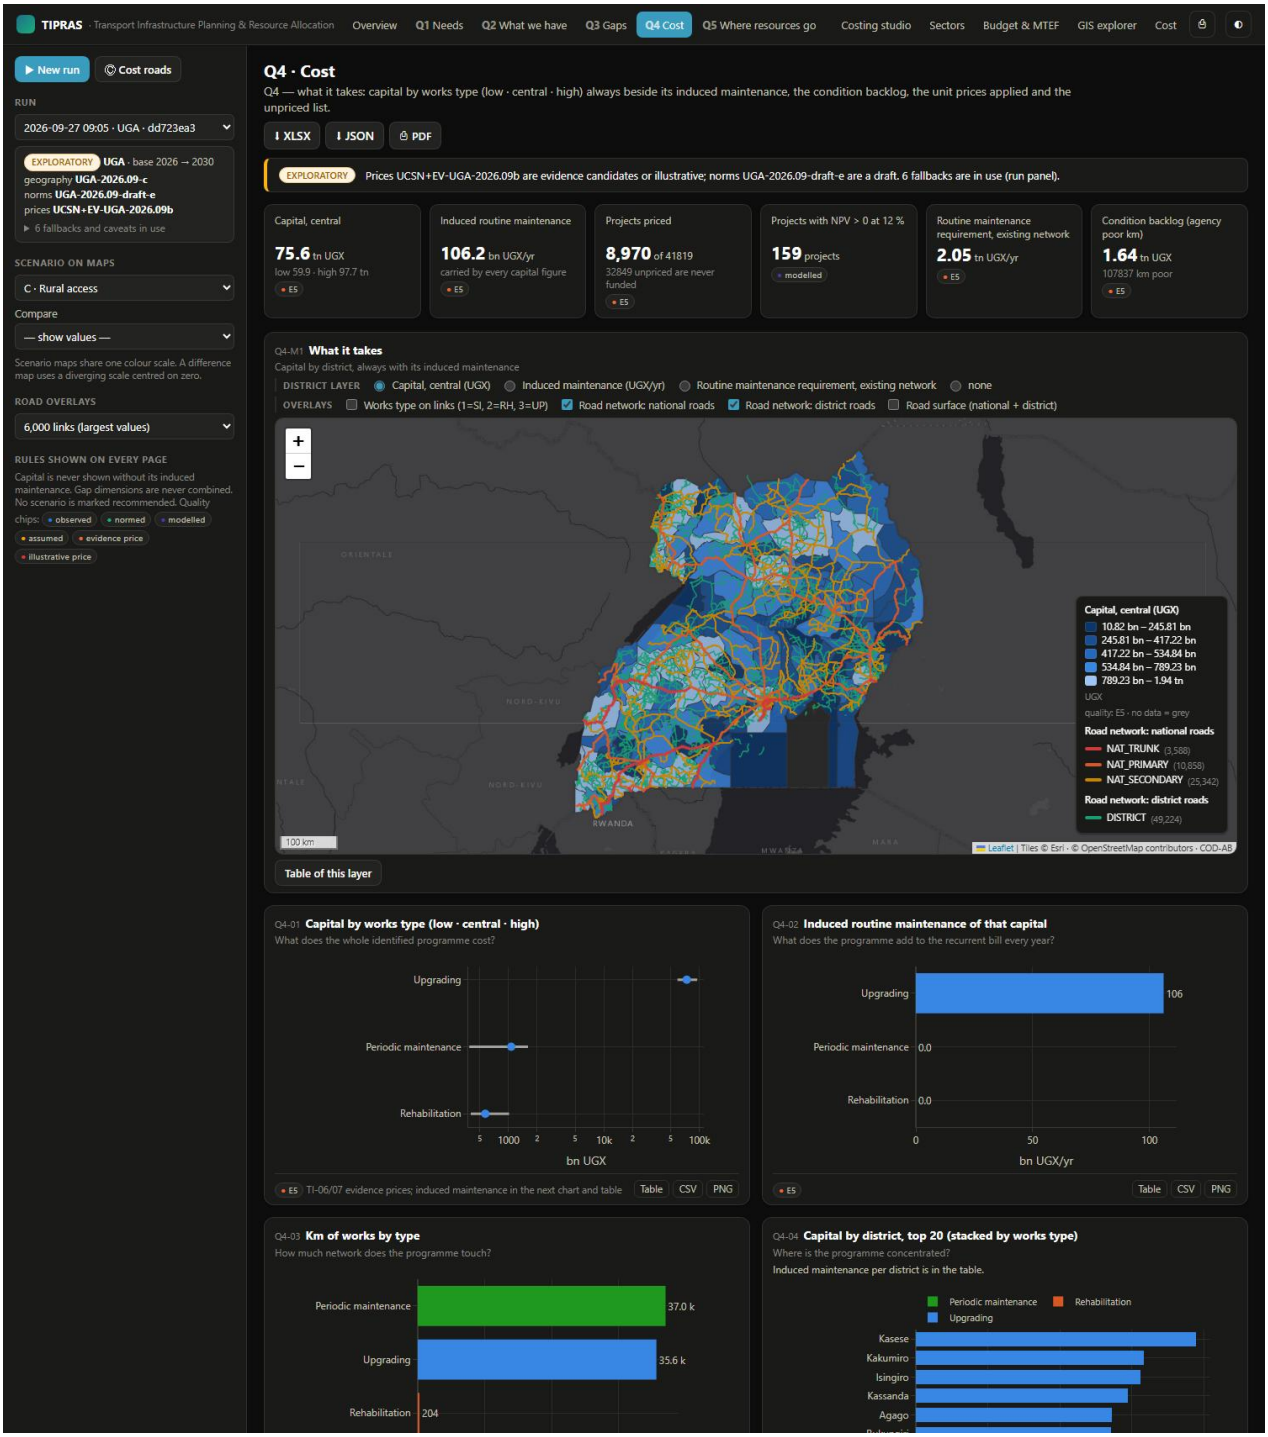

Figure 15.4: Q4 Cost: capital by district always beside its induced maintenance

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

The cost page shows the candidate programme's capital by works type as a low · central · high range plot, and beside it the induced routine maintenance each works type commits the budget to. It also shows the maintenance

requirement, the condition backlog by network, the evidence prices applied and the unpriced items. The rule that no capital figure appears without its maintenance is enforced by the page layout: the capital and maintenance charts are built from the same table and placed together.

15.3 The costing studio

Figure 15.5 shows the costing studio (`web/Costing.jl`), added on 24 September 2026.

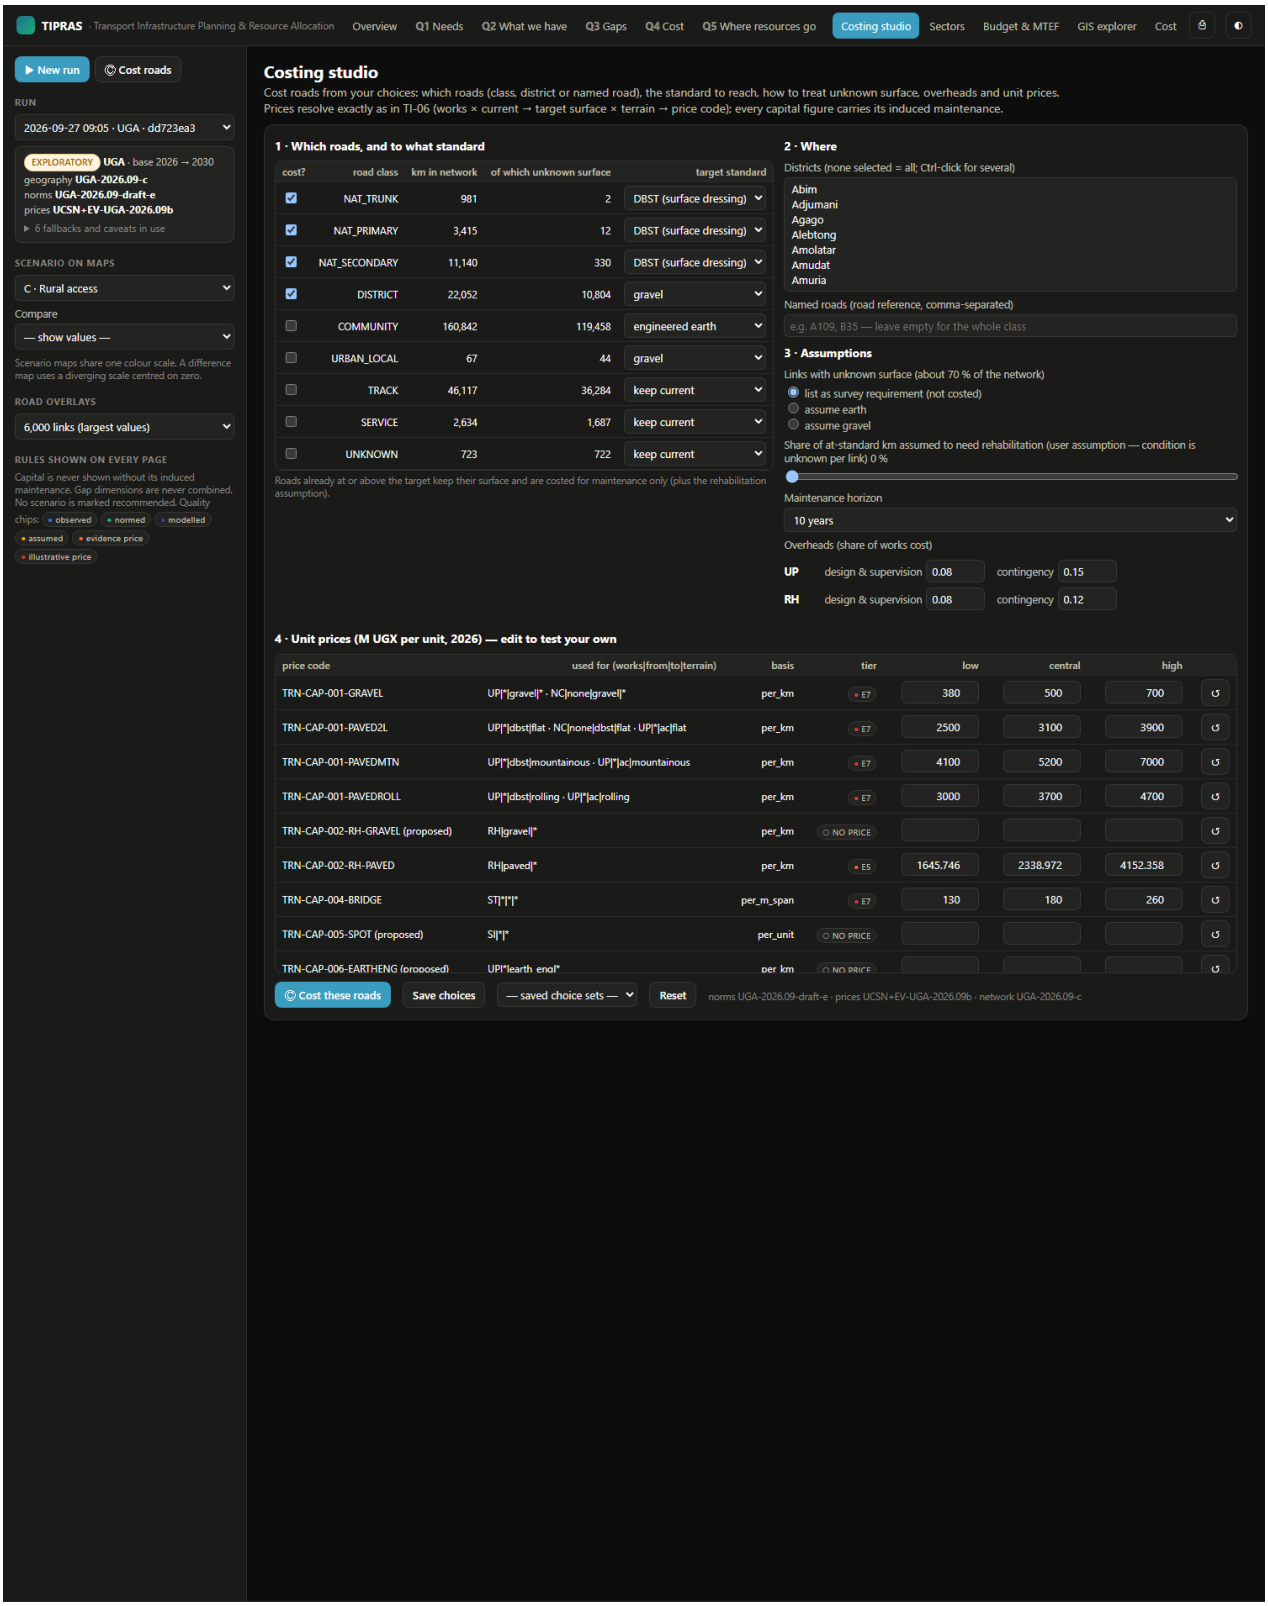

Figure 15.5: The costing studio: cost roads from choices of scope, standard, assumptions and unit prices

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

The studio lets a planner cost a set of roads from their own choices, without running the whole engine:

- **Scope:** road classes, districts, and named roads by their OpenStreetMap reference. A named road is costed along its whole length.
- **Target standard per class:** keep, engineered earth, gravel, low-volume sealed road (LVSR), double bituminous surface treatment (DBST) or asphalt concrete (AC).
- **Unknown surface:** listed as a survey requirement (the default), or assumed earth, or assumed gravel. Because surface is unknown on about 68 % of network km, this choice changes the answer more than any other.
- **Rehabilitation assumption:** a user share of at-standard km needing rehabilitation. It is labelled a user assumption, because condition is unknown per link.
- **Maintenance horizon, overheads and unit prices.** Prices can be edited; a changed price is flagged as a user override.

The studio resolves prices exactly as TI-06 does: works type, current and target surface and terrain give a unit-cost key, the run's norm set maps the key to a price code, and the run's price vintage gives the price. Induced routine maintenance follows the same equation as the engine, and periodic maintenance is annualised over the norm cycle.

Its output is a page of its own: KPI tiles, capital by class (low · central · high) beside its induced maintenance, the transitions costed, capital by district as a chart and a map, the cumulative cost over the chosen number of years, the unit prices applied, the full detail table and the survey list. Choice sets can be saved in the browser.

The studio states one known limit. The unit-cost map has no paved-to-paved code, so upgrading DBST to AC is costed at the full upgrading price. The studio says so and suggests "keep" or a price override. Its defaults keep trunk roads at DBST.

15.4 The road network on every map

Figure 15.6 shows the road network drawn over a district layer (`web/Network.jl`).

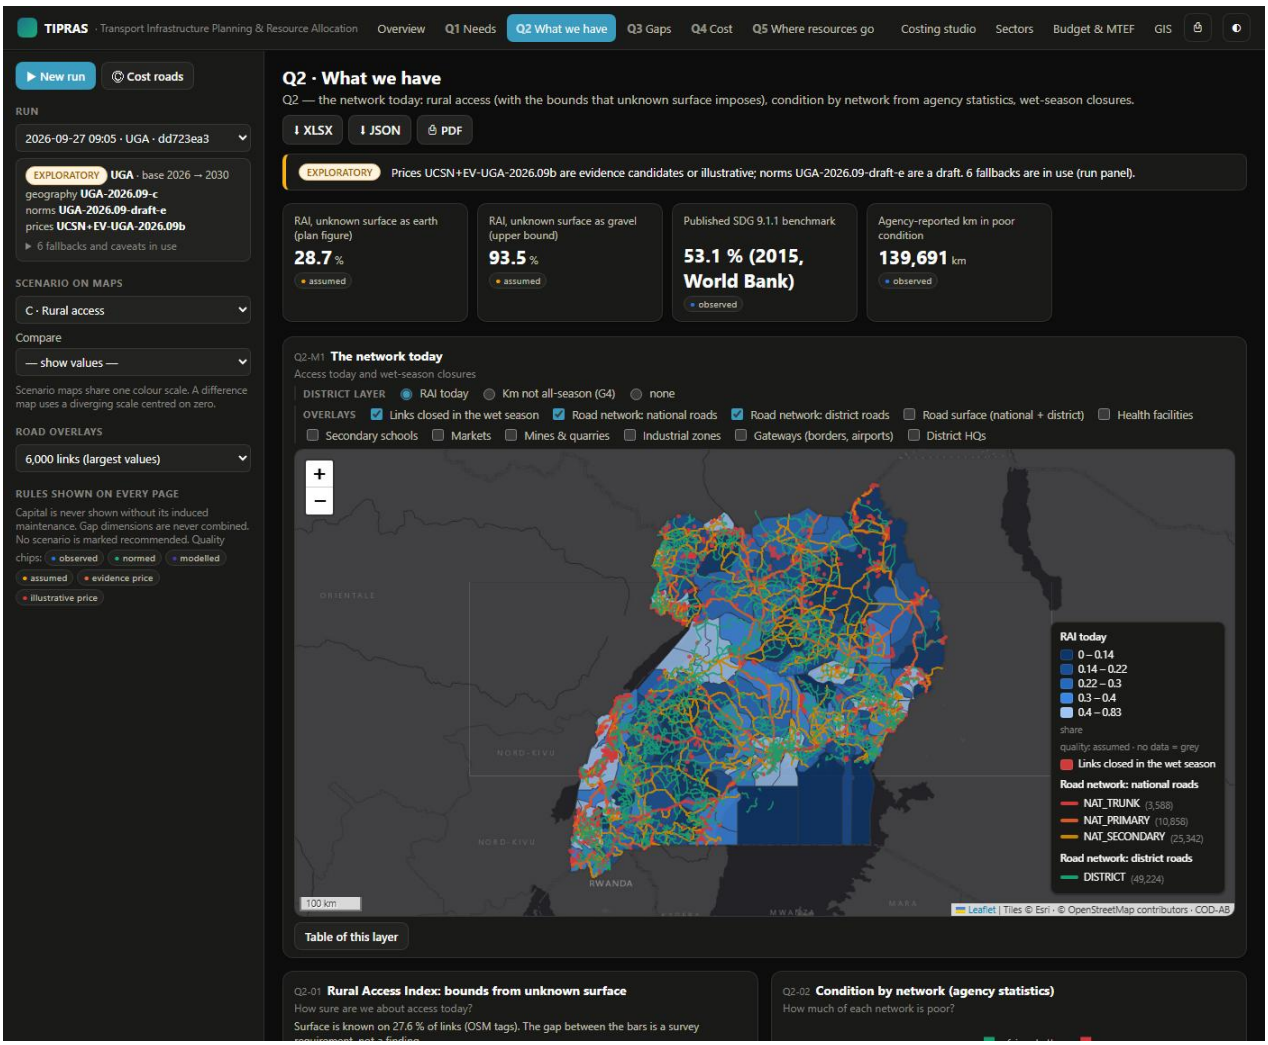

Figure 15.6: Q2 What we have: rural access with the road network and wet-season closures

Source: TIPRAS web interface (app.jl), http://127.0.0.1:8050

Every page map draws the road network in its own layer above the district colours:

- national roads by class, on by default;
- district roads, on by default for Q2 and Q4;
- a surface view (paved, gravel, earth), with **unknown surface drawn dashed**, so the user sees at once where the network data are missing.

In the costing studio, the costed roads are coloured by what the user's choices do to each link: upgrade to a target, maintain, or survey (dashed).

The network is served as grouped GeoJSON: links are combined into one multi-line per category and thinned to points at least 120 m apart, results are cached per geography, and the national layer is prepared when the server starts. At most 160,000 links are drawn nationally, and the legend says which classes were left out; the user picks districts to see them. The full network has 817,354 links, so the national view necessarily shows the main network and leaves community roads and tracks to district views.

15.5 Running the engine from the interface

A "New run" button in the side panel opens a dialog in which the user chooses:

- the geography snapshot, norm version and price vintage, from the versions present in the store and the configuration folders;

- the unknown-surface policy;
- whether to run the optimiser and frontier. Without them a Uganda run takes about 12 minutes; with them it takes 20 to 75 minutes more (TI-08 took 76 minutes in the reference run);
- whether to prune the store to the latest N runs first.

The run starts as a separate process (`scripts/run_job.jl`), so the interface stays usable. The dialog and side panel show progress (modules done out of ten, the current module, elapsed time) and the tail of the log. A run can be cancelled, and when it finishes the new run can be opened at once. Only one run may be active at a time, and a run is refused when less than 4 GB of disk is free unless pruning is chosen, because each national run adds about 2 GB and the data disk is nearly full.

Runs are started outside the server's own process tree and tracked by process id, so they survive a server restart. This was learnt the hard way: on 24 September 2026 a server restart killed a run at TI-05, while runs were still children of the server process, and it had to be relaunched.

15.6 District transport profiles

Clicking a district on any map opens its **transport profile**:

- each gap dimension as its own tile, in its own unit;
- the RAI under every scenario;
- the district's share of the candidate programme, with capital and induced maintenance;
- the allocation the district receives under each scenario.

The profile is the interactive form of the design's District Transport Profile report. It is what a district engineer or chief administrative officer would look at first.

15.7 The rules enforced on screen

The design sets hard rules for the interface. Table 15.2 shows how each is met in the current build.

Table 15.2: The interface's hard rules and how they are met

Rule (docs/14 §14.3)	How it is met
The five questions are the navigation	Q1–Q5 are the core pages of the navigation, after a short overview page
No capital cost without its induced maintenance	capital tiles and charts are always paired with induced maintenance from the same table
Never combine gap dimensions	one district layer at a time; no composite layer exists
Never mark a scenario recommended	no ranking column; optimiser runs are labelled by objective and floors; the P3 check
Every figure shows its quality tier	quality chips on every tile
Modelled values are visually distinct	unknown surface drawn dashed; survey links dashed in the studio
The run banner states exploratory status and fallbacks	the run panel
Every number opens its explanation	partly: the Explain page decomposes a cell, but not every number on every page is clickable yet
The equity exponent is a live slider	**not met**: re-solving scenario F needs the project and coverage objects stored per run

The overview, Q5 and budget pages also show the validation checks the engine passes or fails on the loaded run (Chapter 16). On the reference run they show the optimiser floor checks failing for OPT_RAI_EMONC80 and the frontier's 100 % point, as findings.

15.8 What is still missing

- **The live equity slider.** Re-solving scenario F as the exponent α moves needs the TI-07 project and coverage objects stored per run; today they exist only during a run.
- **The year slider.** TI-09 stores the horizon year only, so there is no network path to step through.
- **The narrative reports** of the design (§14.4), such as the Transport Sector Investment Case, the Road Maintenance Needs Statement and the Programme Budget Submission. This report is produced outside the app.
- **Authentication.** Results routes are read-only; the routes that start or cancel a run and that cost a road set accept requests from anyone who can reach the server, and there is no user model. Norm editing through the interface, which the design reserves for authorised roles, does not exist.
- **The design's versioned API** (`/api/v1/runs/...`). The implemented endpoints serve the same purposes under different paths (`/api/page`, `/api/runs`, `/api/run/<id>`, `/api/table`, `/api/explain`, `/api/network`, `/api/district` and others).

A smoke test, `scripts/check_web_pages.jl`, builds every page payload without the server, so a broken page is caught without opening a browser. The first build of a run's pages takes about 45 seconds; the server builds the latest run's pages in the background when it starts.

15.9 Implications for planning

- **The interface is usable today for exploring a run** with ministry staff: every number in this report can be found on a page, with its quality tier and its table.
- **The costing studio is the quickest route to a planner's own figures.** It answers "what would it cost to bring these roads to this standard?" in seconds, with maintenance, and lists the survey needed where surface is unknown.
- **The unknown-surface choice should always be shown to users.** It moves costs and access more than any price, and the studio and the dashed links make it visible.
- **Before ministry-wide use, authentication and the equity and year sliders are needed.** The sliders are what turn the interface from a viewer into a negotiating tool.
- **The narrative reports should be generated from the same payloads** so that a submission and the screen cannot disagree.

CHAPTER 16

16 Calibration and validation

A planning engine earns trust in two ways. It must be **right by construction**: its accounts must add up, its rules must hold, and it must refuse to produce a number when an input is missing. And it must be **right about Uganda**: its travel times, traffic, access and maintenance costs must match what is observed. The first is validation; the second is calibration. This chapter reports both for the Uganda build. It implements design chapter 15 (`docs/15_CALIBRATION_AND_VALIDATION.md`). The short answer is that the engine is well validated and barely calibrated. Its 2,265 test assertions pass, and its run-time checks hold on the reference run, but none of speeds, vehicle operating costs, traffic or link condition has yet been compared with Ugandan observations. That is the main reason every result in this report is labelled exploratory.

16.1 Pre-run gates

Gates stop a run rather than let it produce a misleading number. The design defines seven, plus six GIS gates.

Table 16.1: Pre-run gates and their behaviour in the reference run

Gate	Aborts when	Reference run
G1 Norms	a needed norm cannot be resolved, or the norms are not approved and the run is not exploratory	passed as exploratory: norms `UGA-2026.09-draft-f` are a draft
G2 Prices	a works item with a requirement has no price, or an illustrative price is used in a non-exploratory run	passed as exploratory: 23 projects are unpriced and are listed, never funded; prices are tier E5 candidates
G3 Geography	the geography version differs between population, nodes, network and boundaries	no abort recorded: one snapshot, `UGA-2026.09-c`
G4 Population	the projection is missing for any district, class and year	no abort recorded
G5 Standards	a target refers to a class without an access standard	no abort recorded
G6 Envelope	the envelope is missing or its pools do not cover the works	passed: the MTEF ceiling file exists for every plan year
G7 Norm approval	norms not approved and the run not flagged exploratory	passed as exploratory
GIS-5 Unknown surface	more than 30 % of network km has unknown surface	**raised as a fallback**: 68.2 % of km is unknown, so every access figure is exploratory

Two principles are visible in the code. A missing norm raises `NormNotFound`; it is never replaced by a default. An unknown surface is treated as earth for access and listed for survey; it is never treated as paved. The design states why: a silently costed zero or an assumed paved road produces a result that looks plausible and is wrong.

16.2 The test suite

The suite is `test/runtests.jl`. On its last recorded run it passed **2,265 of 2,265 assertions** in 3 minutes 12 seconds. It is organised in the layers of the design.

Table 16.2: The test suite by layer

Layer	What is tested
Unit	the standards lattice (the join of two standards covers both; unknown surface ranks as earth); the equivalent standard axle example (an equivalence factor of 4.80 and a design load of 1.89 M axles); a 2 × 2 Furness balance (rows and columns to within 0.1 %, under 100 iterations); the Kingman queueing delay (85.0 time units at 85 % utilisation in the worked example); the Gini coefficient, the concentration index and Moran's I on hand-computed cases
GIS	shortest paths against a brute-force answer on a small graph (24 minutes from A to C via B, not 30 direct); one-way links; snapping a point to a link (300 m, at mid-link); cut edges (only the pendant link is a bridge); refusal of coordinates in degrees
Norms	the most specific scope wins; `NormNotFound` for a missing norm; approved norms cannot be changed; the audit trail cannot be deleted or edited
Golden country	the synthetic country `TST` run end to end, with the identities and rules below
Property	more budget never reduces the optimum, on a small knapsack

16.2.1 The golden country

`TST` is a synthetic country built by `test/fixtures/build_tst.jl`: a grid of 30 by 40 vertices at 2 km spacing, divided into 12 districts. It is deliberately awkward. A river runs through it with one permanent bridge, which is a cut edge in the wet season, and a seasonal drift that closes in the wet season. It has a flood-exposed zone, a mine whose access bridge is rated at 10 t against the mine's 11.5 t axle standard, and the other features listed in the design. It is the only synthetic data in the repository.

The golden test runs the whole pipeline, TI-01 to TI-10, on `TST` and checks:

- **Identity 1:** raked population equals the projection for every district, urban/rural class and year.
- **Identity 2:** served plus unserved equals the requiring population, for every destination class:

$$S_k + U_k = N_k \quad \forall k \tag{16.1}$$

- every link with demand has a required standard at least its function's minimum, and a recorded binding test;
- **no composite column** in the gap matrix: no column name containing "overall", "composite", "index" or "score";
- **Identity 5:** low ≤ central ≤ high cost for every priced project;
- **Identity 6:** every allocation spends no more than each pool's envelope in each year, and every funded item is eligible for its pool:

$$\sum_{j} f_{j,\pi,t} \le B_{\pi,t} \quad \forall \pi, t \tag{16.2}$$

- **unpriced projects are never funded**;
- **the optimiser matches or beats every rule-based scenario** on its own objective (persons served by an all-season road);
- **Identity 7:** the zero allocation (doing nothing) is no better than any funded allocation;
- the golden facts: the bridge isolates more than 100,000 persons, and the mine is not connected at its standard;
- the programme budget for scenario B is written;
- the frontier is non-increasing as the floor tightens, within a 2 % allowance for the MIP gap;
- `explain()` decomposes a population cell into its two factors (the population grid and the raking factor).

Many assertions sit inside loops (identity 1 runs over every district, class and year; identity 6 over every scenario, pool and year and every funded row), which is why the count is in the thousands.

16.2.2 What the suite does not yet test

- **Identities 3, 8 and 9** of the design (freight conservation, district outcomes summing to national, budget views reconciling) are not in the suite. Identity 9 is asserted at run time by TI-10 instead, and identity 10 (every capital row has induced maintenance) is checked by the interface (below).
- **Five of the six property tests** are not implemented: looser access thresholds, added links, deteriorated links, tighter floors and higher unit costs. The one property test present checks budget monotonicity on a toy knapsack, not on the engine's own MILP.
- **The golden outputs are not frozen.** The design asks for a frozen expected file for every result table, with any change reviewed. The suite checks rules and identities, not values.
- **The TST maintenance funding ratio is above 1.** The design requires the golden country to have a ratio below 1 and a binding donor earmark. The last test run reports a ratio of 1.18, and earmarks are not modelled, so those two awkward cases are not exercised.
- **Reproducibility** (identical pins give identical results; a pinned historical run reproduces its numbers) is not tested in CI.

16.3 Validation checks on the Uganda run

The interface runs a set of checks on every loaded run and shows them on the overview, Q5 and budget pages (Chapter 15). On the reference run:

Table 16.3: Validation checks on the reference run

Check	Rule	Result on run `014c9cd7`
P3	no composite score and no "recommended" column in the scenario comparison	passed
Unpriced never funded	every funded project is priced	passed: 0 unpriced projects funded
Induced maintenance	every priced capital item carries induced maintenance	passed: 205.2 bn UGX a year over the programme
No scenario below doing nothing	every scenario's RAI ≥ the do-nothing RAI	passed: the lowest is A at 37.4 %, against 28.7 %
OPT ≥ rules	OPT serves at least as many as any rule-based scenario	passed: 43.34 M against 43.34 M (C), equal
Optimiser floors	every floor met within 0.5 %	**failed for OPT_RAI_EMONC80**: EmONC floor short by 92.3 % (solver time limit)
Budget views reconcile	programme, economic and administrative totals equal	passed: 20,894.9 bn UGX in each

The floor failure is shown as a finding, not hidden. Its cause is known (Chapter 11): the solver reached its 900-second limit and returned the warm start, which does not meet the floor. The frontier's 100 % point failed for the same reason; the Q5 frontier chart leaves it off the curve and says so. OPT_NPV_RAI50's access floor was missed by 0.03 %, inside the solver gap and below the 0.5 % threshold, so it is not flagged by the interface, although the run log records it.

The run also asserts, while running, that the three budget views reconcile per fiscal year and pool. If they did not, the run would stop.

The interaction adjustment of TI-09 (Chapter 12) is itself a validation finding: every allocation's sum of stand-alone project gains overstates the network result by 20.6 to 86.9 %, above the 15 % warning threshold. That confirms the design's choice to recompute the network rather than sum projects.

16.4 Calibration: what has been compared with Uganda

Calibration compares the engine with observation and explains differences; it never silently rescales norms to match. The design lists six comparisons.

16.4.1 The Rural Access Index

The only published national benchmark is the World Bank's estimate of SDG 9.1.1 for Uganda: **53.1 % in 2015**. The engine's base-year RAI is **28.7 %** when unknown surface is counted as earth and **93.5 %** when it is counted as gravel. The benchmark lies inside the bounds, but the bounds are 65 points wide.

The width is the calibration finding. Surface is known on only 31.7 % of network km, and the difference between the bounds comes entirely from how the km of unrecorded surface are counted. The design asks for the difference from the benchmark to be explained by a decomposition: population grid, season rule, network completeness, and here surface. That decomposition has not been done. Three of its parts can be named:

- **Year and population.** The benchmark is for 2015; the engine uses 2026 population carried forward from UBOS 2022 totals.
- **Season rule.** The engine requires a link to be passable in the wet season, so seasonal fords and crossings cut access; the benchmark's season treatment is not documented in the estate.
- **Network completeness.** The engine routes on the full noded OpenStreetMap network of 248,000 km; how complete the network behind the 2015 estimate was is not documented in the estate.

Until surface is recorded, the engine cannot say whether Uganda's RAI is nearer 30 % or 90 %. Every scenario's RAI gain in Chapter 12 is conditional on the lower convention.

16.4.2 Health access

The design compares the engine's health access with SDI's health accessibility layer and with TransportModelling's accessibility outputs (88.1 % of the population within 30 minutes of a health facility, motorised). In the reference run, 93.8 % of the requiring population meets the primary-health standard in the wet season, and the outbound health file reports 79.7 % within 30 minutes and 94.0 % within 60 minutes by mixed modes. These figures use different modes, thresholds and facility registers, and they have not been reconciled. The comparison remains to be done.

16.4.3 Travel times and speeds

Speeds are TransportModelling's generic defaults, marked "NOT Ugandan" in the norms. TransportModelling's own speed calibration against the Nelson et al. travel-time surface was **rejected**: its multipliers pinned at their bounds, producing for example 250 km/h on motorways. The Nelson surface, travel time to cities and ports for 2015 on a hexagonal grid of about 10.7 km² (22,785 cells for Uganda), is on disk in the estate as a validation benchmark. **The comparison of TIPRAS travel times with it has not been done.** It is the cheapest calibration step available, because it needs no new data. TransportModelling's experience gives one lesson in advance: moving from a classified network to the full OSM network removed 56 % of its excess travel time over the Nelson benchmark, so the full noded network TIPRAS uses should already narrow the gap.

16.4.4 Traffic

No traffic counts are in the estate, so modelled traffic cannot be compared with observed counts. The design's test, the GEH statistic per count station,

$$GEH = \sqrt{\frac{2(M-C)^2}{M+C}} \tag{16.3}$$

where M is the modelled and C the counted hourly flow, with the share of stations below 10 reported, cannot be computed. The traffic loading is plausible in places and not in others: the reference run loads 288,725 of 817,382 links, and its largest modelled AADT is 935,037 vehicles a day, a value far above anything a Ugandan road carries. Freight is distributed by gravity because TransportModelling supplies no origin–destination matrices, and grown at a flagged fallback rate of 5 % a year. Traffic matters for the required standard (160,189 links bind on traffic) and for capacity gaps; both should be read with this in mind.

16.4.5 Vehicle operating costs and appraisal

Vehicle operating cost coefficients are TransportModelling values marked UNCALIBRATED, in assumed US dollars converted at an assumed exchange rate. Time savings on links closed in the wet season are valued at three times the improved time, a placeholder until a closure-cost norm exists. Only 574 of 41,757 appraised projects have a positive NPV at 12 %. Whether that is a fact about Uganda's candidate roads or about the coefficients cannot be said until they are calibrated.

16.4.6 Condition

Condition is known on **0 of 817,382 links**. No link-level survey (roughness or visual) exists in the estate; such data are held in the road agencies' asset management systems. Network condition comes from agency statistics by network (139,691 km reported poor). Without link condition, deterioration cannot be projected, rehabilitation is never triggered by the link-level decision rule, and the design's calibration of deterioration against successive surveys cannot start.

16.4.7 Maintenance cost

The design compares the modelled maintenance requirement at the actual funding ratio with the Road Fund's disbursement, and asks for any difference over 15 % to be explained. The modelled requirement is 954.7 bn UGX a year; the approved national road maintenance pool was 327.9 bn in FY2026/27 and the district maintenance transfers about 276.8 bn (176.0 bn through vote 609 and about 100.8 bn through the Road Fund). The Road Fund's annual reports could not be retrieved (the website was unreachable), so disbursement, the network length actually maintained and the force-account share are unknown. The comparison has not been done. The design's warning applies: if actual spending is well below the requirement, that may be a finding about under-maintenance, not an error in the norm.

16.5 What is and is not calibrated

Table 16.4: Calibration status

Item	Benchmark available	Status
RAI	World Bank SDG 9.1.1, 53.1 % (2015)	inside the engine's bounds (28.7–93.5 %); decomposition not done
Health access	SDI layer; TransportModelling 88.1 % within 30 min	not reconciled
Travel time	Nelson et al. surface (2015), on disk	**not done**
Speeds	none Ugandan	generic, "NOT Ugandan"; TransportModelling calibration rejected
Traffic	no counts in the estate	**not possible**
Vehicle operating cost	none	uncalibrated
Link condition	no surveys in the estate	**not possible**
Maintenance cost	Road Fund disbursement	reports unreachable; not done
Unit costs	ROCKS regional ranges, UCSN	compared (Chapter 14); evidence tier E5

The plausibility checks of the design (capital per km against regional ranges, maintenance as a share of asset value, trips and tonnes per person, cost per person gaining access, the spread of EIRR) are not yet produced as a standing report on each run. Some appear in this report's chapters, such as the evidence price comparison and the cost per person gaining access by scenario.

16.6 Implications for planning

- **The engine's accounting can be trusted; its calibration cannot yet.** Identities, pool rules and the no-composite rule are tested and checked on every run. Speeds, traffic, operating costs and condition are not calibrated.
- **Surface data matter more than any model refinement.** The RAI cannot be pinned between 28.7 % and 93.5 % until surface is recorded on the 68 % of km where it is unknown. A surface and condition survey is the first data investment the plan should make in itself.
- **The Nelson travel-time comparison should be done next.** The data are on disk and the method is defined; it would test the speed norms at no data cost.
- **Traffic counts and Road Fund disbursement data should be requested** from the road agencies and the Road Fund, through the data-request letters already drafted.
- **The test suite should grow with the engine**: frozen golden outputs, the missing identities and property tests, and a golden country whose maintenance funding ratio is below 1.
- **The OPT_RAI_EMONC80 and 100 % frontier failures should be fixed before any floor-based result is quoted.** The checks that exposed them are working as designed.

CHAPTER 17

17 Reference catalogues: norms and unit costs

Every number the engine uses to judge a road, a journey or a budget comes from a table of norms, not from the code. This chapter asks what standards, coefficients and prices the Uganda run rests on, and how strong each is. It implements design chapter 16 (`docs/16_REFERENCE_CATALOGUES.md`), which defines two kinds of reference data: **catalogues** of codes (road classes, surfaces, works types, pools) and **norm tables** in `tpl_norm` (standards, unit-cost mappings and coefficients, each with a source, a year and a quality tier). The chapter describes the norm set of run `014c9cd7`, `UGA-2026.09-draft-f`, table by table, then the envelope pool shares and the unit-cost map. The message is plain. The norm set is complete enough to run all ten modules, but 578 of its 646 rows are assumed, so every result in this report is **exploratory**.

17.1 How norms are held

A norm is one row in a long-format table. Each row carries:

- `tbl`, the norm table it belongs to;
- `key`, the case it applies to, for example `HEALTH_EMONC|rural` or `UP|*|gravel|*`;
- `geo_scope`, which is `national`, a region, an urban/rural class or a district;
- `field` and `value`;
- `unit`, `source`, `source_year`, `quality` and a free-text `note`.

The files live in `config/norms/<version>/`, one CSV per table. `tipras norms-import` loads a folder into the store as a new **draft** version. Approved versions cannot be changed: a database trigger refuses any insert, update or delete on them (see the database chapter). The engine resolves a norm by the most specific scope present (district, then region, then urban/rural, then national). If no row matches, it raises `NormNotFound`. It never falls back to zero or borrows a neighbouring category. Every row in `UGA-2026.09-draft-f` has scope `national`, so no district or regional variation is encoded yet.

The quality column takes one of the four tiers of docs/16 §16.13: measured, normed (a cited country standard or policy value), imputed or assumed (a default without country evidence). Prices carry their own evidence tier: **E5 evidence candidate** for prices built from documents, **E7 UCSN illustrative** for the placeholder coefficients published by UnitCostModelling (UCSN).

17.1.1 Version history

There have been six versions of the norm set since the first Uganda run. Each revision is a new version; none has replaced an earlier one in place.

Table 17.1: Norm versions in the store

Version	Tables	Rows	What changed
UGA-2026.09-draft	27	623	First set, from TransportModelling configuration and the manual's defaults
UGA-2026.09-draft-b	28	638	Adds `condition_programme` for the network-level condition backlog; maps rehabilitation, reseal, regravel and LVSR to evidence codes
UGA-2026.09-draft-c	28	648	Adds ferry rows: function standard, wet-season availability and an effective ferry speed
UGA-2026.09-draft-d	28	652	Extends ferry wet-season availability to every condition class

Version	Tables	Rows	What changed
UGA-2026.09-draft-e	28	652	Roads share and pool split replaced by approved-budget evidence (assumed → normed)
UGA-2026.09-draft-f	28	646	Unit-cost map points to the E5 evidence codes; one paved-upgrading code for all terrain

Table 17.1 shows that the set has grown little in size. The changes have been in substance. Draft-e replaced two assumed norms with values from approved budgets: the roads share moved from an assumed 0.6 to 0.545, and the assumed pool split (0.10, 0.45, 0.05, 0.05 and 0.35 for road-fund maintenance, government development, recurrent, district grants and donors) gave way to the evidence split described below. Draft-f re-pointed the unit-cost map from UCSN's illustrative codes to codes built from Ugandan and regional contract evidence. It has six fewer rows than draft-e because the three terrain-specific paved-upgrading keys collapsed into one: no source split paved upgrading costs by terrain. All six versions are still at status `draft`; none has been approved. Two further folders, `draft-f-env70` and `draft-f-env130`, are simulation variants. They differ from draft-f only in the roads share of the transport ceiling, which they set to 0.7 and 1.3 times the evidence value.

17.2 The norm set behind run 014c9cd7

`UGA-2026.09-draft-f` holds 646 rows in 28 tables. A count of the quality column in the CSV files confirms the figures in the run summary: 68 rows are `normed` and 578 are `assumed`. No row is `measured` or `imputed`. That is 10.5 % normed and 89.5 % assumed.

Table 17.2: The norm set UGA-2026.09-draft-f: tables and quality

Norm table	Rows	Quality tiers
access_standard	70	assumed,normed
ag_zone	2	assumed
appraisal	6	assumed
condition_programme	13	assumed
condition_target	1	assumed
design_standard	132	assumed
destination_eligibility	6	normed
deterioration	35	assumed
freight_generation	18	assumed
gravity_beta	2	assumed
group_share	5	assumed
investment_rule	11	normed,assumed
maintenance_norm	6	assumed
node_connection_standard	48	assumed
phasing_profile	18	assumed
pool_eligibility	5	normed
pool_share	6	normed
project_overheads	12	assumed
road_capacity	8	assumed
season_rule	46	assumed
speed_factor	13	assumed
speed_profile	98	assumed
spot_improvement	1	assumed
trip_generation	3	assumed
unit_cost_map	32	normed
vehicle	37	assumed

Norm table	Rows	Quality tiers
vehicle_mix	7	assumed
walk_leg	5	assumed

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_norms.csv).

Table 17.2 lists the 28 tables with their row counts and the quality tiers present in each. Twenty-two tables are wholly assumed. Four are wholly normed: `destination_eligibility` (6 rows), `pool_eligibility` (5), `pool_share` (6) and `unit_cost_map` (32). Two are mixed: `access_standard` (10 normed, 60 assumed) and `investment_rule` (9 normed, 2 assumed). The three largest tables, `design_standard` (132 rows), `speed_profile` (98) and `access_standard` (70), hold half of all rows (300 of 646). Two of these three are wholly assumed. The standards that decide what a road should be, and how fast people travel on it, are therefore the weakest part of the set.

17.2.1 Where the rows come from

The source column of each row names its origin. Table 17.3 groups the 646 rows by that origin.

Table 17.3: Norm rows of UGA-2026.09-draft-f by origin and quality

Origin	Normed	Assumed	Total
TransportModelling configuration, marked UNCALIBRATED or "not Ugandan"	0	172	172
Defaults written in the TIPRAS manual (docs/03–docs/16)	52	147	199
Illustrative values with no cited source, or flagged fallbacks	0	237	237
Approved Budget Estimates FY2024/25–FY2026/27	6	0	6
External published standards (Rural Access Index; WHO/Lancet two-hour EmONC standard)	10	0	10
Sector norms still to be imported from HealthModellingx and EducationModellingx	0	10	10
Phasing shape ("symmetric S-curve")	0	12	12
Total	**68**	**578**	**646**

Two cautions follow. First, only 16 of the 68 normed rows cite a source outside the TIPRAS manual: the six budget-derived pool rows and the ten rows of two external access standards. The other 52 cite the manual itself. They are rules of the method (which facilities count as a destination, the order of works, which pool pays for which works, which price code each works type uses). They are fixed by design, not guessed, but no Ugandan statute or circular stands behind them yet. Second, the 172 rows from TransportModelling (speeds, vehicle operating costs, deterioration, maintenance cycles) are marked there as UNCALIBRATED and, for the HDM-4-form coefficients, "not Ugandan". TIPRAS records them as assumed.

17.3 The tables one by one

This section follows the planning chain: needs, standards, supply, costing, and investment. The two allocation tables are covered in the envelope section.

17.3.1 Needs: who must reach what

`access_standard` (70 rows, 14 keys). This table defines when a person is served. Each key is a destination class and a rural or urban setting. Each has five fields: the measure (distance or travel time), the threshold, the travel mode, the season and the population group that needs the service. Table 17.4 summarises it.

Table 17.4: Access standards in UGA-2026.09-draft-f

Destination class	Rural threshold	Urban threshold	Mode	Group	Quality
ALL_SEASON_ROAD	2.0 km	0.5 km	walk	all	rural normed (Rural Access Index, World Bank 2016); urban assumed

Destination class	Rural threshold	Urban threshold	Mode	Group	Quality
HEALTH_PRIMARY	60 min	30 min	mixed	all	assumed (health policy to be imported from HealthModellingx)
HEALTH_EMONC	120 min	60 min	motorised	women 15–49	rural normed (WHO/Lancet two-hour standard); urban assumed
SCHOOL_PRIMARY	3.0 km	1.5 km	walk	children 6–12	assumed (education norm to be imported from EducationModellingx)
SCHOOL_SECONDARY	60 min	30 min	mixed	youth 13–18	assumed
MARKET	120 min	30 min	motorised	households	assumed (illustrative agriculture norm)
ADMIN_CENTRE	180 min	60 min	motorised	all	assumed

All 14 standards are tested in the wet season. Only two are cited standards: the Rural Access Index rule of 2 km to an all-season road, and the two-hour rule for emergency obstetric care (EmONC). Both apply to rural settings only. The ten normed rows of the table are these two standards' five fields each. The health and education thresholds are placeholders until the sector engines HPRAS and EPRAS supply their own norms. That matters because the access gaps (G1) for health and schools are measured against these thresholds.

`destination_eligibility` (6 rows, normed). This table states which node kinds count as a destination for each class. For example, `HEALTH_EMONC` accepts HC IV, general hospitals, and regional and national referral hospitals; `MARKET` accepts crop, livestock, periodic and urban markets. The source is the manual (docs/03 §1.4 and docs/16 §16.5). One difference from the manual should be noted. Docs/16 §16.4 says EmONC capability is a flag on a health node, not a facility level. The implemented table instead selects EmONC destinations by facility level. A health centre IV that does not in fact provide emergency obstetric care is still counted as an EmONC destination.

Five small tables complete the needs side, all assumed:

- **`group_share`** (5 rows): the population share of each requiring group (households 0.222, women 15–49 0.235, children 6–12 0.21, youth 13–18 0.15), pending the UBOS age structure.
- **`ag_zone`** (2 rows): a production density of 20 t/km² and a minimum of 5,000 t/yr define an agricultural zone (docs/03 §1.3).
- **`trip_generation`** (3 rows): 0.15 motorised trips per person per day in rural areas and 0.6 in urban areas, with 10 % leaving the district; to be replaced from Uganda National Panel Survey (UNPS) or household travel data.
- **`freight_generation`** (18 rows): marketed shares for 15 crops (coffee, tea, sugarcane and cotton 1.0; cassava and sweet potato 0.2), minerals and industry 1.0, and consumption of 0.35 t per person per year.
- **`gravity_beta`** (2 rows): gravity decay of 0.2 per hour for freight (TransportModelling, uncalibrated) and 0.5 per hour for market choice.

17.3.2 Standards: what each road should be

`design_standard` (132 rows, 18 keys, all assumed). This is the largest table. It holds two families of minimum standard. The first is the **function test**: twelve functional classes, each with a minimum surface, width, lanes, all-season flag, axle load, speed and structural capacity (84 rows, docs/16 §16.1). For example, a national trunk road must be DBST, 7 m wide, two lanes, all-season and rated for an 11.5 t axle. A community road must be engineered earth, 4 m wide and all-season. Tracks, service roads and ferries carry no road standard. The second family is the **traffic test**: six traffic bands, set out in Table 17.5 (48 rows, docs/05 §3.4).

Table 17.5: Traffic bands of the design standard (illustrative)

Band (AADT from)	Surface	Width (m)	Lanes	Axle (t)	Design speed (km/h)	Capacity (ESA)
0	engineered earth	5.0	1	8.2	30	50,000
50	gravel	6.0	2	8.2	40	100,000

Band (AADT from)	Surface	Width (m)	Lanes	Axle (t)	Design speed (km/h)	Capacity (ESA)
200	gravel	6.5	2	8.2	50	300,000
400	DBST	7.0	2	10.0	60	1,000,000
1,000	asphalt concrete	7.0	2	11.5	80	10,000,000
8,000	asphalt concrete	14.0	4	11.5	90	30,000,000

The required standard of a link is the highest of the function, traffic, sector-node and access tests. On run `014c9cd7` the function test bound on 656,479 links and the traffic test on 160,189. So these two assumed families decide almost every link's required standard. The source note says they are to be replaced by the national road design manual. Until then, the 242,068 km found below their required surface (gap G4) is a measure against illustrative standards. It is not a measure against Uganda's own.

`node_connection_standard` (48 rows, 8 keys, assumed). This table sets the road a sector node must sit on. For example, an agricultural zone producing at least 50,000 t/yr needs a national secondary road paved to DBST within 2 km. A mine shipping at least 2 Mt/yr needs a national trunk road, or rail, at an 11.5 t axle. Industrial parks and trade gateways need a paved national road. Regional referral hospitals need a paved district road. The thresholds are illustrative (docs/04 §2.3), and the note asks the sector ministries to confirm them. On run `014c9cd7` this test bound on 686 links.

`road_capacity` (8 rows, assumed) holds illustrative capacities of 900, 2,000, 5,600 and 8,400 PCU per hour for 1, 2, 4 and 6 lanes, a design-hour factor of 0.1, a target volume/capacity ratio of 0.8 and the BPR congestion parameters ($\alpha = 0.15$, $\beta = 4$). **`vehicle`** (37 rows) and **`vehicle_mix`** (7 rows), both assumed, hold PCU, equivalent standard axles (ESA) per vehicle, operating costs and payloads for seven classes, against the 16 of docs/16 §16.8. The default mix is 30 % motorcycles, 24 % cars, 18 % minibuses, 3 % buses and 13 %, 8 % and 4 % light, medium and heavy trucks. Most rows come from TransportModelling's uncalibrated HDM-4-form file; its operating costs are taken as US dollars, though the file states no currency.

17.3.3 Supply: how roads perform today

`speed_profile` (98 rows, assumed) gives the speed for each mode, class and surface: 88 rows are TransportModelling's generic sub-Saharan defaults (uncalibrated), 8 are walking speeds and 2 an assumed ferry speed including waiting. **`speed_factor`** (13 rows, assumed) multiplies speed by condition (good 1.0 to bad 0.55), terrain (flat 1.0, rolling 0.9, mountainous 0.7) and wet-season surface (earth 0.7, engineered earth 0.8, gravel 0.9). Unknown condition gets 1.0 and a flag; no IRI is invented. **`walk_leg`** (5 rows, assumed) sets walking at 5 km/h, a 10 km maximum snap distance and terrain factors of 1.0, 1.15 and 1.4, pending slope data from the Spatial Data Infrastructure (SDI).

`season_rule` (46 rows, assumed). This table says whether a link is open in the wet season, for each surface and condition. It also lists the all-season surfaces. Paved surfaces are open in every condition. Gravel is open unless `bad`. Engineered earth is open with drainage. Unengineered earth is open unless it is flood-exposed or a closure is recorded. Scheduled ferries are assumed to run all year. Unknown surface is treated as earth (docs/02 §2.9), which is the conservative choice.

`deterioration` (35 rows, assumed). This table holds the HDM-4-form structural and environmental deterioration coefficients by surface, the IRI ceilings, the IRI given to each condition class (3.5, 5, 7 and 11 m/km) and the IRI bands. The bands are 4, 6 and 10 m/km on paved roads and 8, 12 and 16 on unpaved roads, as in docs/16 §16.3. Because condition is unknown on every link, deterioration is not projected link by link on this run.

`condition_target` (1 row, assumed) and **`condition_programme`** (13 rows, assumed). The target is that every road should be in good or fair condition. The programme table converts the network-level poor-condition statistics into works: rehabilitation for poor national paved roads and community roads, periodic maintenance for national unpaved and district roads. It packages them in 1,000 km tranches. The treatments follow docs/07 §5.3 and are for the road agency to confirm.

17.3.4 Costing: what works cost

`unit_cost_map` (32 rows, 16 keys, normed). This table links each works type and surface transition to a price code and a quantity basis. It is described in its own section below.

`project_overheads` (12 rows, assumed). This table adds design and supervision and a contingency to the base works cost, for six works types. Table 17.6 lists them.

Table 17.6: Project overheads on base works cost (illustrative feasibility-stage rates)

Works	Design and supervision	Contingency	Total uplift
RH rehabilitation	0.08	0.12	0.20
UP upgrading	0.08	0.15	0.23
NC new construction	0.10	0.20	0.30
ST structures	0.10	0.20	0.30
PM periodic maintenance	0.05	0.08	0.13
SI spot improvement	0.08	0.15	0.23

The uplift is applied as one plus the two rates (docs/08 §6.3). On upgrading, which accounts for most of the programme's cost, the overheads add 23 % to the base price.

`maintenance_norm` (6 rows, assumed). This table holds the periodic-maintenance cycles (8 years for paved roads, 5 for gravel and engineered earth) and the IRI to which each type of works resets a road (2.0 after rehabilitation, 2.8 after periodic maintenance, 1.8 after upgrading). The values come from TransportModelling and are marked not Ugandan. These cycles turn periodic maintenance into an annual figure. With the routine-maintenance prices, they give the **induced maintenance** that every capital item carries.

`spot_improvement` (1 row, assumed). This table sets one vented ford per seasonal crossing.

17.3.5 Investment and appraisal

`investment_rule` (11 rows: 9 normed, 2 assumed). The nine normed rows fix the order in which the decision rule considers works: routine maintenance first, then periodic maintenance, spot improvement, rehabilitation, upgrading, structures, capacity expansion, new construction and climate adaptation (docs/09 §7.2). The two assumed rows are the minimum population a new link must serve (2,000 persons) and the maximum length of a procurement package (60 km).

`phasing_profile` (18 rows, assumed). This table spreads each works type's cash over 1 to 3 years with a symmetric S-curve (shape parameters a = b = 2). Rehabilitation and structures take 2 years, upgrading and new construction 3, and maintenance and spot improvements 1. The run notes say this phasing puts little cash in a project's first year. That is one reason some donor money stays unspent within the horizon.

`appraisal` (6 rows, assumed) holds a 12 % discount rate, a 20-year period, a value of time of UGX 3,000 per person-hour, UGX 3,700 per US dollar, an occupancy of 2.2 and a fallback traffic growth of 5 % per year. The discount rate awaits the MoFPED public investment rate and the exchange rate the FINEX path. The run used the growth fallback because sector output trajectories were not supplied. Appraisal results (NPV, EIRR) are therefore exploratory twice over.

17.4 The envelope: roads share and pool split

TI-08 builds the money available each year from the MTEF transport ceiling. Two norms split it: the roads share of transport and the split of roads money across financing pools. The amount in pool p in year t is:

$$E_{p,t} = C_t \cdot \rho \cdot \sigma_p \tag{17.1}$$

where

- C_t is the final-round MTEF transport ceiling for year t (UGX), read from MTEFModelling;
- ρ is the roads share of transport (`pool_share`, key `roads`);
- σ_p is pool p's share of roads money (`pool_share`, key `transport_ceiling`), with the five shares summing to one.

From draft-e on, ρ and σ_p come from the approved budget estimates for FY2024/25 to FY2026/27, read line by line (`data/reference/budget_roads_evidence_UGA.md`). Each line of programme 09, Integrated Transport Infrastructure and Services (ITIS), across votes 016, 113, 118, 122 and 609, was flagged as roads, not roads or unclear, and each roads line was assigned to a pool by funding column and vote.

Table 17.7: Roads share and pool split from approved budgets

Item	FY2024/25	FY2025/26	FY2026/27	Norm used
Transport (ITIS) total, UGX bn	4,968.2	6,602.1	8,827.2	–
Roads, UGX bn	3,005.7	3,572.1	4,813.6	–
Roads share (strict)	0.605	0.541	0.545	**ρ = 0.545** (FY2026/27)
ROAD_FUND_MAINT share	0.099	0.080	0.069	**0.080** (3-year pooled)
GOV_DEV share	0.212	0.247	0.469	**0.333**
GOV_RECURRENT share	0.036	0.015	0.017	**0.021**
DISTRICT_GRANT share	0.109	0.097	0.069	**0.088**
DONOR share	0.544	0.561	0.377	**0.478**

Table 17.7 shows how each norm was chosen. The roads share is the FY2026/27 value, 4,813.6 of 8,827.2 bn UGX. It is the latest approved year and the MTEF base year, and it matches FY2025/26 closely (0.541). The FY2024/25 value, 0.605, comes before the external financing for the standard gauge railway (SGR) ramped up. The alternative definitions give 0.558 (3-year pooled), 0.541 (excluding arrears) and 0.568 (broad, counting unclear lines as roads).

The pool split is the amount-weighted mean of the three years, on a base of 11,160.7 bn UGX of roads spending excluding arrears. A pooled mean was used because two lumpy items dominate single years:

- the timing of project-loan disbursements, which moves the donor share from 0.56 to 0.38;
- Land Acquisition Project II, at 689.5 bn in FY2026/27 against 150.0 and 113.1 bn in the two earlier years, which lifts government development to 0.469 in that year.

By contrast, the maintenance and grant pools are almost flat in money terms, at about 275 to 330 bn UGX each per year.

`pool_eligibility` states which works each pool may fund:

- **ROAD_FUND_MAINT**: routine and periodic maintenance (RM, PM).
- **GOV_RECURRENT**: routine maintenance (RM).
- **DISTRICT_GRANT**: RM, PM, spot improvement (SI) and rehabilitation (RH).
- **GOV_DEV**: SI, RH, upgrading (UP), structures (ST), new construction (NC), capacity expansion (CE), climate adaptation (CA) and PM.
- **DONOR**: RH, UP, ST, NC, CE and CA. No maintenance.

The last rule matters most. Donor money is 47.8 % of the envelope, and it cannot pay for routine or periodic maintenance. This is one reason the steady-state routine maintenance requirement of 1.02 tn UGX/yr meets only 0.59 tn UGX/yr of eligible resources in the final plan year, a funding ratio of 0.58.

Three caveats apply. These are **approved estimates, not outturn**, and external financing usually under-executes. Vote 609's 176.0 bn transfer for district, urban and community access roads (DUCAR) may overlap the Road Fund's DUCAR transfer of about 100.8 bn per year; if so, the district grant share would fall by up to about 0.03. And the

road-fund maintenance pool from FY2025/26 on is an inference made after the Road Fund was mainstreamed into MoWT.

The implemented pools differ from docs/16 §16.11: donor money is one pool, not one `DONOR_<id>` per financier, and `PPP`, `SECTOR_COFIN` and `CONTINGENCY` are absent. The simulation variants `env70` and `env130` keep the label `normed` on a roads share that is a multiplier of the evidence value; only their note says SIMULATION.

17.5 The unit-cost map

`unit_cost_map` is how the engine prices a works item without holding any price in code. Each key has the form works | from-surface | to-surface | terrain, where `*` is an explicit wildcard written in the data. The engine tries the most specific key first and stops at the first match. It never substitutes one specification for another. A missing key or price leaves the item **unpriced**: it is reported and cannot be funded, and it is never costed at zero. Table 17.8 lists the 16 keys of draft-f with their codes and the central prices of the run's price vintage.

Table 17.8: Unit-cost map of UGA-2026.09-draft-f and the prices it resolves to (UGX m, 2026)

Map key	Price code	Basis	Central (E5)	Range (E5)	UCSN illustrative central (E7)
RM paved	TRN-OPS-004-RM-PAVED-EV	per km per year	16.4	9.7–24.3	18 (ROUTINE-PAVED)
RM unpaved	TRN-OPS-003-RM-UNPAVED-EV	per km per year	4.4	1.0–8.9	10 (ROUTINE-GRAVEL)
PM paved	TRN-OPS-002-RESEAL	per km	67.6	64.3–69.8	none
PM unpaved	TRN-OPS-002-REGRAVEL	per km	25.6	9.3–38.4	none
RH paved	TRN-CAP-002-RH-PAVED	per km	2,339.0	1,645.7–4,152.4	none
RH gravel	TRN-CAP-002-RH-GRAVEL	per km	58.6	45.0–132.3	none
RH unpaved	TRN-CAP-002-RH-EARTH	per km	66.1	45.5–77.9	none
UP to engineered earth	TRN-CAP-006-EARTHENG	per km	42.1	33.7–50.6	none
UP to gravel; NC gravel	TRN-CAP-008-GRAVEL-EV	per km	75.1	42.2–193.9	500 (GRAVEL)
UP to LVSR	TRN-CAP-007-LVSR	per km	735.2	222.7–1,075.4	none
UP to DBST or AC; NC DBST (flat)	TRN-CAP-009-PAVED-EV	per km	4,269.0	2,580.8–5,690.3	3,100 / 3,700 / 5,200 (flat / rolling / mountainous)
ST any	TRN-CAP-004-BRIDGE-EV	per m of span	86.2	75.4–86.2	180 (BRIDGE)
SI any	TRN-CAP-005-SPOT	per site	141.0	141.0–141.0	none

Three points stand out in Table 17.8. First, every map row in draft-f points to one of 13 evidence codes (tier **E5, evidence candidate**, confidence LOW). None points to UCSN's own coefficients. The run's price vintage, `UCSN+EV-UGA-2026.09b`, holds 21 codes: the 14 evidence codes and UCSN's 7 illustrative transport coefficients (tier **E7, ILLUSTRATIVE**, confidence VERY_LOW). The 14th evidence code, a pedestrian footbridge (UGX 253 m per bridge), is priced but no map row uses it yet. Draft-e still mapped routine maintenance, gravel and paved upgrading, and bridges to the UCSN codes.

Second, the evidence prices come from documents: 89 of 275 rows extracted from MoWT reports, World Bank and AfDB documents, the ROCKS regional database, Road Fund guidelines, BMAU contract reviews and project records, carried to 2026 UGX with Uganda's GDP deflator. A code's central value is the median of rows in the best evidence tier present. The evidence is thin in places. The bridge price rests on one Bailey bridge, the spot-

improvement price on one force-account site, and engineered earth on two observations. Gravel construction rests on 18 rows and paved upgrading on 14.

Third, the evidence and UCSN prices differ a lot where both exist. Gravel construction is UGX 75.1 m per km on evidence against UCSN's 500 m, about one-seventh. Bridges are 86.2 m against 180 m per metre of span. Unpaved routine maintenance is 4.4 m against 10 m per km per year. Paved upgrading is close: the evidence central of 4,269 m per km sits within UCSN's terrain range of 3,100 to 5,200 m. The comparison has been sent to UnitCostModelling. It shows the range of current knowledge, not a correction of UCSN.

17.5.1 What the map changed

Moving to the evidence codes had a large effect on what could be priced. Under UCSN's codes (run 5, norms draft-e), about 8,000 of 40,844 candidate projects were priced. Most of the rest were engineered-earth upgrades for which UCSN has no code. Under draft-f (run 8, `014c9cd7`), 41,867 of 41,890 projects are priced. The community-road condition backlog of 70,656 poor km, which was unpriced before, is now costed at 5.6 tn UGX. The programme's capital cost is 92.8 tn UGX central (56.6–125.3 tn low–high), with induced routine maintenance of about 205 bn UGX per year. Upgrading accounts for 85.5 tn UGX of it.

17.5.2 What the map still misses

- **Spot improvements on paved and gravel links.** Twenty-three projects stay unpriced: 13 on gravel links and 10 on asphalt links. Their keys, such as `SI|gravel|gravel|flat`, do not reach the map's single spot-improvement row (`SI|*|*`).
- **New paved links outside flat terrain.** The map holds `NC|none|dbst|flat` only. A new paved link in rolling or mountainous terrain would find no row. No new-link candidates arise on this run, because no candidate alignments were supplied.
- **Works types without any row.** New crossings (NX), node expansion (NE), rail (RR), capacity expansion (CE), climate adaptation (CA) and emergency repair (EM) are in the catalogue of docs/16 §16.6 but have no map row. Any such item would be unpriced.
- **Paved to paved.** There is no key for upgrading DBST to asphalt concrete. The costing studio in the web interface notes that such a transition is priced at the full paved-upgrading rate.
- **Right of way.** The proposed UCSN code `TRN-ROW-001` does not exist. Land costs are not in any unit price.

17.6 The catalogues (tpl_ref)

Docs/16 also specifies reference catalogues to be seeded into a `tpl_ref` schema, from functional classes and surfaces to hazards, pools and a budget programme skeleton. In the Uganda build these are not yet seeded tables. The folder `database/seed/` is empty, and the SQLite migrations create no `ref_` tables. The vocabularies are applied in other ways:

- the surface rank order (none 0, earth 1, engineered earth 2, gravel 3, LVSR 4, DBST 5, AC and concrete 6) is a constant in `src/Core/Types.jl`;
- the functional-class rank is a constant in the needs module;
- surface and class codes are mapped from SDI and OSM tags when the geography snapshot is built;
- node kinds, destination classes, works types and pools appear as keys in the norm tables.

A rank is an ordering, not a country parameter, so these constants do not break the rule against coefficients in code. But the catalogue rows with their ATDGR source codes are not loaded, and the seed load that fails on an unmapped code (docs/16 §16.14) does not yet exist.

17.7 Data quality and limitations

- **Assumed dominates.** 578 of 646 norm rows (89.5 %) are assumed. All of the design standards, speeds, deterioration, vehicle and appraisal coefficients are assumed. Every module result that depends on them inherits the assumed tier, and the run is exploratory by rule: norms in `draft` status cannot drive a non-exploratory run.
- **"Normed" is generous for 52 rows.** They cite the TIPRAS manual, not a Ugandan standard. Only 16 normed rows cite an external source.
- **No local variation.** Every row is national. District or regional standards, for example for mountainous districts or urban councils, can be added as scoped rows without code changes, but none exist.
- **Sector norms are placeholders.** The health, education and agriculture access thresholds await the sector engines.
- **Prices are candidates.** All 14 evidence codes are E5 EVIDENCE_CANDIDATE with low confidence. Several rest on one or two observations. UCSN's 7 transport codes are E7 ILLUSTRATIVE. A non-exploratory run would refuse both tiers.
- **Budget norms are estimates, not outturn.** The pool split assumes approved external financing is spent. The district grant share may be overstated by up to 0.03.
- **Catalogues are not seeded.** The `tpl_ref` layer of docs/16 exists only as constants and keys.

17.8 Implications for planning

- The quickest way to raise the standing of the whole plan is to replace the two largest assumed tables with Ugandan standards: `design_standard`, from the national road design manual, and `speed_profile`, from calibrated speeds. Together they decide almost every link's required standard and every travel time.
- The health and education access thresholds should come from the Ministry of Health and Ministry of Education through HPRAS and EPRAS. The G1 gaps for these classes are measured against placeholders until then.
- The 52 manual-cited normed rows (destination eligibility, works order, pool eligibility, unit-cost map) should be reviewed by MoWT and MoFPED and either confirmed by circular or replaced. Their label should then reflect that review.
- UnitCostModelling should be asked to adopt or challenge the 14 evidence codes, especially gravel construction (UGX 75 m against UCSN's 500 m per km), bridges and spot improvements. The thin codes should be strengthened with completed-contract data before any allocation is treated as more than exploratory.
- The unit-cost map needs rows for spot improvements on paved and gravel links (23 projects unpriced), new paved links in rolling and mountainous terrain, and the works types that have no row today.
- The pool split should be re-estimated from releases or outturn when they are available. The vote 609 and Road Fund DUCAR overlap should be confirmed with MoFPED.
- Donor eligibility excludes maintenance, and donors provide 47.8 % of the envelope. Any change to that rule, or a separate financier-level `DONOR_<id>` split, would change what the allocation can fund. Such a change should be made as a new norm version, not as an edit.
- The first norm version put forward for approval should be accompanied by norm cards for the ministries. Approval makes a version immutable, so corrections after approval need a new version.

CHAPTER 18

18 Build status against the build plan

How far has the engine been built, how much of it is tested, and what is still missing before a Uganda result can be used for more than exploration? This chapter answers that question for the build plan in design chapter 17 (`docs/17_BUILD_PLAN.md`). The plan sets nine phases, from foundations (phase 0) to the interface and reports (phase 8), each with a gate that must pass before the next phase starts. The record behind this chapter is `IMPLEMENTATION_STATUS.md` (status as of 23–24 September 2026), the eight Uganda run logs in `outputs/`, the test suite in `test/runtests.jl` and the source tree in `src/`. Every Uganda run so far is **exploratory**: the norms are a draft (`UGA-2026.09-draft-f`, 646 rows, of which 578 are of quality *assumed* and 68 *normed*) and the prices are evidence candidates (tier E5). Nothing in this chapter changes that status.

18.1 What the build plan asks for

The build plan is short. It names nine phases and one gate for each. It also names a critical path of four items that sit outside the code: norms, unit costs, surface and condition data, and sector feeds. Table 18.1 restates the plan so that the status tables that follow can be read against it.

Table 18.1: The nine build phases and their gates (docs/17)

Phase	What is built	Gate to pass
0 Foundations	repository, migrations 0001–0003, `Core/`, `Data/DB.jl`, CI	migrate and test pass on a clean database; `explain()` returns a trace for a toy cell
1 Geography and GIS core	geography snapshot from SDI and ATDGR; graph, snapping, travel time, catchments, spatial statistics; the TransportModelling ports; the `TST` geography	GIS-1 to GIS-5 pass for Uganda; baseline travel time on the full Uganda graph within target; GIS unit tests pass
2 Norms	norm tables, seed catalogues, `resolve_norm`, XLSX norm workbook, versioning and approval	every norm table needed by TI-01 to TI-06 resolves for `TST` and for Uganda; norm cards print
3 Needs chain	TI-01, TI-02, TI-03; sector inbound contracts; gravity fallback; assignment	identity tests 1–4 pass; required standard and binding test on every loaded link
4 Supply, gaps, cost	TI-04, TI-05, TI-06; condition projection; season rules; criticality; price resolver	Q2–Q4 answerable; GIS-6 reconciled or explained; every capital figure shows induced maintenance; no composite column
5 Investment and allocation	TI-07 and TI-08 with scenarios A–G	seven scenarios under one envelope, with earmarks, dependencies and maintenance-first binding on `TST`
6 Optimiser and outcomes	`Optimise/`, coverage sets, ε-constraint, frontier; TI-09; ranges; sector outbound	the optimiser matches or beats every rule-based scenario on its objective; outbound files validate against contracts
7 Budget and integration	TI-10; FINEX, MTEF, CDCGE and MEL exchanges	a budget submission that reconciles across three views; FINEX round trip converges or reports oscillation
8 Interface and reports	Genie screens, map component, year slider, reports, atlas	a new user answers Q1–Q5 for Uganda in under ten minutes and can trace any number to its source

Two features of the plan matter for reading the status. First, the gates are mostly about *behaviour* (an identity holds, a check passes), not about code existing. Second, several gates are written for Uganda, not only for the

golden test country. A phase can therefore be "built" in code and still fail its gate because the Uganda data cannot satisfy it.

18.2 Status by phase

Table 18.2 sets out the state of each phase as recorded in `IMPLEMENTATION_STATUS.md`, with the gate result read from the code, the tests and the run logs.

Table 18.2: Build status by phase, with the gate result

Phase	State recorded	Where in the code	Gate result
0 Foundations	done	`src/Core/*`, `src/Data/DB.jl`, `database/migrations/sqlite/0001–0003`; `explain()` walks `audit_factor_trace`, which is append-only by trigger	passed on SQLite, not on PostGIS; no CI configuration in the repository
1 Geography and GIS core	done	`src/GIS/*` (CSR graph, super-source Dijkstra on the reversed graph, link snapping, Tarjan cut edges, Moran's I, LISA, concentration index, Gini); `scripts/build_geo_snapshot.jl`	GIS unit tests pass; **GIS-5 fails for Uganda** (68.2 % of network km has unknown surface, against a 30 % limit)
2 Norms	done (draft)	`scripts/seed_norms.jl`, `src/Data/Norms.jl`; `resolve_norm` raises `NormNotFound`; approved versions are immutable by trigger	every table resolves for `TST` and Uganda; no XLSX workbook or printed norm cards found in the code
3 Needs chain	done	`TI01`–`TI03`: raking, agricultural zones, access requirement, node standards, generation, Furness gravity, two-stage agricultural chain, all-or-nothing loading, AADT and ESA, the required-standard join with its binding test	binding test runs on every loaded link; identities 1 and 2 are tested, 3 and 4 are not named in the suite
4 Supply, gaps, cost	done	`TI04`–`TI06`: season rules, do-nothing condition, current access, the seven gaps with the access-cause decomposition, prices (low, central, high), maintenance and induced maintenance	Q2–Q4 answered for Uganda; schema test for no composite column passes; GIS-6 not coded as a check
5 Investment and allocation	done	`TI07`–`TI08`: the ordered decision rule, packaging, coverage sets, appraisal (NPV, EIRR, BCR), sector tags, pools and eligibility, scenarios A–G	scenarios, pools and maintenance-first tested on `TST`; project dependencies not found in the allocation or optimiser code
6 Optimiser and outcomes	done (core)	`src/Optimise/Portfolio.jl` (JuMP/HiGHS MILP, money scaled to bn); `TI09` re-simulates the network per allocation	optimiser beats every rule-based scenario on `TST` (tested); on Uganda run 8 it equals scenario C; outbound contracts not validated (`config/interfaces/` is empty)
7 Budget and integration	partial	`TI10`: three reconciling views, fiscal sustainability statement, outbound files for MTEF, FINEX, HPRAS, EPRAS and Agriculture	three views reconcile (asserted in code); no FINEX inbound reader, so no round trip
8 Interface and reports	partial	`app.jl`, `web/` (12 pages), `scripts/render_maps.py` (PDF atlas), `bin/tipras.jl` CLI	Q1–Q5 pages exist; year slider, equity slider, narrative reports and authentication missing; the ten-minute user test has not been run

Read across the rows, the engine is complete in code from phase 0 to phase 6, and partial in phases 7 and 8. That is a fast build: the logs show the first Uganda run on 23 September 2026 and the eighth on 24 September. The gates tell a more cautious story. The phase 1 gate fails for Uganda on data, not on code. The phase 3 and phase 5 gates are met only in part. The phase 7 gate cannot be judged because FINEX does not yet send an envelope back.

18.2.1 Phase 0: foundations

The store is a single SQLite file per workspace, with table-name prefixes standing in for the specified PostgreSQL schemas (`ref_`, `norm_`, `obs_`, `geo_`, `res_`, `audit_`, `stage_`). Three forward-only migrations create the schemas, the norm tables and the observation and result tables. Large geometry is held in Arrow files under `data/geo/<ISO3>/<geo_version>/`. Every derived number writes a factor trace, and `explain()` walks that trace back to its inputs. The audit table is append-only: a database trigger refuses updates and deletes.

The CLI (`bin/tipras.jl`) offers seven commands: `migrate`, `norms-import`, `norms-approve`, `prices-load`, `run`, `prune` and `explain`. The geography snapshot is built by a script (`scripts/build_geo_snapshot.jl`), not by a `tipras geo snapshot` command as the plan names it. No continuous-integration configuration was found in the repository; the tests are run by hand.

18.2.2 Phase 1: geography and GIS core

The planning graph is a compressed sparse row (CSR) structure held twice, forward and reverse, so that travel time can be computed *to* the nearest destination on one-way links. Travel time uses one Dijkstra search per destination class from a "super-source" joined to every destination. Settlements and destinations are snapped to the nearest point on a link, not to the nearest vertex. Cut edges (links whose loss splits the network) come from Tarjan's bridge algorithm, run iteratively so that it does not overflow the stack on 817,354 links.

The Uganda geography snapshot has gone through three versions (`UGA-2026.09-a`, `-b`, `-c`). Version c is the one used in runs 4 to 8. It carries 817,354 links and 247,972 km of network, 24,952 sector nodes and 135 districts. Its main weakness is the one GIS-5 detects: surface is known on 31.7 % of network km and unknown on 68.2 %. Above the 30 % limit, the engine flags every access figure as exploratory. This is a data finding, and the gate will stay failed until a road inventory arrives.

One item of the phase 1 plan was not built: the degree-2 simplification of the planning graph. The engine routes on the full noded network. This keeps every link's identity, but it is one reason why a national run takes between 8 and 90 minutes rather than a few minutes (see Table 18.5).

18.2.3 Phase 2: norms

`scripts/seed_norms.jl` builds the draft norm set from the TransportModelling configuration files and the manual's defaults. The first draft held 623 rows in 27 tables. The current draft-f holds 646 rows in 28 tables. Each row carries a source, a year and a quality tier. `resolve_norm` looks for the most specific match (district, then region, then urban or rural, then national) and raises `NormNotFound` when nothing matches. It never defaults to zero and never borrows a neighbouring category. An approved norm version cannot be changed, by trigger.

The norm set has gone through six versions in four days (`draft`, `draft-b` to `draft-f`), plus two envelope sensitivity variants (`draft-f-env70`, `draft-f-env130`). The changes are data, not code: the budget evidence for pool shares entered in draft-e, and the evidence price map in draft-f. That is how the plan intends norms to evolve.

The plan's XLSX norm workbook and printable "norm cards" were not found in the code. Norms are edited as CSV files under `config/norms/<version>/`.

18.2.4 Phases 3 and 4: needs, supply, gaps and cost

Modules TI-01 to TI-06 run end to end on Uganda. TI-03 produces a required standard for every loaded link and the reason it binds: 656,479 links bind on function, 160,189 on traffic, 686 on a sector-node standard and 28 on mode. TI-04 reports the rural access index (RAI) as a pair of bounds, because unknown surface may be earth or gravel: 28.7 % to 93.5 %, against the published benchmark of 53.1 % (2015). TI-05 reports the seven gap dimensions, G1 to G7, each in its own unit and never combined. TI-06 prices works through a map from works type and surface transition to a price code.

Two parts of phase 4 are present in code but idle on Uganda data. The deterioration projection (an HDM-4 form ported from TransportModelling) runs only on links with a known condition, and no link has one. The rehabilitation trigger needs link condition for the same reason. Condition enters instead as agency statistics by network (for example 35,320 poor km of 41,194 km on district roads, FY2022/23), which TI-07 turns into condition programmes in 1,000 km tranches.

18.2.5 Phases 5 and 6: investment, allocation and the optimiser

TI-07 builds a programme of 41,890 projects in run 8, of which 41,867 are priced. Its central capital cost is 92.8 tn UGX (56.6 tn low, 125.3 tn high), and the routine maintenance it would induce once complete is 205.2 bn UGX a year. Upgrading accounts for 85.5 tn of the central capital. TI-08 allocates the envelope under seven rule-based scenarios (A to G) and four optimiser specifications, each named by its objective and floors.

The optimiser is a mixed-integer programme solved by HiGHS. Each specification declares one objective in native units (persons served, NPV in bn UGX, or km of poor road treated) and turns every other measure into a soft floor with a reported shortfall. There is no weighted sum. A specification that cannot meet its floor reports the shortfall as a finding.

The plan's phase 5 gate names project dependencies. No dependency constraint was found in `TI08_Allocation.jl` or `Optimise/Portfolio.jl`. The TransportModelling portfolio model has one (chapter on TransportModelling reuse), but it was not carried over.

18.2.6 Phase 7: budget and integration

TI-10 produces the budget in three views (administrative, economic and programme) and asserts that they sum to the same total: 20.9 tn UGX in run 8 for scenario B. It writes a fiscal sustainability statement and outbound files for MTEF, FINEX, HealthModellingx, EducationModellingx and Agriculture. The inbound side is thinner. The envelope comes from MTEFModelling's ceiling file, multiplied by a roads share (0.545) and pool shares that are norms of quality *normed*, from the approved budget estimates FY2024/25 to FY2026/27. There is no FINEX reader, so the fiscal loop of the plan (FINEX sends an envelope, TIPRAS returns a requirement, repeat until it converges) has not been closed. The folders `src/Integration/` and `src/Report/` exist but are empty.

18.2.7 Phase 8: the interface and reports

The first status record said "no Genie web interface yet". It was rebuilt on 24 September 2026 on the EnergyModellingx pattern. It has 12 pages, key-figure tiles with quality chips, charts with their tables, district maps with a scenario-difference mode, a GIS explorer, district transport profiles, exports, and a set of validation checks shown on the pages. A costing studio lets a user cost a chosen set of roads at a chosen standard, and a run button launches a full run as a separate process. Still missing are the live equity slider, the year slider, the narrative reports of design chapter 14 §14.4 and authentication. The plan's acceptance test (a new user answers Q1–Q5 in under ten minutes) has not been run.

18.3 Deviations from the specification

Some differences from the specification are deliberate and recorded. Table 18.3 lists them with the reason and the step that would close each.

Table 18.3: Deliberate deviations from the specification and how to close them

Specification	Built instead	Reason and next step
PostgreSQL with PostGIS as the primary store	SQLite (the specification's offline mode) plus Arrow geography snapshots	the local PostgreSQL on port 5432 needs credentials that were not available; SDI's cluster on port 5433 was used read-only to export the network. Port the migrations to PostGIS when a TIPRAS database is provisioned
Genie interface (docs/14)	rebuilt 24 September 2026 on the EnergyModellingx pattern, still on Genie	live equity slider, year slider, narrative reports and authentication still missing
Candidate alignments for new links (G5 "network missing")	not supplied	needs alignment output from SDI or TransportModelling; the G1 cause "network missing" is reported as not assessable
Rail, ports and waterways as links	nodes only (ferries are links since run 4)	phase-2 scope in docs/00 §0.6
Year-by-year outcome path	horizon-year state only	TI-09 recomputes the network at 2030, not every year
Uncertainty ranges on outcomes	cost low and high carried; not propagated to NPV	add the corner runs of docs/23 §23.21
Access class weights and ε-floors in the optimiser	done 24 September 2026	specifications in `config/scenarios/optimiser_UGA.yml`

The SQLite store is the deviation with the largest practical effect. A national run adds about 2 GB to the store, and disk D: is nearly full, so `tipras prune --keep N` was added and runs warn below 6 GB free. A PostGIS store would remove this constraint and allow several users to work against one database.

18.4 The test suite

The tests live in one file, `test/runtests.jl` (185 lines). They are grouped in one top-level testset with nine testsets inside it. The file has 55 `@test` lines, which expand in loops to many more checks. The last recorded run of the suite (`outputs/test_run.log`, 24 September 2026, 07:19) reports **2,265 assertions passed out of 2,265**, in 3 minutes 12 seconds. Table 18.4 lists the testsets.

Table 18.4: The testsets in test/runtests.jl

Testset	What it checks
unit: standards lattice (docs/23 §23.7)	the join of two road standards takes the higher of each attribute; equal ranks for AC and concrete; unknown surface ranks as earth for access
unit: ESA fourth-power example (§23.6)	the equivalent-standard-axle factor of the manual's example (4.80) and its design total (1.89 million)
unit: Furness 2×2 (§23.5)	the gravity balancing reproduces row and column totals within 0.1 % in fewer than 100 iterations
unit: Kingman delay example (§23.9)	the queueing delay of the manual's example (85.0)
unit: spatial statistics (§23.16)	Gini is 0 for equal values; the concentration index has the right sign; Moran's I is high on a clustered line
GIS: Dijkstra, snapping, cut edges (docs/20 §20.4a)	shortest paths against brute force, one-way handling, snapping, cut edges, the projected-coordinate guard (GIS-4)
norms: resolution and refusals (docs/02 §2.4)	most-specific resolution, `NormNotFound`, immutability of approved versions, the append-only audit
golden: TST end to end (docs/15 §15.4)	identities 1, 2, 5, 6 and 7; the no-composite schema check; pool and eligibility compliance; unpriced items never funded; the optimiser beats every rule-based scenario on its objective; the cut-edge and mine-axle golden facts; `explain()`
property: monotone budget (docs/15 §15.2.2)	more budget never reduces the optimiser's objective

18.4.1 The golden country

The golden test country `TST` is built by `test/fixtures/build_tst.jl`. It is the only synthetic data in the repository, and it lives under `test/`. It is a grid of 30 by 40 vertices at 2 km spacing, split into 12 districts. It is deliberately awkward. A river runs between columns 20 and 21, with one permanent bridge, so that bridge is a cut edge in the wet season. A seasonal drift crosses the river further north and closes in the wet season. A flood-exposed zone covers one corner. Terrain changes from flat to rolling to mountainous across the grid. A fixed random seed (15) makes every golden output repeatable. Each of these features exists to test one rule: the cut edge tests G5, the drift tests the season rule, and the flood zone tests G7.

The `TST` end-to-end run in the test log shows the whole chain working on a small case. The seven scenarios fund between 27 and 143 projects from an envelope of 1.02 tn over four years. The optimiser's frontier between rural access and primary-health access moves from 253,931 persons served at one end to 83,152 at the other, with health access rising to its attainable 8,355. The budget views reconcile, with a funding ratio of 1.18.

18.4.2 What the tests do not cover

The suite has three gaps that matter.

- **It predates the last optimiser change.** `src/Optimise/Portfolio.jl` was last modified on 24 September at 09:04, after the recorded test run at 07:19. That change (round-robin candidate lists and a shared warm start) produced the run 8 results. The suite should be re-run, and its result recorded, before run 8 is quoted as tested.
- **Identities 3 and 4 are not named.** The plan's phase 3 gate asks for identities 1 to 4. The golden testset names 1, 2, 5, 6 and 7.
- **Nothing tests Uganda data.** The tests prove the method on `TST`. Plausibility on Uganda is reported by each run's log and fallback list, not tested. The regression benchmark that docs/22 proposes (TransportModelling's health access within 30 minutes, 88.1 % of population) has not been added.

18.5 The Uganda run history

Eight Uganda runs were made on 23 and 24 September 2026. Each one changed one block of inputs. Table 18.5 summarises them from the run logs and the status record. "Module time" is the sum of the TI-01 to TI-10 times in the log.

Table 18.5: Uganda runs 1–8: inputs, run time and headline result

Run	Run id	Geography / norms / prices	Module time	What changed	Headline result
1	a2711deb	a / draft / UCSN illustrative	14.5 min	first full national run	62,298 projects appraised, 1,760 with NPV > 0; 25 unpriced codes. Superseded 15 minutes later
2	bb2c3251	a / draft / UCSN illustrative	11.0 min	first results recorded in the status file	RAI 20.2 % to 93.6 %; 8,116 projects priced of 40,844; optimiser RAI at 2030 0.594
3	a2c3d78d	b / draft-b / UCSN + evidence	7.6 min	HeiGIT and GRIP4 surface; agency condition by network; 95 evidence cost rows	RAI lower bound 20.2 % → 28.6 %; condition programmes appear; optimiser RAI 0.661
4	33c4dad6	c / draft-d / UCSN + evidence	492 min (logged)	ferries, fords, culverts, road numbers, airports, markets, IWI wealth	RAI by wealth quintile: do-nothing 0.100 (poorest) to 0.521 (richest)
5	929bc009	c / draft-e / UCSN + evidence	56.1 min	envelope from approved budgets (roads share 0.545); four optimiser specifications and a frontier	OPT RAI 0.657; every rule-based scenario leaves 7.9–12.8 tn of the 25.8 tn envelope unallocated

Run	Run id	Geography / norms / prices	Module time	What changed	Headline result
6	b32f70f5	c / draft-f / UCSN + evidence 09b	19.5 min	evidence prices for 14 codes	41,757 projects appraised; OPT_RAI_COND50 treated 0 km against its floor: a pre-screen fault
7	—	c / draft-f / evidence 09b	—	candidate cap raised to 60,000	every solve hit the 900 s limit; stopped, empty log
8	014c9cd7	c / draft-f / UCSN + evidence 09b	87.4 min	optimiser fix	41,867 of 41,890 projects priced; RAI at 2030 under C 0.715; the run used in this report

Three things stand out from the history.

The evidence moved the results more than the code did. Between run 2 and run 8 the lower RAI bound rose from 20.2 % to 28.7 %, because 4,637 km gained a surface from HeiGIT and 4,848 km from GRIP4. The number of priced projects rose from 8,116 to 41,867, because the evidence prices filled codes that UCSN lacks. Scenario C's RAI at 2030 rose from 0.645 (run 5) to 0.715 (run 8) for the same reason: engineered-earth upgrades could now be funded.

Run time is dominated by the optimiser. Without the optimiser, a national run takes 8 to 15 minutes of module time (runs 1–3). TI-08 alone took 2,853 s in run 5 and 4,569 s in run 8, because it now solves four specifications, their attainable maxima and five frontier points. Run 4's logged 492 minutes is an outlier: its TI-01 and TI-03 times (1,915 s and 27,219 s) are many times those of every other run (TI-03 took 69–140 s elsewhere), and the log does not say why.

The faults were found by the engine's own checks. Run 6 reported that OPT_RAI_COND50 met 0 % of its condition floor. Because a shortfall is a reported finding, not a silent result, the cause was traced to the candidate pre-screen. It had filled its 14,000 places with access projects and dropped every condition programme. The fix merges the per-measure candidate lists in turn, raises the cap to 20,000 plus the warm-start projects, and warm-starts the attainable maximum from the same rule-based set. Run 7 tried a 60,000 cap and was stopped when every solve hit the 900 s limit.

18.5.1 Run 8 in brief

Run 8 is the run used throughout this report. Table 18.6 gives its headline results, each with the tier it should be read with.

Table 18.6: Run 8 (014c9cd7) headline results and their quality

Result	Value	Read with
Network	817,354 links, 247,972 km; surface known on 31.7 % of km	surface from OSM, HeiGIT and GRIP4; GIS-5 fails
RAI today (wet season)	28.7 % (unknown = earth) to 93.5 % (unknown = gravel)	benchmark 53.1 % (2015)
Programme	41,890 projects; 41,867 priced; 23 spot improvements unpriced	prices E5 evidence candidates
Capital, central (low–high)	92.8 tn UGX (56.6–125.3 tn), with 205.2 bn UGX a year of induced routine maintenance	draft norms; one code for paved upgrading on all terrain
Condition backlog, central	7.25 tn UGX across national paved, national unpaved, district and community networks	agency poor km by network, FY2022/23
RAI at 2030	do-nothing 0.287; A 0.374; B 0.698; C 0.715; D 0.446; E 0.465; F 0.707; G 0.408	horizon-year state only
Optimiser	OPT equals C (0.715); OPT_RAI_COND50 0.709 while treating 53,977 km; OPT_NPV_RAI50 0.474 with NPV 2,265 bn	OPT_RAI_EMONC80 and the frontier's 100 % point hit the time limit and miss their floors

Result	Value	Read with
Routine maintenance	requirement 1.02 tn UGX a year at steady state against 0.59 tn of eligible resources; funding ratio 0.58	scenario B; pool shares from approved budgets, not outturn

No scenario is marked recommended. The run log lists eight fallbacks. Among them: freight is grown at a flagged 5 % a year because sector output paths were not supplied; the vehicle operating cost coefficients are uncalibrated TransportModelling values; and two optimiser specifications do not meet their floors. Each of these is a statement about the data, and each is printed with the result it affects.

18.5.2 Later runs and simulations

The output folder holds further runs after run 8, launched from the web interface's run button, and a scripted set of seven one-at-a-time simulations (`scripts/run_simulations.jl`): the base case, prices at their low and high bounds, the roads share of the envelope at 70 % and 130 % of its evidence value, unknown surface treated as earth, and the UCSN-first price map. These are not recorded in `IMPLEMENTATION_STATUS.md`. Their results, where used, are reported in the outcomes chapter.

18.6 What is missing

Table 18.7 lists the items of the plan that are not built, or built but not working on Uganda data, with the reason and what would close each. It separates code work from data work, because the two have different owners.

Table 18.7: Missing items, by kind, with what would close each

Item	Kind	Why it is missing	What would close it
Per-link condition, and with it deterioration and rehabilitation triggers	data	no link-level survey exists in the estate	UNRA and DUCAR condition surveys (request letters 01 and 02, drafted, not sent)
Surface on about 68 % of network km (GIS-5)	data	open sources cover arterial roads only	road inventory from UNRA and DUCAR
Calibrated speeds	data	speeds are TransportModelling generic defaults, marked "NOT Ugandan"	speed survey or GPS probe data
Sector magnitudes (minerals, trade, agriculture below region)	data	minerals and trade exchanges are synthetic demonstration data and are not used	live channels from MineralModelling, TradeModelling and AgricultureModelling
Candidate alignments for missing links	data from another model	not supplied by SDI or TransportModelling	alignment output written to the exchange
PostGIS store	code and infrastructure	no credentials for a TIPRAS database	provision a database; port migrations 0001–0003
Project dependencies in allocation	code	not carried over from the TransportModelling portfolio model	add the dependency constraint and a `TST` test
FINEX inbound reader and the fiscal loop	code	not built	a FINEX reader and a convergence test
Outbound contract validation	code	`config/interfaces/` is empty	write the contracts of docs/21 §21.10 and validate against them
Year-by-year outcomes and the year slider	code	TI-09 stores the horizon only	recompute the network for each plan year
Outcome ranges	code	low and high costs are not carried into NPV and outcomes	corner runs (docs/23 §23.21)
Narrative reports, equity slider, authentication	code	interface phase partial	persist TI-07 objects per run; add reports and log-in

Item	Kind	Why it is missing	What would close it
Graph simplification	code	full noded graph is routed	degree-2 simplification with a link map
Approved norms	institutional	every norm is draft	review by MoWT, MoH, MoES and MAAIF (letter 05)

The table is weighted towards data. Of the 14 items, six are data or institutional and cannot be closed by writing code. The most consequential is per-link condition. Without it the engine cannot place rehabilitation on a road, cannot project deterioration, and must treat network condition as a set of national statistics.

18.7 Data quality and limitations

The status in this chapter is itself evidence of a particular kind, and it has limits.

- **The record is a status file, not an audit.** `IMPLEMENTATION_STATUS.md` was written by the build. Where this chapter states a gate result, it has been checked against the code, the test file and the run logs. Where the code could not show a result (for example the ten-minute user test), the chapter says it has not been run.
- **Gates that depend on data cannot pass yet.** GIS-5 fails because surface is unknown on 68.2 % of network km. The condition parts of phases 4 and 5 are idle. These failures are correct behaviour: the engine refuses to invent attributes.
- **The test record is older than the result.** The 2,265 passing assertions were recorded before the optimiser change that run 8 uses.
- **Run times are indicative.** They are module times from single runs on one machine with 8 threads. Run 4's time is unexplained.
- **The whole run is exploratory.** Norms are draft (578 of 646 rows *assumed*), prices are E5 evidence candidates, and the vehicle operating cost coefficients are uncalibrated.

18.8 Implications for planning

- The engine can already answer the five planning questions for Uganda end to end, in under 90 minutes with the optimiser and under 15 without it. Its answers are exploratory, and every output says so.
- The binding constraint on the build is now data, not code. Six of the 14 missing items need agency or model data, or approval. The drafted request letters are the next step on the critical path, and none has been sent.
- A link-level condition survey is the single largest gap. It would switch on the deterioration projection and the rehabilitation trigger, which are already coded and tested.
- The test suite should be re-run and its result recorded before run 8 is quoted as tested, because the optimiser changed after the last recorded test run.
- Two code gaps affect how results can be used in budgeting: there are no project dependencies in the allocation, and no FINEX reader to close the fiscal loop. Until both exist, allocations should be read as a comparison of scenarios, not as a budget submission.
- Moving the store to PostGIS would remove the 2 GB-per-run disk limit and allow shared use. It needs a provisioned database, not new design.
- The norms are the other half of the critical path. Moving them from *assumed* to *normed* or *observed* needs the sector ministries to confirm the access and design standards, which letter 05 asks them to do.

CHAPTER 19

19 The database

A plan that a ministry can defend must be able to answer three questions about any figure in it. Where did the number come from? Which inputs, norms and prices produced it? Can the same result be produced again? This chapter describes the store that holds the engine's norms, prices, results and audit trail, and how far it meets those needs for Uganda today. It implements Part A of design chapter 18 (`docs/18_DATABASE_SPECIFICATION.md`): the database infrastructure, migrations, roles and audit, offline mode, links to other stores, and backup and retention. Parts B to D of that chapter, the source datasets and their acquisition, are covered in the chapter on existing data assets. The short answer is that the Uganda build runs on a single SQLite file, not the PostgreSQL/PostGIS cluster the specification asks for. This is a deliberate and recorded deviation. The schema keeps the specification's structure and its main protections: approved norms cannot be edited, and the audit trail cannot be rewritten. But several of the audit tables are defined and still empty, and only one module writes factor traces.

19.1 What the specification asks for

Docs/18 §18.1 specifies one PostgreSQL instance dedicated to transport planning, following the estate's pattern of one cluster per domain. Table 19.1 summarises it.

Table 19.1: The database specified in docs/18 Part A

Item	Specification
Engine	PostgreSQL 16 or later, with `postgis` 3.4+, `pgcrypto`, `btree_gist` and `pg_trgm`
Not installed	`postgis_raster` and `pgrouting`: rasters stay in SDI, routing runs in Julia
Port and database	a dedicated port (for example 5437, beside SDI on 5433), database `tipras`
Schemas	`tpl_ref`, `tpl_norm`, `tpl_geo`, `tpl_obs`, `tpl_result`, `tpl_audit`, `tpl_stage`
Storage CRS	EPSG:4326 for geometry; measurement in the country's projected CRS
Sizing (Uganda)	120–270 GB working total for 20 retained runs; 1 TB SSD provisioned
Partitioning	`tpl_audit.factor_trace` and `tpl_result.gis_layer` by `run_id`
Roles	reader, analyst, norm custodian, loader, engine, admin; no role may update or delete audit rows
Row-level security	on `country`, so one ministry's data cannot leak to another
Backup	nightly base backup for 30 days, continuous WAL archiving, quarterly restore test
Offline mode	SQLite plus GeoPackage for district offices, with a 50 km routing buffer

The specification estimates that the factor trace alone would reach 60 to 150 GB for 20 runs, and results 40 to 80 GB. These estimates explain why it asks for a server database, partitions and a retention policy rather than a single file.

19.2 What is built: an SQLite store

The Uganda build uses the specification's **offline mode** as its main store. The deviation is recorded in `IMPLEMENTATION_STATUS.md`, as Table 19.2 shows.

Table 19.2: The database deviation as recorded

Specification	Now	Why, and next step
PostgreSQL + PostGIS primary store	SQLite offline store (the specification's offline mode) plus Arrow geography snapshots	The local PostgreSQL (port 5432) needs credentials that were not available. SDI's cluster (port 5433) was used read-only to export the network. The migrations are to be ported to PostGIS when a TIPRAS database is provisioned

The store is one file, `outputs/tipras.sqlite` by default. Every CLI command can point to another file with `--store`, and each simulation keeps its own store under `outputs/simulations/`. The golden test country `TST` uses `outputs/tipras_tst.sqlite`.

The data layer is `src/Data/DB.jl`. Its `open_store` function:

1. creates the file if needed;
2. turns on foreign keys;
3. sets the journal to write-ahead logging (WAL), so the web interface can read while a run writes;
4. applies any pending migrations.

The specification's seven schemas become table-name prefixes, because SQLite has no schemas. Table 19.3 maps one to the other.

Table 19.3: Specified schemas and their form in the SQLite store

Specified schema	SQLite form	Tables	Status
`tpl_ref` (catalogues)	`ref_`	none	not built; vocabularies live in norm keys and code constants
`tpl_norm`	`norm_`	`norm_version`, `norm_value`	built; long format, one table for all 28 norm tables
`tpl_geo`	`geo_` plus Arrow files	`geo_version` (registry)	registry table empty; geometry in `data/geo/<ISO3>/<geo_version>/`
`tpl_obs`	`obs_`	`obs_source`, `obs_envelope`, `obs_price`	prices loaded; sources and envelope empty
`tpl_result`	`res_`	seven result tables	built and written by every run
`tpl_audit`	`audit_`	`audit_run_manifest`, `audit_run_event`, `audit_factor_trace`, `audit_exchange_receipt`	manifest and factor trace written; events and receipts empty
`tpl_stage`	`stage_`	none	not built; raw sources are staged as files

19.3 Migrations

The schema is built by three forward-only migration files in `database/migrations/sqlite/`. `migrate!` applies every file named `NNNN_*.sql` that is not yet recorded in a `schema_migrations` table, in order, each inside one transaction. It never edits or reverses one. A small parser keeps trigger bodies (`BEGIN … END;`) whole when it splits a file into statements. All three migrations were applied to the Uganda store on 2026-09-23. Table 19.4 lists what each creates.

Table 19.4: Migration files

File	Content
`0001_schemas_audit.sql`	the audit tier: run manifest, run events, factor trace, exchange receipts; four append-only triggers

File	Content
0002_norms.sql	the norm tier: norm_version with its status, norm_value in long format, a lookup index, three immutability triggers
0003_ref_obs_results.sql	the geography registry, observations (sources, envelope, prices) and the seven result tables

The store therefore holds 16 working tables plus schema_migrations, seven triggers and one explicit index (norm_value_lookup). Docs/18 §18.2 reserves numbering ranges by schema, for example 0100–0199 for norms and 0400–0499 for results, in blocks of ten per module. The SQLite set uses a compact 0001–0003 numbering instead. Porting to PostGIS is the natural point to adopt the specified ranges, because the PostGIS migrations will be a new series in any case.

Docs/18 §18.2 also says catalogue seeds are loaded by tipras seed from database/seed/, with a checksum recorded in the audit tier. That folder exists but is empty, and there is no seed command.

19.4 The audit tier

The audit tier is what makes a figure traceable. It has four tables.

19.4.1 The run manifest

audit_run_manifest holds one row per run. It is written when a run starts and completed when it finishes. It pins everything the run depended on:

- country, base year and horizon;
- geography version, norm version and price vintage;
- target set, scenario, envelope version and season;
- whether the run is exploratory;
- code version, start and finish times;
- the list of fallbacks the run used, and the gate results, both as JSON.

On 5 October 2026 the Uganda store held 13 manifests. Ten runs had finished; three had not. Two of the unfinished runs match those IMPLEMENTATION_STATUS.md records as stopped or killed: run 7, stopped when every optimiser solve hit its time limit, and a web-launched run killed by a server restart at TI-05. The third, on norms draft-c, has no recorded cause. All 13 are flagged exploratory, because every norm version is still in draft. Run 014c9cd7, the run behind this report, pins geography UGA-2026.09-c, norms UGA-2026.09-draft-f and prices UCSN+EV-UGA-2026.09b, and its manifest carries the eight fallbacks listed in the run summary.

Three weaknesses should be noted:

- The envelope_version field is empty on all 13 manifests. The envelope is read from the MTEFModelling ceiling file at run time, not from a pinned version in obs_envelope. If that file is re-estimated, the run cannot be reproduced exactly from its manifest.
- The gate_results field holds an empty object on every finished run. Gate failures stop a run, so a finished run has passed its gates, but the passes themselves are not recorded.
- The manifest is updated in place when the run finishes. No trigger protects it. Docs/18 §18.3 asks that the whole audit schema be append-only.

19.4.2 Run events and exchange receipts

audit_run_event is meant to log events within a run, and audit_exchange_receipt to record every file exchanged with another engine, with its path, SHA-256 checksum and whether it was a snapshot or a fallback. Both tables have append-only triggers. Neither is written by the code today, and both were empty on 5 October 2026. Run

events go instead to `outputs/<run_id>/run_log.txt`. Exchange files are written under `data/exchange/outbound/` without a receipt row. Until receipts are written, the store cannot prove which version of a sector feed a run consumed.

19.4.3 The factor trace and explain()

`audit_factor_trace` is the backbone of the provenance design (docs/02 §2.7). Each row says that one result cell depends on one factor. It records:

- the run and the cell;
- the factor's name, value and unit;
- a `source_ref`, which is a norm (`norm:<table>:<key>@<scope>`), an observation (`obs:…`), a geography record (`geo:…`), an exchange file (`xchg:…`) or another result cell;
- the factor's quality tier and basis.

The primary key is run, cell and factor. CHECK constraints limit quality to measured, normed, imputed or assumed, and basis to observed, estimated, modelled or assumed. Two triggers refuse any update or delete. The test suite checks this: an attempt to delete from or update the table must fail.

`explain(db, run_id, cell)` walks the trace. It lists every factor of the cell, then follows any `source_ref` that points to another result cell, to a depth of six levels. It stops at norms, observations, geography and exchange files, which are leaves. A companion function returns the cell's **inherited quality**: the worst tier among all its leaves. The web interface serves this as `/api/explain?run&cell`, and the CLI as `tipras explain --run <id> --cell <ref>`. The golden-country test checks that a raked population cell decomposes into its factors.

The design is complete, but the coverage is not. In the current code only TI-01 writes factor traces. For each settlement cell and year it records two factors: the gridded population and the raking factor that scales it to the UBOS district total. Each Uganda run writes 701,468 trace rows: 350,734 cell-years with two factors each. Across the 13 runs the table held 9,119,084 rows. Gaps, costs, allocations and outcomes are not traced. For those results, `explain()` returns nothing. Their provenance is carried instead by the quality tag on each result row, the manifest pins, the fallbacks and the CSV tables.

One inconsistency also shows in the trace. The raking factor is stored with quality `normed` and basis `estimated`. Docs/16 §16.13 allows `estimated` only with `measured` or `imputed`. The code checks that each tag is a valid value. It does not check that the pair is valid, which is a load error under the specification.

19.5 The norm tier

`norm_version` holds one row per version, with a status of draft, proposed, approved, superseded or rejected, and who approved it and when. `norm_value` holds every norm row in long format, keyed by version, table, key, scope and field. A value is stored as a number or as text, never neither. Each row has a source, a source year and a quality tier (CHECK-constrained).

Three triggers make an approved version immutable: they refuse any insert, update or delete of its rows. `tipras norms-import` loads a folder of CSVs as a new **draft** version and refuses to overwrite an existing one. `tipras norms-approve --version <v> --by <name>` moves a draft or proposed version to approved. From then on it can only be replaced by a new version. Gate G7 in the engine refuses a non-exploratory run on norms that are not approved. The test suite checks that resolution picks the most specific scope, that a missing norm raises `NormNotFound`, and that an approved version cannot be updated.

The Uganda store holds six norm versions, from `UGA-2026.09-draft` (623 rows) to `UGA-2026.09-draft-f` (646 rows). All six are drafts; none has been approved. The catalogues chapter describes their contents.

19.6 Observations

`obs_price` holds unit prices by vintage, code, geographic level and year, with low, central and high values, the evidence tier and the status. `tipras prices-load` reads a UCSN-format CSV and loads its transport rows. Prices are resolved only through the norm unit-cost map. A duplicate code within a vintage raises an error. A missing price raises `PriceNotFound`, so the item is reported as unpriced and never costed at zero. An ILLUSTRATIVE or EVIDENCE_CANDIDATE price is refused in a non-exploratory run. The store holds the three vintages in Table 19.5.

Table 19.5: Price vintages in the store

Vintage	Codes	Content
`UCSN-UGA-2026.09-illustrative`	7	UCSN transport coefficients, tier E7 ILLUSTRATIVE
`UCSN+EV-UGA-2026.09`	14	first evidence set (runs 3 to 5)
`UCSN+EV-UGA-2026.09b`	21	the 7 UCSN coefficients and the 14 E5 evidence-candidate codes (run 8)

`obs_source` (the source register) and `obs_envelope` (pool amounts by year) are defined but empty. Sources are registered in `data/raw/SOURCES.md` and `data/reference/`, not in the store. The envelope is computed at run time from the MTEF ceiling file and the pool-share norms.

19.6.1 Geography outside the database

The specification puts links, vertices, settlements and nodes in `tpl_geo` with PostGIS geometry. The build keeps them in Apache Arrow files, one folder per geography version, for example `data/geo/UGA/UGA-2026.09-c/`. That folder is 343 MB. It holds links (168 MB), link coordinates (143 MB), vertices, settlements, nodes, districts, district adjacency, population projections, agricultural cells and a `manifest.json`. The manifest records the projected CRS (EPSG:32636), the link identifier scheme, the source paths, the parameters and caveats of the build, and a content hash of the links. Docs/20 §20.3 allows this file form, and it loads far faster into the Julia graph than a database query would. But the `geo_version` registry table that should point to each snapshot is empty. A run checks that the snapshot's manifest matches the version it pins (gate G3), and fails if it does not.

19.7 Result tables

Seven `res_` tables hold a run's results (Table 19.6). Each has `run_id` in its primary key, and each is written once per run.

Table 19.6: Result tables

Table	Grain	Notes
`res_gis_layer`	run × module × layer × feature × year × scenario	every mapped value, with unit and quality; feature kind constrained to link, structure, node, settlement, cell, hex or district
`res_access`	run × settlement × destination class × year × season × state	travel time, distance, whether the standard is met, population in need
`res_gap`	run × dimension × object × year × measure	dimension constrained to the seven gap dimensions G1 to G7
`res_candidate`	run × project	works, km, low/central/high cost, induced maintenance per year, appraisal, access gain, sector tags
`res_allocation`	run × scenario × project × pool × start year	amount by pool
`res_outcome`	run × scenario × indicator × geography × year	low/central/high
`res_budget`	run × scenario × view × code × fiscal year × pool	view constrained to programme, economic or administrative

Two rules of the engine are written into these tables. First, `res_gap` accepts only the seven named dimensions, `G1_access` to `G7_resilience`, and has no column that could hold a combined score. The test suite also checks that no gap-matrix column is named like an index, score or composite. Second, `res_candidate` stores a missing cost as NULL, not zero. A NULL cost marks an unpriced item, which is reported and cannot be funded. Each candidate also carries its induced maintenance per year beside its capital cost.

On 5 October 2026 the store held results for two runs, `e37494c6` and `dd723ea3`. Both were run on norms draft-e. Each had 2,616,172 GIS-layer rows, 1,505,141 gap rows, 1,227,569 access rows, 44,063 allocation rows, 41,819 candidates, 773 budget rows and 167 outcome rows. The results of run `014c9cd7` have been pruned from the store. Its manifest and factor traces remain, and its output tables remain on disk. This report reads that run's results from `outputs/014c9cd7-024b-4b62-8a40-e72bfe960683/tables/`.

19.8 Run storage under outputs/

Each run also writes a folder `outputs/<run_id>/`, which holds:

- `tables/`: CSV tables per module, 21 files for run `014c9cd7`, from `ti01_raking_warnings.csv` to `ti10_programme_budget_B.csv` and `unpriced_codes.csv`;
- `gis/`: the GIS summary and the LISA cluster table (and, where rendered, the PDF atlas per module);
- `run_log.txt`: the module timings, messages and fallbacks.

The folder for run `014c9cd7` is 27 MB. Runs launched from the web interface also leave their specification, log and status under `outputs/jobs/`. The bulk of the storage is in the SQLite file, not in the run folders. `IMPLEMENTATION_STATUS.md` puts a national run at about 2.5 GB in the store, after settlement- and link-level duplicate map layers were moved to district level. `CLAUDE.md` gives about 2 GB. On 5 October 2026 the store file was about 4.6 GB, with a WAL file of about the same size beside it. Drive D:, which holds the store, is nearly full.

19.9 The prune command

Because each national run adds gigabytes and the disk is nearly full, the build has a retention command:

```
julia --pkgimages=no --project bin/tipras.jl prune --keep 1
```

It works as follows:

1. It lists the finished runs in the manifest, newest first, and keeps the most recent N (default 1).
2. For every other run it deletes the rows in all seven `res_` tables.
3. It never deletes audit rows. A pruned run's manifest and factor traces remain, so the run can be identified and regenerated from the same pins.
4. It checkpoints the WAL and runs `VACUUM` to return the space to the disk.
5. It reports how many result rows it removed and which runs it kept.

Three guards surround it:

- A run started with less than 6 GB free on the store's disk records a fallback that tells the user to prune.
- The web interface refuses to start a run with less than 4 GB free unless the user chooses "prune to N runs" in the run dialog. The run job then prunes before it starts.
- `run.ps1 prune` from the repository root does the same from Windows.

Pruning is the practical version of the specification's retention rule (docs/18 §18.1), under which traces of superseded runs may be purged after 180 days but never those of an approved run. The current command goes the other way: it keeps traces and deletes results, which suits a store where traces are small and results are large. It does not delete the `outputs/<run_id>/` folders. Those hold the CSV tables and atlases of pruned runs and are

the only remaining copy of their results. Nor does it check whether a run is approved, which is safe today only because no run is.

Two small gaps sit in the CLI. Its usage line lists `validate`, which is not implemented, and it omits `prune` and `explain`, which are.

19.10 How the store is PostGIS-ready

The SQLite store was built so the move to PostgreSQL/PostGIS changes the storage, not the engine. Docs/18 §18.4 asks for the data layer to sit behind `Data/DB.jl` from the start, so both backends can pass the same tests. The build follows that in these ways:

- **One access layer.** The Julia modules and the web interface reach the store only through `open_store`, `migrate!`, `query`, `execute!`, `write_table!` and the norm, price and trace functions. The one exception is the Python atlas script, `scripts/render_maps.py`, which reads the SQLite file directly and would need a new connection string.
- **Schema by prefix.** Each prefix maps one-to-one to a specified schema (`norm_` → `tpl_norm`, `res_` → `tpl_result`, `audit_` → `tpl_audit`).
- **Portable SQL.** Primary keys, CHECK constraints and triggers have direct PostgreSQL equivalents. The triggers become PL/pgSQL functions, `datetime('now')` becomes `now()`, and `REAL` becomes `double precision`.
- **UUID run identifiers.** Runs are UUIDs, so runs made offline can be uploaded to a central store without collisions, as §18.4 requires.
- **Geometry by key.** Results refer to features by key (link, settlement, node, district), not by geometry. Moving the Arrow snapshots into `tpl_geo` with GiST indexes would not change any result table.

A port would still need the following work:

- PostGIS migrations in the specified number ranges, including the missing `tpl_ref` and `tpl_stage` schemas;
- the six roles and their grants, with no UPDATE, DELETE or TRUNCATE on `tpl_audit` for any role, including the owner;
- row-level security on `country`;
- LIST partitioning of the factor trace and GIS layers by `run_id`;
- a backend switch in `open_store` and a second run of the test suite against PostgreSQL;
- a nightly backup, WAL archiving and a restore test.

Until then, the SQLite triggers are the only protection for the audit tier. Any user who can write the file can also delete it.

19.11 What is built, tested and missing

Table 19.7 sets the build against each requirement of docs/18 Part A.

Table 19.7: Database status against docs/18 Part A

Requirement	Status
Dedicated PostgreSQL/PostGIS cluster	missing: SQLite offline store used instead (recorded deviation)
Forward-only migrations	built: three SQLite migrations, applied transactionally and recorded
Seven schemas	partial: norm, observation, result and audit tiers built as prefixes; `ref` and `stage` missing; geography as Arrow files
Approved norms immutable	built and tested (three triggers)

Requirement	Status
Audit append-only	built and tested for the factor trace and run events; the manifest is updated in place
Factor trace and `explain()`	built and tested; written only by TI-01 population raking
Run events, exchange receipts	tables built, not written
Source register, envelope versions, geography registry	tables built, empty
Roles and row-level security	missing (not possible in SQLite)
Partitioning, backup, restore test	missing; retention handled by `prune`
Offline mode with routing buffer	national runs use the whole network, so no buffer is needed today; district extracts for offline offices not built
Seed load of catalogues	missing (`database/seed/` empty)

The test suite (2,265 assertions in `test/runtests.jl`) covers the store's protections directly. It checks that an approved norm cannot be updated, that the factor trace cannot be deleted or updated, that a missing norm raises rather than defaults, and that `explain()` decomposes a population cell in the golden country.

19.12 Data quality and limitations

- **Single file, no access control.** The audit trail is protected by triggers only. There are no roles, no row-level security, no backups and no restore test.
- **Thin provenance.** Only population raking is traced. For gaps, costs, allocations and outcomes, provenance rests on the quality tags on result rows, the manifest and the run log. `explain()` cannot decompose them.
- **Reproducibility gaps.** The envelope version is not pinned, exchange files are not receipted, and the source register is not in the store. A run can be re-run from its manifest only if the MTEF file and sector feeds have not changed.
- **Pruned runs.** Results of pruned runs, including run `014c9cd7`, survive only as CSV tables and atlases under `outputs/`. They can be regenerated from the manifest pins, subject to the gaps above.
- **Quality-basis pairs are not validated.** The raking factor is stored as normed and estimated, a pair the specification does not allow.
- **Disk pressure.** The store and its WAL took about 9 GB on a nearly full drive on 5 October 2026. Without regular pruning, new runs will fail or warn.

19.13 Implications for planning

- Provisioning a TIPRAS PostgreSQL/PostGIS database, with credentials, is the precondition for any multi-user or ministry deployment. The migrations and data layer are ready to be ported, but roles, row-level security and backups cannot exist in a single SQLite file.
- Before any result is presented as more than exploratory, factor tracing should be extended beyond TI-01 to gaps, unit costs, allocations and outcomes. Then every published figure can be decomposed with `explain()`.
- The envelope should be loaded into `obs_envelope` as a pinned version, and every exchange file should get a receipt with its checksum. Then a run's manifest is enough to reproduce it after counterpart models are re-estimated.
- The manifest should be made append-only, with completion recorded as an event, and gate results should be stored. Then the audit trail shows that each gate passed, not only that the run finished.
- Retention should be agreed as policy, not left to disk space. A run that supports a submitted plan should be exported as a signed bundle (docs/18 §18.6) and exempt from pruning. `prune` should refuse to touch an approved run.

- Until the store moves to a server, each national run should be preceded by `prune --keep 1`. The `outputs/<run_id>/tables/` folders should be backed up, because for pruned runs they are the only copy of the results.
- The quality-basis pair rule of docs/16 §16.13 should be enforced when traces and observations are written, so a mislabelled tier is caught at load time rather than in review.

CHAPTER 20

20 Existing data assets and acquisition

What data does the engine have for Uganda, where did it come from, and what is still missing? Every number in the earlier chapters rests on the answer. This chapter implements design chapter 19 (`docs/19_EXISTING_DATA_ASSETS.md`). It describes the six parts of the EDIP estate that hold transport-relevant data, the 60 sources that were assessed for the Uganda run, how the acquisitions changed the results, and the data that only government agencies can supply. Its main finding is simple. The estate and the open sources supply geometry, population and a routable network, and they supply them well. They supply almost none of the road *attributes* that the planning questions turn on: surface beyond the arterial roads, condition, traffic, standards and calibrated costs. That is why every Uganda run is exploratory. The records behind the chapter are `data/reference/ESTATE_INVENTORY.md`, `data/reference/TRANSPORT_INFRA_DATA_GAP_ASSESSMENT.md`, `data/raw/SOURCES.md`, `data/raw/MANIFEST.csv` and the draft request letters in `docs/requests/`.

20.1 The estate at a glance

The EDIP estate is the set of data repositories and models that AICTEM maintains on one machine. Six parts of it hold transport data. Table 20.1 summarises what each holds for Uganda, as verified on disk on 23 September 2026.

Table 20.1: The six parts of the estate and their Uganda state (verified 2026-09-23)

Estate part	What it is	Size on disk	Uganda state
ATDGR (Africa transport data repository)	schema, vocabularies, source registry and raw acquisitions	7.4 GB	**nothing canonical**: the Uganda exports carry 0 rows
SDI (spatial data infrastructure)	PostGIS backbone: boundaries, population, noded roads, health access	4.0 GB	loaded and routable; the database was not running
AID (Africa Infrastructure Database, Krantz 2024) and the Uganda OSM extract	continental infrastructure lines and points, wealth surfaces	3.6 GB	complete; Uganda subset extracted to CSV
TransportModelling	tested Julia engine with the Uganda network, DEM and UNPS extracts	2.1 GB	methods implemented; publishes no exchange files
UCSN (unit-cost system)	unit-cost coefficients, budget crosswalk, evidence register	1.3 GB (data)	7 transport coefficients, all ILLUSTRATIVE (tier E7)
Sector models	minerals, trade, agriculture, health, education and MTEF exchanges	—	health registers real; minerals and trade exchanges synthetic

The pattern is the same as for the health engine (HPRAS): the estate removes the geospatial work, and none of the institutional work. Design chapter 19 put the scale of the problem in four numbers measured by TransportModelling on its own Uganda network: speeds are assumed on 96.1 % of network km, condition is assumed on 96.9 % of links, no link carries a traffic count, and every transport unit cost is illustrative. The work recorded in this chapter has moved two of those numbers, surface and cost. It has not moved condition or traffic.

20.1.1 ATDGR: a framework with no Uganda records

ATDGR (version 0.1.0) is a complete framework: a 41-file database schema, 22 vocabularies, a quality engine and an export writer. For Uganda it holds no data. Its canonical and staging folders are empty, and its three Uganda

exports (CSV, GeoPackage and GeoParquet, dated 8 August 2026) each report 0 layers and 0 rows. Its availability file registers 34 sources for Uganda: 16 downloadable and 18 blocked or needing manual acquisition. Every Uganda agency source, including UNRA road condition, UNRA traffic counts, the MoWT project pipeline, URA border posts and URC rail, is registered with access by e-mail request, an unknown licence and no web address. ATDGR's 14 blank import templates (road segment, bridge, traffic count, border crossing and others) are useful as ready-made forms for the data requests. Its geoBoundaries files are Git LFS pointer stubs of about 130 bytes and cannot be read.

20.1.2 SDI: the spatial backbone

SDI supplies the geography that every module uses. Its Uganda content is:

- **Boundaries:** 4 regions, 135 districts, 208 counties and 1,520 sub-counties from the OCHA/UBOS administrative boundaries, keyed by P-code.
- **Population:** WorldPop 2020 constrained at 100 m, totalling 42.17 million, which the engine rakes to the UBOS 2022 district totals.
- **Settlements:** OSM places (74,927 points in the Geofabrik places layer, of which only 2,178 carry a population), so settlement population comes from WorldPop.
- **Network:** the noded road network, with 817,354 edges and 695,307 nodes. This is the network the engine routes on.
- **Hazards:** JRC flood depth for return periods of 10 to 500 years, CHIRPS monthly rainfall, HydroRIVERS and HydroBASINS.

SDI has defects that the engine carries forward. Surface is `unknown` on every SDI road, because the Geofabrik shapefile it was built from carries no surface tag. Its road class is OSM's, not Uganda's: the "national" class covers 4,398 km against about 21,000 km of gazetted national roads. Its noding logs stop at districts 13 and 15 even though the final graph reports all 817,354 edges. Its database runs as a user cluster on port 5433 and was not running on 23 September 2026, and no file copy of the noded network existed before the engine exported one into its own geography snapshot.

20.1.3 AID and the OpenStreetMap extract

The Africa Infrastructure Database is stored in R's `.qs` format. TransportModelling had already extracted the Uganda subset to CSV, and the engine reads those CSVs, so it carries no R dependency. The AID Uganda extracts include 419,789 points of interest, 25,273 wealth-index (IWI) points and 25,360 relative-wealth (RWI) points, and a population and travel-time hexagon grid of 22,785 cells. The AID classified road layer is not routed on: it covers only tertiary roads and above (about 38,200 km), and routing on it would overstate travel times. The Geofabrik Uganda OSM file (371 MB, dated 6 August 2026) is the source of surface, smoothness, bridge, ford, ferry and road-number tags.

The `.qs` aggregate files (10 km hexagons, neighbour lists) could not be read on this machine. The installed R 4.6 cannot load the `qs` package built for R 4.5, and a rebuild from source failed. Those files are described from their documentation only.

20.1.4 TransportModelling: code and seeds, no results

TransportModelling is a tested engine (3,536 assertions at the time of the survey) with Uganda inputs that TIPRAS uses: the full OSM road tag table (404,529 ways, surface tagged on 22.0 %), the Copernicus 30 m elevation model (49 tiles, 1.71 GB) and the UNPS 2019/20 community road-condition module. It publishes nothing to an exchange: there is no origin–destination matrix, no deterioration-parameter file and no calibrated speed table. TIPRAS takes code, configuration seeds and measured findings from it, as the chapter on TransportModelling reuse describes.

One finding from TransportModelling shaped the engine directly. Moving from the classified AID network to the full OSM network removed 56 % of the excess travel time over the Nelson (2019) benchmark and cut the mean

access leg from 45.9 to 16.1 minutes. The engine therefore routes on the full noded network and never on a classified subset.

20.1.5 UCSN: seven illustrative coefficients

UCSN publishes seven transport coefficients for Uganda, all at tier E7 (ILLUSTRATIVE): paved two-lane construction on flat, rolling and mountainous terrain (central 3.1, 3.7 and 5.2 bn UGX per km), gravel construction (0.5 bn per km), bridges (180 m UGX per metre of span), and routine maintenance of paved (18 m UGX per km per year) and gravel roads (10 m UGX per km per year). It has no code for periodic maintenance, rehabilitation, engineered-earth roads, culverts, drifts or footbridges. Its budget crosswalk also maps rail and water transport onto the paved-road coefficient, which TI-10 must not inherit.

20.1.6 Sector models

The sector models are uneven. HealthModellingx holds 3,169 geocoded MoH facilities, of which 2,416 (76 %) have coordinates to two decimal places or fewer (about 1.1 km precision). EducationModellingx holds UBOS 2022 population for 135 districts but has no outbound channel for transport. AgricultureModelling holds production for 4 regions only, and 568 OSM market points. The minerals exchange (115 rows, 65.3 Mt) is marked "DEMONSTRATION DATA — synthetic values" and the trade exchange is synthetic and for The Gambia. The engine does not use either: mines and special economic zones enter as nodes without tonnages. MTEFModelling supplies the transport ceiling that TI-08 uses as its envelope.

20.2 Sources acquired for the Uganda run

The estate survey showed what was missing. The acquisition work of 23–24 September 2026 then looked for each missing item in open sources, before any request to an agency. Every source assessed is recorded in `data/raw/MANIFEST.csv` with its web address, date, files, size, checksum, licence, citation and status. Where the estate already held a dataset, the manifest records its path and nothing was downloaded again.

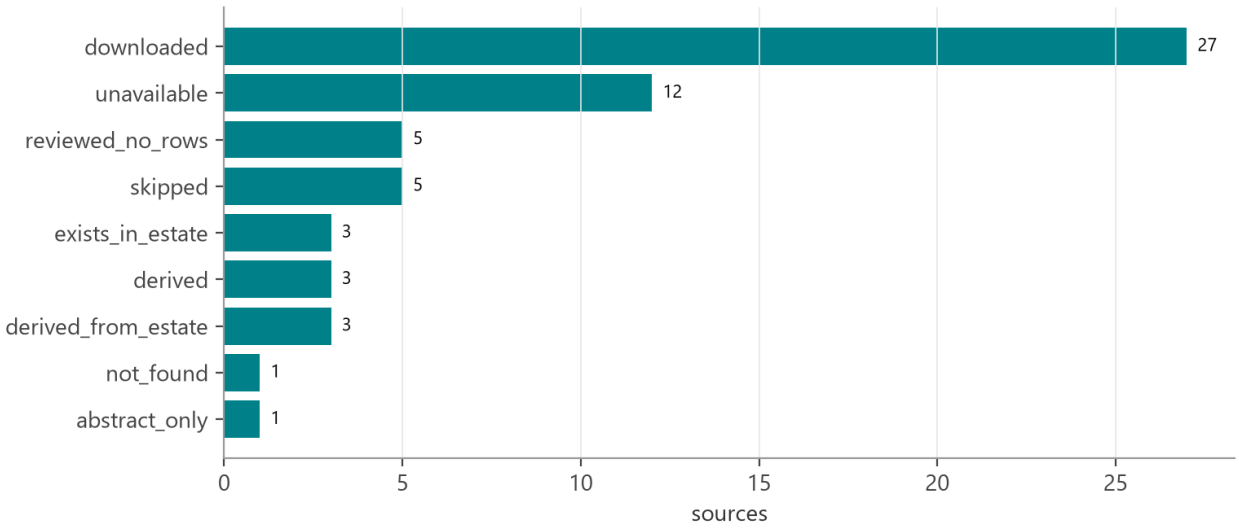

Figure 20.1: Sources assessed for the engine, by acquisition status

Source: data/raw/MANIFEST.csv

Figure 20.1 counts the 60 sources assessed by their acquisition status. Twenty-seven (45 %) were downloaded from the original open source. Twelve (20 %) were unavailable: the site refused the connection, demanded a browser challenge, returned an access error or had withdrawn the dataset. Five were skipped because the estate already held the same data or the file exceeded the 400 MB single-file limit, and five more were downloaded, read and found to contain no usable rows, so they were not kept. Nine came from the estate: three already present and used in place, three small Uganda extracts cut from estate files, and three tables derived from the downloaded

documents (the two rounds of extracted unit costs and the extracted condition table). One source yielded only an abstract and one was not found at all. Together, the files recorded in the manifest total about 643 MB. The largest are GRIP4 (277 MB), the MoWT reports (108 MB) and the JRC flood tiles (103 MB).

The figure shows that a search of open sources was worth doing: the downloaded documents alone yielded 275 unit-cost rows and 203 condition rows that the estate did not hold. It also shows its limit. A fifth of the sources could not be reached by a script, and the ones that matter most for condition and traffic are not open at all.

Table 20.2: Sources acquired for the Uganda run, by status

Source	Content	Status	Licence
jrc_flood_hazard_v2	JRC/CEMS-GloFAS Global river flood hazard maps v2.1.2, RP100 depth (3 arcsec ~90 m), 4 tiles intersecting UGA	downloaded	CC-BY-4.0 (European Union 1995-2026, see copyright.txt)
nasa_landslide_susceptibility	NASA Global Landslide Susceptibility Map (1 km, 2023-02-27 release), clipped to UGA bbox 29.5,-1.5,35.1,4.3	downloaded	NASA open data (US Government work, no copyright; attribution requested)
mapspam_2010	MapSPAM 2010 v2r0 physical production (P), all technologies (_A), 42 crops, 5 arcmin, clipped to UGA bbox	downloaded	CC-BY-4.0 (IFPRI Dataverse terms of use)
grip4_africa	GRIP4 Global Roads Inventory, Region 3 (Africa) vector shapefile + UGA bbox extract (GPKG, 17,317 segments)	downloaded	CC0 (per globio.info download page)
sdg_9_1_1_rai	SDG 9.1.1 Rural Access Index (SP_ROD_R2KM, % rural pop within 2 km of all-season road), Uganda	downloaded	UN SDG Global Database terms (free reuse with attribution)
hydrorivers_uga	HydroRIVERS v1.0 Africa, clipped to UGA bbox (19,390 reaches)	derived_from_estate	CC-BY-4.0 (HydroSHEDS licence)
ourairports_uga	OurAirports airports.csv + runways.csv filtered to iso_country=UG (42 aerodromes, 15 runways)	derived_from_estate	Public Domain (PDDL-1.0 as recorded by ATDGR)
osm_border_crossings	OSM border crossings / customs points for Uganda (barrier=border_control, government=customs), 32 points	derived_from_estate	ODbL-1.0
osm_geofabrik_uga	Geofabrik Uganda OSM PBF (roads, bridges, fords, ferries, railways, aerodromes, amenities, waterways)	exists_in_estate	ODbL-1.0
admin_boundaries_uga	Uganda admin boundaries ADM0-ADM4 (COD-AB) and geoBoundaries gbOpen ADM0-ADM2	exists_in_estate	COD-AB: HDX licence as published (verify); geoBoundaries: CC-BY-4.0
worldpop_uga_2020	WorldPop Uganda 2020 constrained 100 m population	exists_in_estate	CC-BY-4.0
healthsites_uga	healthsites.io Uganda export	skipped	ODbL-1.0
mapspam_2020	MapSPAM 2020 v2r2 global production	unavailable	CC-BY-4.0
wb_rai_dataset	World Bank Rural Access Index dataset (subnational)	unavailable	CC-BY-4.0 (WB)
heigit_planet_road_surface_uga	HeiGIT Uganda Planet Road Surface, Width and Passability (AI from PlanetScope 2020/2024) on OSM arterial roads	downloaded	Other: non-commercial use only; contains derivatives of PlanetScope imagery; redistribution subject to Planet Labs PBC terms (per HDX)
ocha_uganda_roads_2009	OCHA ROSEA Uganda roads Feb 2009 shapefile (2,214 segments; northern Uganda focus; road type/status, rehabilit	downloaded	HDX legacy hrinfo terms (https://data.humdata.org/about/license/legacy_hrinfo)
wb_rocks	World Bank ROCKS Road Costs Knowledge System v2.3 (full package zip) + Uganda extract (202 records)	downloaded	CC BY 4.0 (World Bank Data Catalog)

Source	Content	Status	Licence
wb_project_docs_uganda_roads	World Bank Uganda road project documents: TSDP PAD 2009 and ICR 2017 (P092837), RDP2 ICR 2009 (P065436), ARSDP	downloaded	World Bank public disclosure documents
mofped_bmau_its	MoFPED BMAU Integrated Transport Infrastructure Services Annual Budget Monitoring Report FY2023/24 (106 pp)	downloaded	Government of Uganda public document (no explicit licence)
mowt_uganda_reports	MoWT Uganda: ITIS Programme Performance Reports FY2020/21, FY2021/22, FY2022/23; Annual Sector Performance Rep	downloaded	Government of Uganda public documents (no explicit licence)
urf_uganda_road_fund	Uganda Road Fund Budgeting and Operational Guidelines FY2018/19 (hosted on Amuria District LG website)	downloaded	Government of Uganda public document (no explicit licence)
lpcb_uganda_lowcost_seal_2013	Pinard and Overby (2013) Review of Low-cost Seal Technology Options on Low Volume Roads in Uganda (69 pp)	downloaded	Public document library (no explicit licence)
ubos_statistical_abstract	UBOS Annual Statistical Abstract 2024 (447 pp)	downloaded	UBOS public document (reuse with attribution)
obeti2024_unpaved_maintenance_cost	Obeti et al. (2024) Investigating the cost of mechanized unpaved road maintenance operations in Uganda - abstr	abstract_only	CC BY-NC (per OpenAlex); full text not downloaded (publisher HTTP 403 to scripted access)
unit_costs_extracted	Extracted Uganda road works unit costs (95 rows)	derived	As per each source row
road_condition_extracted	Extracted Uganda road condition and network length figures (203 rows)	derived	As per each source row
hotosm_uga_roads	HOT OSM Uganda roads export (GPKG 196 MB)	skipped	ODbL
heigit_mapillary_road_surface_uga	HeiGIT Uganda road surface from Mapillary street imagery (GPKG 475 MB)	skipped	ODbL
hdx_uganda_road_network_unra	HDX Uganda Road Network UNRA	unavailable	
urf_annual_reports	Uganda Road Fund annual reports / One-Year Road Maintenance Plans	unavailable	
afdb_uganda_road_docs	AfDB Uganda Road Sector Support Project appraisal; Road Maintenance and Upgrading PCR 2011	unavailable	
ifad_prelnor_supervision_2023	IFAD PRELNOR Supervision Report June 2023 (bridges/culverts on CARs)	unavailable	
mowt_ric_summaries	MoWT Research Information Centre summaries (low-cost seals, URF allocation)	unavailable	
wb_wp8746_rural_access_2019	World Bank Policy Research WP 8746: Assessing Rural Accessibility and Rural Roads Investment Needs Using Open	downloaded	CC BY 3.0 IGO (World Bank Policy Research Working Papers)

Source	Content	Status	Licence
afdb_road_costs_2014	AfDB (2014) Study on Road Infrastructure Costs: Analysis of Unit Costs and Cost Overruns of Road Infrastructur	downloaded	AfDB publication (no explicit open licence)
wb_beyond_the_gap_2019	Rozenberg & Fay (eds.) (2019) Beyond the Gap, Vol. 1 - World Bank plain-text rendition	downloaded	CC BY 3.0 IGO
wb_project_docs_eastafrica_roads	World Bank East Africa rural/feeder road project documents: Rwanda Feeder Roads Development Project PAD 2014 a	downloaded	World Bank public disclosure documents
afcap_geddes_2016_road_asset_mgmt	AfCAP GEN2018A (2016) Economic Growth through Effective Road Asset Management - Final Formulation Phase Report	downloaded	UK aid / ReCAP research output (no explicit licence)
afdb_uganda_road_appraisals_2015_2025	AfDB Uganda road appraisal reports: RSSP4 PAR (2012/13), RSSP5 PAR (May 2014), Kampala City Roads Rehabilitati	skipped	AfDB public documents
wb_aicd_uganda_infrastructure_2012	World Bank WPS5963 (2012) Uganda's Infrastructure: A Continental Perspective (AICD country report)	skipped	CC BY 3.0 IGO
jica_uganda_bridges	JICA / World Bank / ReCAP bridge cost per metre span for Uganda or East Africa	unavailable	
ifad_prelnor_design_2014	IFAD PRELNOR Detailed Design Report (main report + appendices, 828 pp), 15 Oct 2014 - CAR unit costs (Appendix	downloaded	IFAD document, public disclosure (no explicit licence)
afdb_caiip_uganda	AfDB CAIIP-1 results/PCR summary (appraisal assumptions for district road rehab and CAR/district RM) and CAIIP	downloaded	AfDB public disclosure (no explicit licence)
urf_wayback_2019_2023	Uganda Road Fund Budgeting & Operational Guidelines FY2019/20 and FY2020/21 (unit cost Tables 6a/6b) and URF A	downloaded	Government of Uganda public documents (no explicit licence)
mofped_bmau_roads_2019_2023	MoFPED BMAU Roads Sub-Sector Annual Monitoring Report FY2019/20, Roads Semi-Annual FY2020/21, ITIS Programme S	downloaded	Government of Uganda public documents (no explicit licence)
obeti2024_fulltext_kyu	Obeti et al. (2024) full text (Kyambogo University repository copy) - mechanised unpaved maintenance cost per	downloaded	CC BY-NC 4.0 (per publisher/OpenAlex)
dfid_tiup_cost_roads_africa_2007	DFID TI-UP Enquiry 'Cost of Roads in Africa' (Feb 2007), secondary compilation incl. Uganda rural road costs	downloaded	UK Open Government Licence (GOV.UK R4D)
wb_rwanda_feeder_roads	World Bank Rwanda Feeder Roads Development Project P126498: PAD 2014 and ICR 2026 (East African comparator for	downloaded	World Bank documents: CC BY 3.0 IGO

Source	Content	Status	Licence
ifad_prelnor_supervision_2023_retry	IFAD PRELNOR Supervision Report June 2023 (retry with full document UUID URL + Wayback)	unavailable	
ioe_ifad_caiip1_pcrv	IFAD IOE Project Completion Report Validation, CAIIP-1	reviewed_no_rows	
mofped_bmau_caiip_brief_2016	BMAU Briefing Paper 10/16 on CAIIP implementation challenges	reviewed_no_rows	
wb_nusaf3_icr_2021	World Bank NUSAF3 ICR (P149965) - LIPW community access roads	reviewed_no_rows	
wb_drdip_icr_2024	World Bank DRDIP ICR (P152822, Djibouti/Ethiopia/Uganda)	unavailable	
wb_tanzania_rise_pad_2021	World Bank Tanzania RISE PAD (P164920) - rural roads / community-based routine maintenance	reviewed_no_rows	
ilo_labour_based_uganda_2006	ILO (2006) Increased Application of Labour-Based Methods Through Appropriate Engineering Standards - Uganda Co	reviewed_no_rows	
research4cap_uganda	ReCAP/AfCAP Uganda low-volume road reports (spot improvement, ENS, LVR costs)	not_found	
mowt_ducar_manuals	MoWT DUCAR force-account guidelines 2017 / DUCAR maintenance manuals	unavailable	
afdb_caiip2_appraisal	AfDB CAIIP-2 Appraisal Report	unavailable	
wb_wdi_deflators	World Bank WDI: GDP deflator Uganda and USA (NY.GDP.DEFL.ZS), official exchange rate UGX/USD (PA.NUS.FCRF), 19	downloaded	CC-BY-4.0
unit_costs_extracted_2	Extracted unit costs round 2: rocks_region_unit_costs.csv (ROCKS Eastern & Southern Africa), international_uni	derived	As sources

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_sources.csv).

Table 20.2 lists each source with its content, status and licence. It falls into four groups.

1. **Hazard, agriculture and validation layers.** JRC flood hazard v2.1.2 at about 90 m (four tiles intersect Uganda), the NASA landslide susceptibility map clipped to Uganda, MapSPAM 2010 crop production, and the SDG 9.1.1 rural access index. The RAI benchmark is a single observation: 53.1 % in 2015. MapSPAM 2020 was unavailable because the repository requires a guestbook form, so the 2010 edition was used.
2. **Road surface and inventory layers.** HeiGIT 2024 road surface (paved or unpaved, predicted from PlanetScope imagery) covers 4,269 segments and 15,346 km of motorway, trunk, primary and secondary roads only. GRIP4 gives a road type and surface for 17,317 segments in the Uganda extract. The OCHA 2009 roads layer has 2,214 segments, mostly in the north, and no surface or condition field.
3. **Documents for unit costs and condition.** MoWT programme and sector performance reports, the UBOS Statistical Abstract 2024, World Bank project documents and ROCKS, AfDB studies, the Road Fund guidelines, PRELNOR and CAIIP documents, BMAU monitoring reports and several East African comparators. From them, 275 unit-cost rows and 203 condition and network-length rows were extracted, each with its page and a short quote.
4. **Price indices.** World Bank WDI GDP deflators for Uganda and the United States and the UGX/USD exchange rate, used to bring every cost to UGX 2026.

The licence column matters for publication. Most layers are open (CC BY 4.0, CC0, ODbL or public domain). HeiGIT road surface is **non-commercial only** and subject to Planet Labs terms, so it is kept for internal planning use. Government documents carry no explicit licence. OSM-derived link layers must stay under ODbL and carry the credit "© OpenStreetMap contributors".

20.2.1 Why sources were unavailable

The twelve unavailable sources failed for reasons that a person at a browser, or an agency, can overcome. Table 20.3 groups them.

Table 20.3: Why sources could not be acquired by script

Reason	Sources	What would work
Browser challenge or bot check (Cloudflare, IFAD)	AfDB Uganda road documents, AfDB CAIIP-2 appraisal, IFAD PRELNOR supervision report 2023 (two attempts)	download by hand in a browser
Access refused or timed out	Uganda Road Fund annual reports and maintenance plans (roadfund.ug); World Bank DRDIP completion report	retry by hand; request from the Road Fund (letter 03)
Rate limit	World Bank subnational RAI dataset (HTTP 429)	retry later; the national figure came from the SDG series
Gated by a form	MapSPAM 2020	fill in the guestbook by hand
Withdrawn or not resolvable	HDX "Uganda Road Network UNRA" (404); MoWT Research Information Centre (DNS fails); DUCAR manuals (404)	request from MoWT (letters 01 and 02)
No document found	bridge cost per metre of span (JICA, World Bank, ReCAP searched)	request from MoWT and UCSN (letters 02 and 04)

Two of these failures matter more than the rest. The Road Fund's annual reports would give maintenance unit rates and disbursements by agency, which the engine now takes from guidelines recovered through web archive copies. And no open document gives a bridge cost per metre of span: the engine's bridge price rests on a single Bailey bridge (Gem Farm, 30.4 m, FY2018/19).

20.3 What the acquisitions changed

The acquired data entered the engine in runs 3 to 8. Table 20.4 links each gap in design chapter 19 to what now fills it, and to what it did to the results.

Table 20.4: Gaps of docs/19 §19.12 and what now fills them

Gap	Now in use	Effect on the Uganda run
Surface	HeiGIT 2024 (non-commercial) on 4,637 km; GRIP4 refines "unpaved" to gravel on 4,848 km; OSM `tracktype` on 859 links	RAI lower bound 20.2 % → 28.6 % (run 3); surface now known on 31.7 % of network km
Condition	MoWT and UBOS poor km by network, FY2022/23; OSM `smoothness` carried as a proxy and never used	condition programmes by network; 139,691 poor km in total, none placed on a link
Unit costs	275 extracted rows; 89 used by the 14 evidence codes, all E5 EVIDENCE_CANDIDATE	priced projects 8,116 (run 2) → 41,867 of 41,890 (run 8)
Structures	5,039 culverts snapped to 3,676 links; 1,958 river crossings on 1,849 links with no recorded structure	a structure-survey list, not a closure; no span, load or condition
Seasonal crossings	3,943 ford points, 250 ford ways and 12 low-water crossings; 85 flood-prone ways; JRC flood v2	2,248 seasonal-crossing links; 3,006 flood-exposed links (1,368 km)
Speeds	OSM `maxspeed` as an upper cap only	speeds remain generic defaults below the cap
Gateways and markets	28 of 33 OSM ferry routes as links; airports 1 → 6; markets +106 from OSM	island and lake crossings routable; ferry speed and availability ASSUMED
Equity	AID IWI and RWI on about 172,000 settlement cells	RAI reported by wealth quintile
Agriculture	MapSPAM 2010	agricultural production cells for freight generation

Three of these acquisitions changed the headline results. Surface from HeiGIT and GRIP4 raised the lower bound of the rural access index by 8.4 points. The evidence prices made it possible to price engineered-earth upgrading, which had blocked 32,630 projects in run 2 and 32,826 in run 4. And the agency condition statistics created a network-level condition gap where there had been none. None of the three removes the exploratory flag. Surface is still unknown on 68.2 % of network km, condition is still not known for any link, and every evidence price is a candidate.

The engine records each acquisition's effect but never stretches it. The 5,039 culverts are an inventory without condition. The 1,958 unrecorded crossings may be mapping gaps rather than missing structures, so they go to a survey list, not to a closure. OSM `smoothness` is volunteered and clusters in residential streets (248 of its 379 "bad" ways), so it is held as a proxy and does not trigger works.

20.4 The two estate folders in depth

A separate assessment examined the two estate folders closest to transport: `D:\EDIP\Data\Transportdata` (7.4 GB, 661 files) and `D:\EDIP\Data\Infrastructuredata` (3.6 GB, 484 files). It compared them, gap by gap, against geography snapshot `UGA-2026.09-b`. Table 20.5 gives its verdicts.

Table 20.5: Verdicts of the gap assessment for the two estate folders

Gap	Best source in the folders	Verdict
Link condition	OSM `smoothness`: 4,994 links, 989 km via the engine's tag table; `flood_prone` on 74 ways	partial: a weak proxy, no survey
Surface beyond what is used	OSM `tracktype` on 840 unknown-surface links (313 km)	partial: about 0.1 % of links
Lanes and width	nothing new	no

Gap	Best source in the folders	Verdict
Bridges, culverts, fords	5,423 culvert waterways; 4,158 unstructured road–water intersections; 250 ford ways; 12 low-water crossings	use for inventory; no condition or load rating
Traffic and axle load	3 OSM weighbridges, 111 toll ways, a blank ATDGR template	no
Speeds	OSM `maxspeed` on 9,770 links (2,182 km); Nelson 2015 travel times for validation	partial: posted limits, not observed
Unit costs and contracts	an empty ATDGR price schema	no
Project pipeline	OSM construction (59 ways, 69 km) and proposed (18 ways, 11 km)	partial, tiny
Rail, ports, airports, borders	1,000 km of rail, 33 ferry routes, 61 ferry terminals, 42 aerodromes, 25 border-control and 7 customs points	partial: locations, no capacity
Freight and passenger flows	blank templates	no
Markets, mines, industry, tourism	884 OSM markets, of which 100 are new; 311 mining points; 23 SEZs	partial: locations, no magnitudes
Population, poverty, travel surveys	WorldPop/GPW4 hexes; IWI and RWI; 9.9 million building footprints	partial: no trip rates
Hazards and land cover	OSM wetland and flood-prone tags only	partial
Outputs for validation	Nelson 2015 travel times to cities and ports	use for validation only

The bottom line of the assessment is that the two folders do not close the four decisive gaps: link condition, traffic, unit costs and the project pipeline. ATDGR holds only registry entries and blank templates for those. What the folders did add, and what runs 4 to 8 now use, is a set of OSM-derived structure and multimodal layers: culverts, unstructured crossings, ford ways, ferry routes, road numbers on 11,118 links (3,224 km) and about 100 more markets. Road numbers matter beyond their face value. They are the only key that will let agency condition or traffic data, which come by road number and chainage, be joined to the engine's links.

20.5 What only agencies can supply

Table 20.6 lists what is still missing after the open-source search, and who holds it.

Table 20.6: Data still missing after the open-source search, and the holder

Missing	Holder	What it would change
Link-level condition (IRI or visual, with survey date)	UNRA (now within MoWT); MoWT DUCAR	rehabilitation needs by road; the deterioration projection; G2 by link
Surface on about 130,000 km of district, urban and community roads	MoWT DUCAR; district inventories	the GIS-5 test; the RAI bounds would narrow
Structure span, load rating and condition	UNRA bridge system; DUCAR	structure works and criticality; the 1,958-crossing survey list
Traffic counts and axle loads	UNRA; weighbridge operators	traffic standards and G3 capacity gaps; pavement design loads
Unit costs for engineered earth, gravel rehabilitation, culverts, drifts, vented fords and bridges	MoWT, URF, UCSN	thin evidence (one or two observations) replaced
Project pipeline with chainages and contract sums	MoWT; MoFPED	committed works netted out of the programme
Road Fund maintenance plans and disbursements	URF	maintenance rates and the pool split
Rail, port and border-post capacity and volumes	URC, URA, NCTTCA	gateway capacity gaps
Design, access and maintenance standards	MoWT, MoH, MoES, MAAIF	norms move from *assumed* to *normed*

Missing	Holder	What it would change
District population projections and travel survey data	UBOS	replaces population carried forward at an assumed 3 % a year

20.5.1 The request letters

Six request letters have been drafted in `docs/requests/`, one per holder, each as a Markdown file with a Word copy. **None has been sent.** Each carries placeholders for the sender, date and reference that must be filled in, and the addressee must be confirmed before sending: the sources record that UNRA was merged into MoWT in 2024, so letter 01 is addressed to UNRA or whichever MoWT unit now holds the national road asset system. Table 20.7 lists them.

Table 20.7: Draft data-request letters (drafted, not sent)

Letter	Recipient	Main content	Priority
01	UNRA or the MoWT national roads unit	condition and surface on national roads; bridge inventory and load ratings; traffic counts and axle loads; ferry schedules; project pipeline with chainages	1
02	MoWT, DUCAR directorate	surface and condition on district, urban and community roads; culverts, drifts and footbridges; unit costs for engineered earth, gravel rehabilitation and structures	1
03	Uganda Road Fund	annual reports, maintenance plans, unit rates, disbursements and outturn	2
04	UCSN (internal memo)	new codes for engineered earth, gravel rehabilitation, spot improvement, culverts, footbridges and earth-road maintenance; review of the illustrative bridge price and the evidence prices	1
05	MoH, MoES, MAAIF and MoWT	confirm or replace the access standards and the road design, maintenance and condition rules	2
06	URA, URC, NCTTCA and UBOS	border volumes and clearance times; rail status and capacity; corridor transit times; population projections and travel data	3

Priority 1 means the gap blocks the core results; priority 2 that it materially improves costs or results; priority 3 that it refines or validates. The letters ask for specific fields and formats, propose join keys (district P-code, road number and chainage, OSM id), and commit to planning use only, no redistribution without consent, and attribution. Some figures quoted in the letters come from earlier runs (for example letter 03 cites the maintenance funding ratio of about 0.47 from run 1, and letter 04 cites 32,826 unpriced engineered-earth projects from run 4). Since run 8, those projects are priced from evidence and the ratio is 0.58. The letters should be updated before they are sent.

The priority order follows design chapter 19 §19.13, which ranks acquisitions by how many gaps each unlocks. The road inventory and condition surveys come first, because one data-sharing agreement with MoWT could deliver surface, condition, structures and traffic together.

20.6 Licences carried into outputs

The licences of the sources constrain what the engine may publish. Table 20.8 summarises them.

Table 20.8: Licences and what they require of published outputs

Source	Licence	Constraint
OSM and Geofabrik (network, tags, points of interest)	ODbL 1.0	published link layers are share-alike; maps credit "© OpenStreetMap contributors"
HeiGIT road surface	non-commercial only; Planet Labs terms	internal planning use; no redistribution of the layer
JRC flood, MapSPAM, HydroRIVERS, WorldPop, WDI	CC BY 4.0	attribution; MapSPAM products carry IFPRI's notice
GRIP4	CC0	none (credit OSM where derived)
OurAirports	public domain (PDDL)	none
UNPS 2019/20	World Bank LSMS licence, no redistribution	aggregates only
Government documents (MoWT, URF, BMAU, UBOS)	no explicit licence	cite; treat as public documents
Agency data when received	unknown; permission required	per data-sharing agreement; nothing redistributed until terms are recorded

20.7 Data quality and limitations

- **The estate state is a snapshot.** The surveys were made on 23 and 24 September 2026. Since then, the TransportModelling folder at `D:\EDIP\models\TransportModelling\AICTEM_Transport_Infrastructure` is no longer present on disk. The current geography snapshot (`UGA-2026.09-c`) and norm versions were built while it was present and are stored with the engine, so runs still work. But `scripts/build_geo_snapshot.jl` (OSM tag table, AID points of interest, elevation model) and `scripts/seed_norms.jl` (the TransportModelling configuration) both read from that path and cannot be re-run until it is restored.
- **Open sources are biased towards arterial roads.** HeiGIT covers motorway to secondary roads only. Surface remains unknown on about 119,500 km of community roads, 36,500 km of tracks and 10,800 km of district roads in the assessed snapshot.
- **Volunteered data is not a survey.** OSM `smoothness`, ford points and market locations reflect mapping effort. Northern Uganda has more mapped markets; Kampala has few mapped health facilities in the MoH register.
- **Several evidence prices rest on one or two observations.** Engineered earth, spot improvement and bridges are the thinnest.
- **Condition is network-level.** The 139,691 poor km come from agency statistics by network. They cannot be placed on a road.
- **Some data could not be read.** The AID `.qs` aggregates need R 4.5, and the ATDGR boundaries are LFS stubs.

20.8 Implications for planning

- The open-source search has gone about as far as it can. Of 60 sources assessed, 27 were downloaded and they moved the RAI lower bound and the number of priced projects. The remaining gaps (link condition, traffic, standards and the pipeline) are held only by agencies.
- A single data-sharing agreement with MoWT, covering the national and DUCAR road inventories, condition surveys, bridges and traffic, would close more gaps than any other step. Letters 01 and 02 ask for exactly that. They have been drafted and not sent.

- The draft letters quote results from runs 1 and 4. They should be updated to run 8 before they go out, so that agencies see the current state.
- The access standards and road rules in letter 05 are what turn a draft norm into a government norm. Until the ministries confirm them, every gap and cost in this report is measured against a placeholder.
- The 1,958 river crossings with no recorded structure and the 169,343 km on the survey list are cheap to verify compared with the works they might trigger. A field survey programme would pay for itself by removing unknowns before money is committed.
- The HeiGIT surface layer is non-commercial. Any public release of link-level surface must exclude it or obtain permission.
- The TransportModelling folder should be restored to its estate path, or its Uganda inputs copied into the engine's own data folder, so that the geography and the draft norms can be rebuilt.

CHAPTER 21

21 GIS analysis across modules

Every chapter of this report has used maps. This chapter looks across all of them and asks three questions. Which needs, gaps and results are geographically clustered? Where are the persistent hot spots and cold spots? And what does the spatial pattern mean for how a transport plan is delivered? It implements design chapter 20 (`docs/20_GIS.md`) and the spatial statistics of `docs/23_MATHEMATICAL_MODELS.md` §23.16. In TIPRAS, unlike in the health and education engines, GIS is not a reporting layer added at the end. It is the computational core of the engine: needs are distances, supply is a network, gaps are places and links, costs depend on terrain and length, and allocation decides locations. The chapter therefore first describes the spatial machinery that every module uses, and then reports the cross-module statistics of the reference run 014c9cd7.

21.1 What the GIS subsystem does

21.1.1 Vector inside, raster outside

The engine draws a hard line between two kinds of spatial work (Table 21.1). Anything that needs a raster or a digital elevation model is done once, outside the engine, and arrives as an attribute. Anything that needs the network and must be repeated for every scenario or every candidate project is done inside, in `src/GIS/`. The reason is speed. The optimiser and the appraisal re-evaluate the network after each candidate; a round trip to an external system per candidate would make a national run infeasible.

Table 21.1: Where each kind of spatial work is done

Work	Where	Why
Boundaries, gridded population, noded road network, hazard rasters, terrain	SDI (`D:\EDIP\SDInfrastracture`)	The estate's spatial data infrastructure, already loaded for Uganda
Asset inventory, attributes, condition, traffic counts, projects	ATDGR (read-only)	The transport repository; no modelling there by design
Raster friction, off-road walking cost, alignment of new links, flood modelling, slope	SDI geoprocessing or TransportModelling	Heavy raster work, done once per geography version
Graph build, snapping, shortest paths, catchments, assignment, criticality, overlays	`src/GIS/` in this repository	Must be repeated per scenario and per candidate
Map rendering	`scripts/render_maps.py` (PDF atlas); Leaflet in `web/`	Draws values already in the store; computes nothing

The rule has a practical consequence for Uganda. Terrain arrives as a class per link (flat, rolling or mountainous) computed from the Copernicus GLO-30 elevation model before the run. Flood exposure arrives as a yes/no flag at the link midpoint from the JRC 100-year flood hazard map. Neither is recomputed inside a run.

21.1.2 Coordinates and the geography snapshot

All geometry is stored in geographic coordinates (EPSG:4326), as in ATDGR and SDI. Every length, area, buffer and snap distance is measured in the projected system WGS 84 / UTM zone 36N (EPSG:32636), the same as TransportModelling, so that lengths agree between the two systems. A length computed in degrees is a defect, and a unit test (gate GIS-4) catches it.

Each run pins one immutable geography snapshot. The reference run uses `UGA-2026.09-c`, a set of Arrow files built by `scripts/build_geo_snapshot.jl`: links, link coordinates, vertices, settlements, nodes, districts, district adjacency, the population projection and the agricultural production cells, with a manifest that records sources, parameters and caveats. The snapshot holds 817,354 network links and 247,972 km of road, 175,367 settlement cells of 1 km², 24,952 sector nodes and 135 districts. The link identifier is the SDI road id plus a sub-id, because ATDGR's canonical road table is still empty. The manifest states eight caveats, of which three matter most for this chapter:

- condition is unknown on every link, because the estate holds no link-level condition data;
- surface comes from OpenStreetMap tags, then from the HeiGIT 2024 paved/unpaved estimate, then from GRIP4, and 68.3 % of network km still has unknown surface;
- the population projection carries the UBOS 2022 district totals forward at an assumed 3.0 % a year.

21.1.3 The core operations

Table 21.2 lists the spatial operations that the ten modules call. They are implemented in six files (`Network.jl`, `Snap.jl`, `TravelTime.jl`, `Criticality.jl`, `Geometry.jl` and `SpatialStats.jl`), 586 lines in all.

Table 21.2: Core spatial operations in `src/GIS/`

Operation	What it does	Used by
Snap a point to the network	Nearest link by a projected KD-tree over link vertices, then exact projection onto the segment; returns link, offset and off-network distance	TI-01 settlements and nodes; TI-05 G6
Build the planning graph	Edge weight = length / speed (surface, condition, terrain, season); closed edges removed	TI-03, TI-04, TI-09
Multi-source travel time	One Dijkstra run per destination class from a super-source joined to all destinations, on the reversed graph	TI-02, TI-04, TI-07, TI-09
Off-network walking leg	Straight-line distance × terrain factor ÷ walking speed	TI-04
Catchments	Nearest-destination allocation of settlements	TI-04 (descriptive)
Assignment	All-or-nothing loading along shortest paths	TI-03
Criticality	Remove a link and recompute; cut edges first by Tarjan's bridge algorithm	TI-05 G5
Overlay	Link length by polygon (district, hazard zone)	TI-01, TI-05 G7
Spatial weights and statistics	Queen contiguity; Moran's I, LISA, concentration index, Gini	All modules (this chapter)

The multi-source Dijkstra is the operation that makes a national run possible. Rather than computing the distance from each of 175,367 settlements to each of 3,169 health facilities, the engine joins every facility to a single artificial source and runs one shortest-path search backwards through the network. The label at each vertex is then the travel time from that vertex to the nearest facility. One search per destination class replaces several million searches.

21.1.4 Layers in the result store

Every module writes its spatial results as rows of one table, `res_gis_layer`, keyed by run, module, layer, geometry kind (link, node, settlement, cell, district), geometry key, year and scenario, each with a value, a unit and a quality tier. At national scale the layers dominate the store, so the engine keeps them lean. Link layers are sparse: only non-zero values are written. Snapshot attributes such as surface are not copied into layers. Most module layers are held at district level. This is why a national run adds about 2 GB to the store and why `tipras prune --keep 1` is run before each new national run.

The atlas is drawn by `scripts/render_maps.py`, which reads `res_gis_layer` directly and writes one PDF per module. The web interface's GIS explorer shows the same district layers with their LISA classes. Neither computes anything; both draw values already in the store.

21.2 Method

21.2.1 Spatial weights

Two districts are neighbours if their boundaries touch, at a point or along an edge (queen contiguity), with a tolerance of about 50 m to absorb small slivers in the COD-AB boundary file. The snapshot stores the result as a list of neighbour pairs. In `UGA-2026.09-c` there are 360 pairs among the 135 districts. Each district has between 2 and 8 neighbours, 5.3 on average. No district is left without neighbours. Kalangala, the island district in Lake Victoria, has seven, because the district boundaries extend into the lake: Buvuma, Kalungu, Kyotera, Masaka, Mpigi, Mukono and Wakiso.

The weights are row-standardised: each neighbour of district d receives weight 1 divided by the number of d's neighbours. A district's *spatial lag* is then simply the average value among its neighbours.

21.2.2 Global Moran's I

The global statistic tests whether a layer is clustered across the whole country:

$$I = \frac{n}{S_0} \frac{\sum_d \sum_e w_{de}\, z_d\, z_e}{\sum_d z_d^2} \tag{21.1}$$

where

- n = 135 is the number of districts;
- z_d = x_d − x̄ is district d's value of the layer minus the national mean;
- w_de is the row-standardised weight of neighbour e for district d (zero if they are not neighbours);
- S_0 is the sum of all weights, which equals n after row-standardisation.

Under no spatial pattern the expected value is:

$$E[I] = -\frac{1}{n-1} \approx -0.0075 \tag{21.2}$$

A value of I well above zero means that similar values sit next to each other. A value near zero means no pattern. A negative value means a checkerboard, with high values next to low ones.

21.2.3 Permutation inference

Significance is tested without assuming a distribution. The engine shuffles the district values at random 999 times with a fixed seed (20260923), recomputes I each time and counts how many shuffled values lie at least as far from E[I] as the observed one (N in the equation below is that count):

$$p = \frac{1 + N(|I^{perm} - E[I]| \geq |I - E[I]|)}{999 + 1} \tag{21.3}$$

The smallest attainable p-value is therefore 0.001. The test is two-sided. Because the seed is fixed, the same run always gives the same p-values; the validation chapter relies on this.

21.2.4 Local Moran (LISA) and the cluster classes

The local statistic measures how much each district contributes to the global pattern:

$$I_d = \frac{z_d}{m_2} \sum_e w_{de}\, z_e, \qquad m_2 = \frac{1}{n} \sum_d z_d^2 \tag{21.4}$$

Its significance is tested by *conditional* permutation. District d's own value is held fixed while its neighbours are drawn at random, 999 times, from the other 134 districts. A district that is significant at α = 0.05 is classed by the sign of its own deviation and of its spatial lag:

- **HH (hot spot):** a high value among high neighbours;
- **LL (cold spot):** a low value among low neighbours;
- **HL and LH (spatial outliers):** a high value among low neighbours, or a low value among high neighbours.

Districts that are not significant are shown in grey on every LISA map. For a gap layer, "high" means a large gap, so an HH cluster is a block of districts with large gaps. For the Rural Access Index (RAI), "high" means good access, so the cold spot (LL) is the problem area.

21.2.5 Other spatial statistics in the engine

Two further statistics are implemented in `SpatialStats.jl`. The concentration index tests whether a quantity (supply, or allocation) is concentrated where need is highest:

$$CI = \frac{2}{\bar{x}_w} \operatorname{cov}_w(x_d, R_d) \tag{21.5}$$

where R_d is district d's population-weighted fractional rank by need, from least to most needy, and x̄_w is the population-weighted mean of x. A positive CI is pro-need. The Gini coefficient of travel time measures inequality of access across the population:

$$G = \frac{\sum_i \sum_j P_i P_j \, |tt_i - tt_j|}{2 \left(\sum_i P_i\right)^2 \bar{t}} \tag{21.6}$$

where P_i is the population of settlement i, tt_i its travel time and t̄ the population-weighted mean travel time. TI-04 computes the Gini of travel time to primary health care. Neither statistic is a ranking or a weight; both describe a distribution.

21.2.6 What the method does not do

Two choices limit what the LISA classes can be used for. First, each district is tested at α = 0.05 without a correction for testing 135 districts at once. About 7 districts per layer can be expected to come out significant by chance alone, so a single district flagged in a single layer is weak evidence; a block of neighbouring districts flagged in several layers is strong evidence. Second, the statistics are computed on district totals, not on rates. A district with a large population tends to have a large unserved population, so some of the clustering of the G1 layers reflects the clustering of people rather than of poor access. The design chapter asks for counts to be normalised per capita or per km; the district layers of this run are not yet normalised. Both points are taken up again in the section on data quality.

21.3 Results: which layers cluster

Figure 21.1 and Table 21.3 give the global Moran's I for the 25 headline district layers of TI-01 to TI-06 and TI-10. Of the 25 layers, 23 are significantly clustered at the 5 % level. Two are not: G1 unserved population for the district administrative centre (I = 0.086, p = 0.051), which misses the threshold narrowly, and G7 flood-exposed critical population (I = −0.007, p = 0.793), which shows no pattern at all.

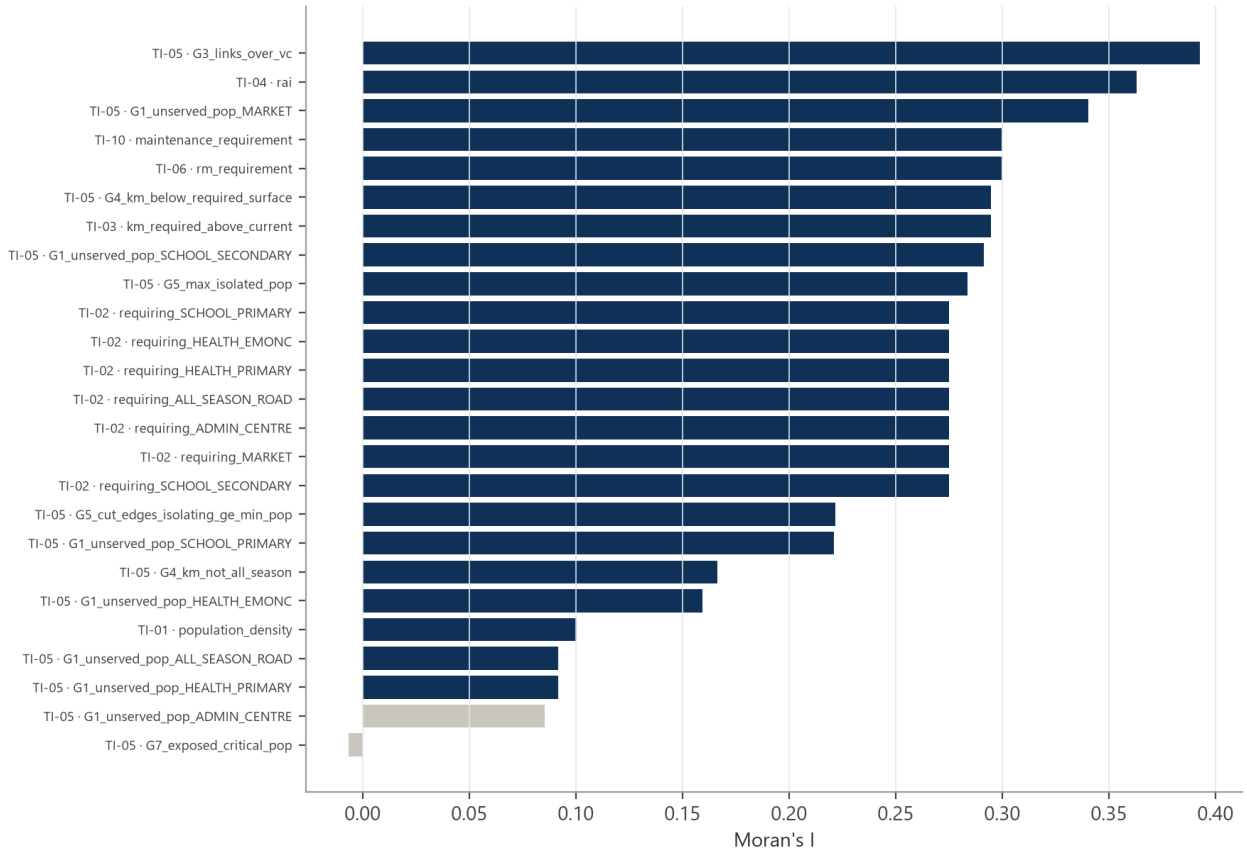

Figure 21.1: Global Moran's I of every district layer (blue: significant at 5 %)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 21.3: Global Moran's I of the district layers, with the number of districts in each LISA class

Module	Layer	Moran's I	p	HH	LL	HL	LH
TI-05	G3_links_over_vc	0.393	0.001	5	0	0	3
TI-04	rai	0.363	0.001	9	10	0	2
TI-05	G1_unserved_pop_MARKET	0.340	0.001	8	0	0	4
TI-10	maintenance_requirement	0.300	0.001	8	5	0	5
TI-06	rm_requirement	0.300	0.001	8	5	0	5
TI-05	G4_km_below_required_surface	0.295	0.001	10	6	0	5
TI-03	km_required_above_current	0.295	0.001	10	6	0	5
TI-05	G1_unserved_pop_SCHOOL_SECONDARY	0.291	0.001	14	3	0	1
TI-05	G5_max_isolated_pop	0.284	0.001	5	0	0	5
TI-02	requiring_SCHOOL_PRIMARY	0.275	0.001	4	0	0	4
TI-02	requiring_ALL_SEASON_ROAD	0.275	0.001	4	0	0	4
TI-02	requiring_HEALTH_PRIMARY	0.275	0.001	4	0	0	4
TI-02	requiring_ADMIN_CENTRE	0.275	0.001	4	0	0	4
TI-02	requiring_HEALTH_EMONC	0.275	0.001	4	0	0	4
TI-02	requiring_MARKET	0.275	0.001	4	0	0	4
TI-02	requiring_SCHOOL_SECONDARY	0.275	0.001	4	0	0	4
TI-05	G5_cut_edges_isolating_ge_min_pop	0.222	0.001	3	0	0	5
TI-05	G1_unserved_pop_SCHOOL_PRIMARY	0.221	0.001	13	4	0	3
TI-05	G4_km_not_all_season	0.167	0.004	9	0	0	2
TI-05	G1_unserved_pop_HEALTH_EMONC	0.159	0.004	8	0	0	3
TI-01	population_density	0.100	0.001	7	11	0	1
TI-05	G1_unserved_pop_ALL_SEASON_ROAD	0.092	0.033	2	2	0	6
TI-05	G1_unserved_pop_HEALTH_PRIMARY	0.092	0.039	4	0	0	4
TI-05	G1_unserved_pop_ADMIN_CENTRE	0.086	0.051	7	0	0	5
TI-05	G7_exposed_critical_pop	-0.007	0.793	0	0	0	3

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_moran.csv).

The values are moderate. The highest is 0.393, for G3 links over capacity. None reaches the 0.5 range that wind speed or electricity access reaches in the energy engine. Transport gaps in Uganda are clustered, but they are not organised in a few very large regional blocks. They form several smaller clusters in different parts of the country, each driven by a different cause.

The layers fall into three groups.

Moderately clustered (I from 0.25 to 0.40). Thirteen layers:

- G3 links over the volume-to-capacity threshold (0.393);
- the RAI today (0.363);
- G1 unserved population for markets (0.340);
- the routine maintenance requirement (0.300, identical in TI-06 and TI-10);
- G4 km below the required surface (0.295, identical to TI-03 km required above current);
- G1 unserved population for secondary schools (0.291);
- G5 largest isolated population (0.284);
- the seven TI-02 requirement layers (0.275 each).

Weakly clustered (I from 0.09 to 0.25). Eight layers: G5 cut edges (0.222), G1 for primary schools (0.221), G4 km not all-season (0.167), G1 for emergency obstetric care (EmONC, 0.159), population density (0.100), and G1 for all-season roads and for primary health care (0.092 each).

Not clustered. G1 for the administrative centre ($p = 0.051$) and G7 flood-exposed critical population ($p = 0.793$).

21.3.1 Why some layers have identical values

Three sets of layers in Table 21.3 carry exactly the same Moran's I. This is not an error; it shows how the layers are built.

- **The seven TI-02 requirement layers (0.275).** The population that requires each destination class is the district population multiplied by a single national share for the group the standard applies to (all persons, women aged 15–49, children 6–12 and so on; `group_share.csv`, quality *assumed*). Multiplying every district by the same constant does not change Moran's I. The seven layers are therefore one layer — district population — seen seven times. The statistic will separate them only when HouseholdModelling or UBOS supplies age and sex structure by district.
- **TI-03 km required above current and TI-05 G4 km below required surface (0.295).** These are the same quantity counted by two modules: the km where the required surface outranks the existing one.
- **TI-06 routine maintenance requirement and TI-10 maintenance requirement (0.300).** TI-10 carries the TI-06 requirement into the budget unchanged.

Counting distinct quantities, the 25 layers contain 17 independent ones, of which 15 are clustered.

21.3.2 What the clustering says

Population density is only weakly clustered (0.100). Dense districts — Kampala and the Elgon highlands around Mbale — are separated by sparse ones. Yet the requirement layers, which are proportional to population, cluster more strongly (0.275). The difference is that density divides by area while the requirement layers do not. Large, populous districts around Kampala (Wakiso, Mukono, Luwero) form a block of high totals, and the requirement layers pick that up.

The service-access gaps cluster where the services are thin. G1 for markets (0.340) and for secondary schools (0.291) are the most clustered of the G1 classes. These are the classes where the destination network is sparsest and where whole sub-regions lack destinations. G1 for primary health care and for all-season roads cluster only weakly (0.092), because unserved people for these classes are spread across almost every district.

The network gaps cluster where the network is busiest or hardest. G3 congestion (0.393) is concentrated around Kampala, the only place with traffic heavy enough to exceed capacity. G5 isolation (0.284) is concentrated in mountain districts. G4 not all-season (0.167) is concentrated along the western rift and in the north.

G7 flood exposure does not cluster because only three districts have any exposed critical population: Ntoroko (298,306 persons), Nakaseke (12,089) and Napak (6,745). A layer that is zero in 132 districts cannot form a cluster. This result reflects the narrow G7 test — a link must be flood-exposed, critical and without a short detour — and the single hazard layer used (JRC river flood, 100-year return period), not a finding that flood risk is evenly spread.

21.4 Results: where the hot spots and cold spots are

Figure 21.2 maps the LISA classes of six key layers side by side: the RAI today, G1 all-season road access (beyond 2 km), G1 EmONC, G4 not all-season, G5 isolated population and G7 flood-exposed population. Each panel is a separate gap dimension in its own unit; the panels are placed together for comparison and are never added. Red is HH, dark blue LL, orange HL and light blue LH. Grey districts are not significant.

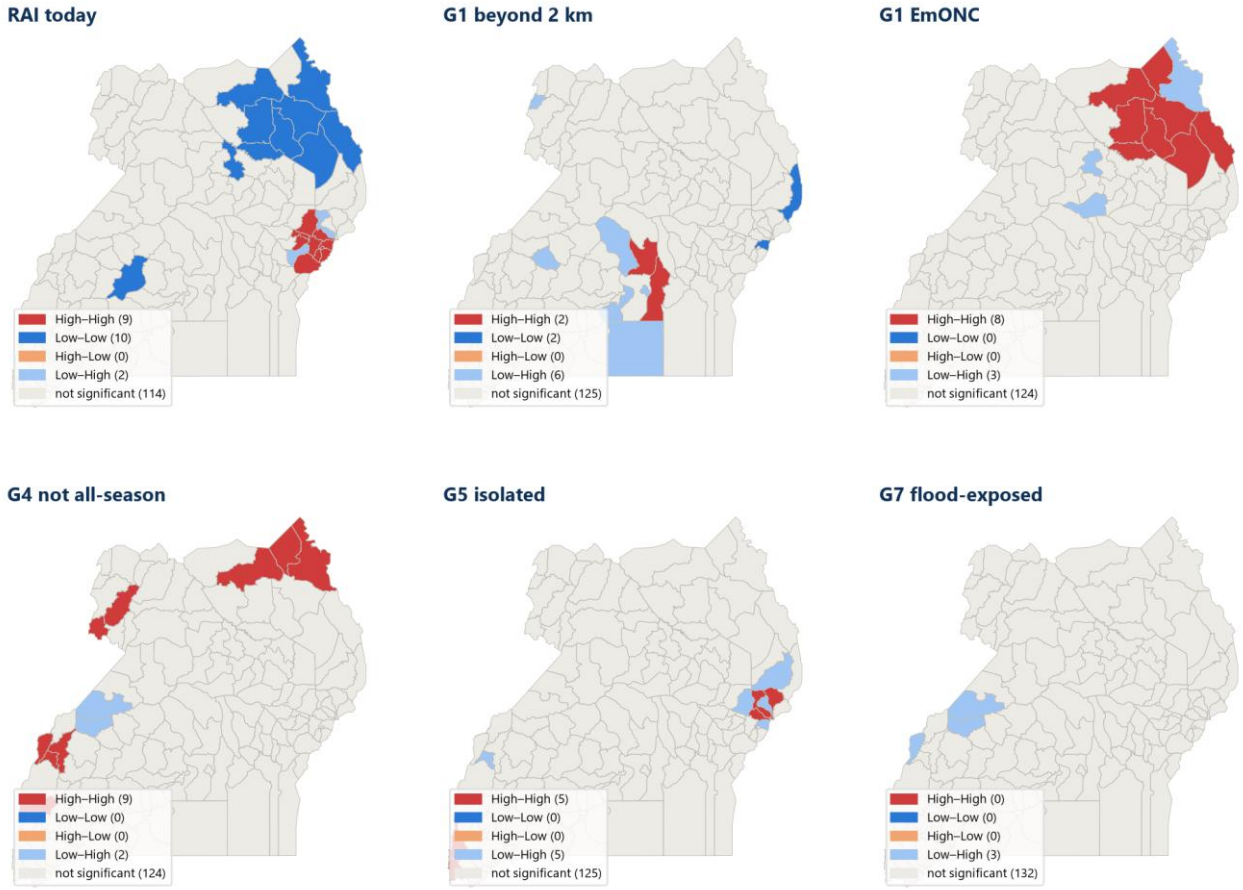

Figure 21.2: Hot spots and cold spots of the key gaps (LISA, queen contiguity, 999 permutations)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

21.4.1 The Rural Access Index

The RAI panel shows the clearest pattern in the report. The cold spot (LL) is a block of ten districts across the north-east and north: Abim, Agago, Kaabong, Kitgum, Kotido, Lira, Moroto, Napak and Otuke, with Mubende as a separate cold spot in the centre-west. Their mean RAI is 14.6 %, against a national district mean of 28.2 %. The hot spot (HH) is a block of nine districts in the east around Mount Elgon and in Bukedi: Budaka, Bududa, Bukedea, Butebo, Manafwa, Mbale, Namisindwa, Sironko and Tororo. Their mean RAI is 56.6 %. Bulambuli and Butaleja are low-access outliers (LH) inside the eastern hot spot.

The eastern hot spot coincides with the dense-population hot spot of TI-01 (Budaka, Bududa, Butaleja, Kampala, Manafwa, Namisindwa, Tororo). That is expected: where people are dense, more of them live within 2 km of any road, including an all-season one. But the RAI here is the *lower bound*, computed with unknown surface treated as earth. Its pattern therefore also reflects where OpenStreetMap contributors have tagged road surfaces. Part of the contrast between the east and the north-east may be a contrast in mapping rather than in roads. A surface survey in Karamoja and Acholi would settle the question.

21.4.2 G1 access gaps, class by class

Each access class is tested on its own. The classes are not combined, and they cluster in different places.

All-season road (2 km), 33.3 million unserved. This is the largest gap by far, but it barely clusters (I = 0.092). The HH districts are Luwero and Mukono; the LL districts are Amudat and Bududa. Six districts are LH outliers, among them Kampala, Kalangala and Mpigi. The unserved population is spread across almost every district. The two HH districts most likely appear because their large rural populations sit next to Wakiso, which has the largest absolute gap (1.71 million).

Primary health care, 3.21 million unserved. The HH cluster is four districts in Acholi and Karamoja: Agago, Kitgum, Moroto and Napak, with a mean unserved population of 43,893 against a national mean of 23,768. Kotido has the largest gap of any district (227,078) but sits next to districts with smaller gaps.

EmONC, 1.04 million unserved. The EmONC panel of Figure 21.2 shows a solid red block across the north-east: Abim, Agago, Karenga, Kitgum, Kotido, Moroto, Napak and Otuke. These eight districts hold 171,416 unserved persons, 16.6 % of the national EmONC gap, at a mean of 21,427 per district against 7,667 nationally. Amolatar, Kaabong and Kole are LH outliers. Of the national EmONC gap, 614,403 persons are unserved because there is no facility within reach even on an upgraded network (cause *no_destination*), and 420,690 because the network is below standard. The north-east cluster is therefore partly a facility siting question for HealthModellingx, not only a road question.

Primary school, 2.94 million unserved. Thirteen HH districts form a belt across the mid-west and Lango: Kakumiro, Kamwenge, Kibaale, Kiryandongo, Kole, Kyankwanzi, Kyegegwa, Kyenjojo, Lira, Mubende, Nwoya, Omoro and Oyam. The LL cold spot is in the south-west and east: Bushenyi, Masaka, Mbale and Rukungiri. Of the national gap, 2.86 million is *no_destination*. A primary school gap is almost entirely a gap in schools within walking distance, not in roads.

Secondary school, 0.86 million unserved. This is the largest HH cluster of any layer: fourteen districts across Acholi, Karamoja and Lango — Abim, Agago, Amuru, Gulu, Kaabong, Karenga, Kitgum, Kotido, Lamwo, Lira, Moroto, Omoro, Otuke and Pader. They hold 201,545 unserved persons, 23.4 % of the national total. The cold spot is in Bukedi: Butaleja, Butebo and Kibuku.

Market, 2.33 million unserved. Eight HH districts form a block in the mid-west: Agago, Kagadi, Kakumiro, Kibaale, Kyegegwa, Kyenjojo, Mubende and Ssembabule. They hold 581,690 unserved persons, a quarter (24.9 %) of the national market gap, at a mean of 72,711 per district against 17,280 nationally. Mubende (141,370) and Kakumiro (140,790) have the largest market gaps in the country. Part of this may be a mapping effect: the market register is thin in the mid-west and dense in West Nile (Arua alone has 149 of the 814 market nodes).

Administrative centre, 4.22 million unserved. Not significant overall (p = 0.051), but seven districts are HH: Abim, Agago, Amuru, Bundibugyo, Karenga, Kitgum and Moroto.

The same sub-regions recur across classes. Acholi and Karamoja appear in the HH clusters of primary health, EmONC, secondary school and administrative centre. The mid-west (Kibaale, Kakumiro, Kagadi, Kyenjojo, Kyegegwa, Mubende) appears in the HH clusters of primary school and market. That recurrence is a finding in itself. It is reported here class by class. It is not a reason to build a combined index.

21.4.3 G3 capacity

Five districts form the congestion hot spot: Kampala, Mityana, Mpigi, Mukono and Wakiso. They hold 2,255 of the 3,795 links over the volume-to-capacity threshold (59.4 %); Wakiso alone has 1,014. Only 42 of 135 districts have any link over the threshold. Kalangala, Luwero and Nakaseke are LH outliers: low congestion next to the metropolitan cluster. The congestion problem is metropolitan, not national. Because lanes are unknown on 276,431 loaded links, G3 counts only links where capacity can be estimated. These links are listed as *capacity unknown*; they are never assumed adequate.

21.4.4 G4 standard gaps

Two G4 layers are mapped, and each keeps its own unit.

Km below the required surface (242,068 km). Ten HH districts centre on Kampala and the central region — Gulu, Kampala, Kyegegwa, Lamwo, Luwero, Mityana, Mpigi, Mukono, Nakaseke and Wakiso. Six LL districts are in Karamoja and on Mount Elgon: Amudat, Bulambuli, Kumi, Kween, Moroto and Nakapiripirit. This is mostly a map of network length. The standard gap is measured in km, and large districts with dense community networks have the most km. Wakiso leads with 9,055 km.

Km not all-season (3,299 km). This layer measures something different: km of required all-season links that close in the wet season. The G4 panel of Figure 21.2 shows three separate hot spots, nine HH districts in all, with a mean of 58.5 km against 24.4 km nationally: the Rwenzori region (Bundibugyo, Bunyangabu, Kabarole, Rubirizi), West Nile (Madi Okollo, Zombo) and the north (Kaabong, Karenga, Kitgum). Ntoroko has the largest single value (463 km). It is not in the HH class, because some of its neighbours have low values, but it lies next to the Rwenzori cluster. These are the places where the wet season, not the network itself, cuts access.

21.4.5 G5 network redundancy

Two G5 layers are mapped.

Largest population isolated by one link failure. The HH districts are Bulambuli, Kween and Sironko on Mount Elgon, and Kanungu and Kisoro in Kigezi. Their mean is 403,519 persons; Sironko's largest isolate is 716,686 persons, the largest in the country. These are mountain districts where the network is tree-like: one road up a valley serves everything above it. The G5 panel of Figure 21.2 shows the Elgon cluster in red, with Bududa, Bukedea, Kapchorwa and Nakapiripirit as LH outliers around it.

Cut edges isolating at least the minimum population (3,738). The hot spot is Kampala, Mukono and Wakiso, with 723 cut edges between them. This is a different pattern, and it is a matter of scale. Dense peri-urban networks have many short dead-end spurs, each a cut edge with enough people behind it. The two layers answer different questions, and both are reported.

21.4.6 G7 climate exposure

The G7 panel of Figure 21.2 is almost empty. No district is HH; three — Bundibugyo, Kagadi and Kikuube — are LH outliers next to Ntoroko, which holds 94.1 % of the exposed critical population (298,306 of 317,140). Ntoroko lies on the floodplain of the Semliki river and Lake Albert. Its network has few alternative routes, and the JRC map marks it as flood-exposed.

21.5 Results: how much of the country is in a cluster

Figure 21.3 shows, for each of the 23 significant layers, how many districts fall into each LISA class.

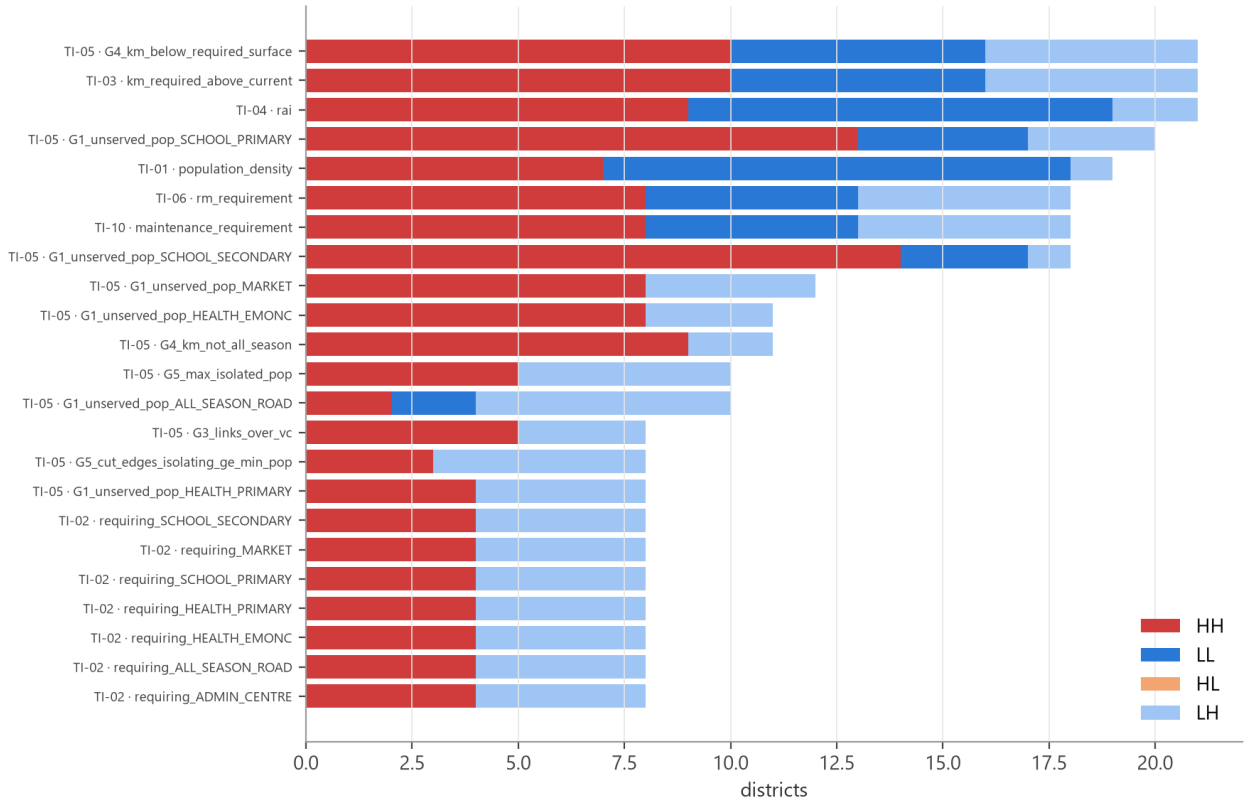

Figure 21.3: How many districts sit in each LISA cluster, layer by layer (significant layers)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Three features stand out.

- **Clusters are small.** The largest number of significant districts in any layer is 21 of 135 (16 %), for the RAI and for km required above current. For most gap layers, between 8 and 12 districts are significant. Most of the country is grey on every map. That is the right reading: in most districts, a district's gap is not predicted by its neighbours' gaps.
- **Gap layers have hot spots but few cold spots.** G1 EmONC, G1 market, G3, G4 not all-season and both G5 layers have no LL districts at all. Large gaps cluster; small gaps do not form blocks of their own. The exceptions are G1 primary school (4 LL), G1 secondary school (3 LL), G1 all-season road (2 LL) and the km-based G4 and maintenance layers (5 or 6 LL), where low values in Karamoja and on Mount Elgon form blocks.
- **There are no HL outliers in the base layers.** No district with a large gap is surrounded by districts with small gaps. Every LH outlier is a district with a small gap next to a cluster with large gaps. Most LH districts are small or island districts (Kalangala, Buvuma), urban districts (Kampala) or districts that are themselves served by a regional centre (Gulu, Lira). They are the edges of the clusters, not separate problems.

21.6 Results: how the programme and the scenarios cluster

The engine computes the same statistics for every district layer written by TI-08, TI-09 and TI-10. These are not in Table 21.3 but are in the run's `gis/gis_summary.csv`. They show how each allocation scenario redistributes spending and access in space. Scenarios are compared here, not ranked.

21.6.1 Allocation per capita

The allocation per capita clusters in every scenario except one, but in different places:

- **B (preserve first), I = 0.423.** The strongest pattern. The HH cluster is thirteen districts in the Rwenzori and Kigezi regions (Bundibugyo, Bunyangabu, Kabarole, Kagadi, Kamwenge, Kanungu, Kasese, Kisoro, Kitagwenda, Kyenjojo, Ntoroko, Rubanda, Rubirizi). The LL cluster is nine districts in Busoga and around Jinja (Buikwe, Buyende, Iganga, Jinja, Kaliro, Kamuli, Luuka, Mayuge, Mukono). B funds the condition programmes, which go where the network is long and hilly.
- **F (regional equity), I = 0.301.** HH across Lango, Acholi and Teso: Abim, Agago, Alebtong, Kapelebyong, Katakwi, Kiruhura, Kole, Kotido, Lira, Moroto, Otuke and Pader.
- **D (productive sectors), I = 0.297.** HH in the mid-west market-gap cluster: Kagadi, Kakumiro, Kibaale, Kikuube, Kyegegwa, Kyenjojo and Omoro.
- **E (social services), I = 0.228.** HH in Karamoja and Acholi: Abim, Agago, Amudat, Bukwo, Kitgum, Kole, Kween, Moroto and Nakapiripirit — close to the EmONC and secondary-school hot spots.
- **A (current trajectory), I = 0.192.** HH in the north: Abim, Alebtong, Amuria, Amuru, Arua, Katakwi, Lamwo, Lira, Otuke and Pader.
- **C (rural access) and OPT (maximising people served), I = 0.153.** A small HH cluster on the Karamoja–Elgon border: Amudat, Bukwo, Nakapiripirit and Otuke. OPT_RAI_EMONC80 gives the same allocation as OPT, because its EmONC floor could not be met.
- **G (climate resilience), I = 0.142.** HH in Kigezi: Kikuube, Mitooma, Ntungamo, Rubanda, Rukiga and Rukungiri.
- **OPT_NPV_RAI50 (maximising net present value with an RAI floor), I = –0.002, p = 0.915.** No spatial pattern. Projects with positive NPV are scattered.

In other words, the choice of scenario is also a choice of region. Spending in B goes to the south-west; in F and E to the north and north-east; in D to the mid-west.

21.6.2 The RAI at the 2030 horizon

The RAI at the horizon is more clustered than today's in every scenario except A (0.292). The highest values are in OPT_NPV_RAI50 (0.531), B (0.459), F (0.442) and C and OPT (0.435), against 0.363 in the do-nothing path. Investment does not dissolve the spatial pattern. It changes it.

Two features persist across scenarios.

- **The Karamoja cold spot persists.** Abim, Moroto and Napak remain in the LL class in every scenario except E, and Kotido in all but D and E. Even scenarios that spend heavily in the north-east leave its RAI below that of its neighbours, because the starting point is so low and the population so dispersed.
- **The hot spot moves west.** In the do-nothing path, the RAI hot spot is Elgon and Bukedi. In B, C, F and the optimiser runs, it extends to the Rwenzori and mid-west (Bundibugyo, Bunyangabu, Kabarole, Kagadi, Kamwenge, Kasese, Kibaale, Kikuube, Kyenjojo, Ntoroko), where these scenarios upgrade most km. Busoga and the Lake Victoria districts (Buikwe, Buvuma, Mukono and others) become a new cold spot in C, F and OPT, because these scenarios spend little there.

21.6.3 Capital in scenario B

The TI-10 capital budget of scenario B is clustered (I = 0.259). The HH districts are Bundibugyo, Bunyangabu, Kabarole, Kagadi, Kampala, Kamwenge, Kibaale, Kitagwenda, Luwero, Mukono and Rubirizi. The steady-state routine maintenance induced by B's assets, 62.7 bn UGX a year, follows the same districts. Every capital map therefore implies a recurrent-cost map, which the district maintenance budgets will have to carry.

21.7 Data quality and limitations

- **Surface is unknown on 68.3 % of network km.** Gate GIS-5 sets a 30 % threshold, so the run is exploratory and every access figure carries the warning. Every layer that depends on surface — the RAI, G1 all-season

road, G4 and the allocation layers — may partly map where OpenStreetMap surface tags exist. The RAI bounds of 28.7 % (unknown surface treated as earth) and 93.5 % (treated as gravel) show how much is at stake.

- **Condition is unknown on every link.** Network condition enters only as agency statistics by network class (G2). No condition layer can be mapped or tested for clustering.
- **The layers are district totals, not rates.** Unserved persons and km are absolute values. Large, populous districts dominate several HH clusters (Wakiso, Mukono, Luwero). Per-capita and per-km versions of each layer, as the design chapter specifies, would separate access from size.
- **The requirement layers are not yet distinct.** Group shares are national constants, so the seven TI-02 layers are one layer. District age and sex structure from UBOS or HouseholdModelling would fix this.
- **No multiple-testing correction.** At $\alpha = 0.05$ per district, about 7 districts per layer may be flagged by chance. Clusters of several contiguous districts that recur across layers are robust. Isolated single flags are not.
- **Destination registers carry mapping bias.** Markets are dense in West Nile and thin in the mid-west. The school register dates from about 2018. 3,631 health facilities in the national register have no coordinates. A G1 cluster may partly reflect a register gap.
- **One hazard layer.** G7 uses only the JRC 100-year river flood map. Landslide exposure on Mount Elgon and in the Rwenzori and Kigezi highlands is not assessed, and an unassessed link is not counted as safe.
- **The concentration index and the Gini are implemented but not yet reported for the district layers.** TI-04 computes the travel-time Gini for primary health care; the run tables do not yet carry concentration indices of allocation against need.

21.8 Implications for planning

- **Plan by sub-region as well as by district.** Acholi and Karamoja recur in the hot spots of four access classes (primary health, EmONC, secondary school, administrative centre). The mid-west recurs in two (primary school, market). Programmes framed at sub-regional scale can address these clusters coherently, while each gap is still reported in its own unit.
- **Separate siting from roads in the north-east.** Most of the EmONC and primary-school gap in the hot spots is *no_destination*. A road programme there will not close it without new facilities or outreach, which is a finding to pass to HealthModellingx and EducationModellingx.
- **Treat congestion as a metropolitan problem.** G3 is concentrated in five districts around Kampala, which hold 59.4 % of over-capacity links. It needs a metropolitan transport answer, not a national road programme.
- **Treat the mountain districts as a redundancy problem.** Mount Elgon and Kigezi carry the largest single-link isolation risks (up to 716,686 persons behind one link in Sironko). Second access routes and structures there are a resilience investment, not an access investment.
- **Read every scenario as a regional choice.** B concentrates spending in the south-west, F and E in the north and north-east, D in the mid-west. The spatial distribution of each scenario, and its maintenance tail, should be shown to districts alongside its national totals.
- **Expect the Karamoja RAI cold spot to persist.** No scenario lifts Abim, Moroto and Napak out of the cold spot except E. Any target for these districts needs its own floor, not a national average.
- **Close the data gaps that drive the maps.** A surface survey in the north and north-east, district age structure, and a complete geocoded facility and market register would change several of the clusters in this chapter. Until then, every map in this chapter is exploratory.

CHAPTER 22

22 Sector linkages: agriculture, minerals, health, education, industry and trade

Transport infrastructure has no purpose of its own. It exists to connect people to services and producers to markets. This chapter asks what the transport plan does for each of the sectors it serves: which sector nodes lack the connection their sector requires, how much of the candidate programme and of each allocation scenario serves each sector, what each sector gains in its own unit, and what TIPRAS sends back to the sector engines. It implements design chapter 21 (`docs/21_SECTOR_LINKAGES.md`). It also records honestly which sector feeds are real today and which are missing or synthetic. In September 2026, most of the sector inputs that the design expects from the sibling engines did not yet exist for Uganda. The engine therefore ran on registers taken from the estate and on draft norms. The linkage machinery is built and exercised, but the sector results in this chapter are exploratory.

22.1 What the module does

22.1.1 One pattern for every sector

Every sector linkage follows the same path through the chain, so that adding a sector is a matter of data and an interface contract, not new code:

1. **Inbound.** The sector model supplies its nodes (facilities, schools, mines, parks, markets, production zones), their magnitudes (enrolment, output, throughput) and its access standards.
2. **Needs (TI-01, TI-02).** The nodes become destinations or production origins. Each node with a magnitude is given a *connection standard*: the road class, surface and axle load its route must meet.
3. **Gaps (TI-05).** Two gap dimensions carry the sector: G1, people beyond the access standard for a destination class, and G6, sector nodes whose route does not meet the connection standard.
4. **Candidates (TI-07).** Every candidate project is tagged with the sectors whose gaps it helps to close.
5. **Allocation (TI-08).** Sector-focused scenarios use those tags: D (productive sectors) for agriculture, minerals and industry; E (social services) for health and education.
6. **Outcomes (TI-09).** Each sector's gain is measured in its own unit: persons, women of reproductive age, youth, households, tonnes.
7. **Outbound.** TIPRAS sends physical changes back — minutes, kilometres, persons within a time band, tonnes connected — and the sector model turns them into its own outcomes.

22.1.2 Three rules

The design sets three rules that keep the coupling honest.

- **Consume, do not model.** TIPRAS never forecasts crop output, mine production, clinic caseloads or enrolment. It takes them from the sector engines.
- **One standard, two users.** Where a sector defines an access standard, such as the 2-hour EmONC standard, both engines read the same norm row. Transport never retypes it.
- **Emit physical changes only.** TIPRAS emits minutes, kilometres, tonnes and UGX per tonne. It never emits yields, deaths averted or learning outcomes. Each sector model owns the response function that turns a physical change into its outcome and cites its source.

Table 22.1 summarises what the design expects from and returns to each sector covered in this chapter.

Table 22.1: What TIPRAS consumes from and returns to each sector (design, docs/21 §21.2)

Sector	Consumes	Gap	Scenario	Emits back
Agriculture	Production surface by crop; markets; marketed share	G1 MARKET; G6 agricultural zones	D	Transport cost per tonne, zone to market; market travel time; wet-season cut-off
Minerals	Mine sites, output, node-to-node flows	G4 axle and surface; G6; G3 rail and ports	D	Route capacity and standard; haulage cost per tonne
Health	Facility register with level and status; access standards; planned facilities	G1 HEALTH_PRIMARY and HEALTH_EMONC; G6 hospitals all-season	E	Population within 30, 60 and 120 minutes by district
Education	School register with enrolment; walking standards; planned schools	G1 SCHOOL_PRIMARY and SCHOOL_SECONDARY	E	School access by settlement, wet season
Industry	Parks, SEZs, agro-processors; freight generated	G6 industrial nodes; G3	D	Logistics cost to gateways; travel time to labour pools
Trade	Gateway flows by commodity; partner shares; growth	G3 at borders and ports; G4 corridors	D	Corridor generalised cost; border time; gateway capacity

22.2 Method

22.2.1 The sector connection test (G6)

For every sector node n with a connection standard, TI-05 finds the route from the node to its target on the current network in the wet season. For agricultural zones the target is the nearest market. For every other node it is the nearest link of the required road class. It then tests every link on that route:

$$NC_n = 1 - \prod_{l \in R(n)} \mathbf{1}[\, r(s_l) \geq r(s_n^*)\,] \cdot \mathbf{1}[\, a_l \geq a_n^*\,] \tag{22.1}$$

where

- NC_n is 1 if node n is *not connected to standard* and 0 if it is;
- R(n) is the set of links on the route from n to its target;
- r(s_l) is the rank of link l's surface in the surface order (earth, engineered earth, gravel, low-volume sealed, surface dressing, asphalt);
- s*_n is the surface the node's connection standard requires;
- a_l is the link's axle-load limit where one is recorded, and a*_n the required axle load.

A node also fails if it cannot be snapped to a wet-season link or has no route at all. The connection standard is read from `node_connection_standard.csv` by node family and magnitude band: the row with the largest minimum magnitude that the node meets. A node without a magnitude, or below the lowest band, has no standard and is not a G6 node. The amount at stake is the node's magnitude summed over the nodes that fail:

$$G6_s = \sum_{n \in N_s} m_n \, NC_n \tag{22.2}$$

where N_s is the set of sector s's nodes with a standard and m_n the node's magnitude (tonnes a year for agricultural zones).

22.2.2 Sector tags on candidate projects

TI-07 tags each candidate project j with every sector whose gap it helps to close. A project carries a tag if it belongs to at least one *enabling set* — the minimal set of projects whose joint completion brings a settlement within the access standard for a destination class — or if it lies on a link whose required standard is set by a sector node:

$$\text{tags}(j) = \left\{ \hat{s}(k) : j \in \omega,\ \omega \in \Omega_{i,k} \right\} \cup \left\{ \hat{s}(n) : l \in L_j,\ \text{binding}(l) = n \right\} \tag{22.3}$$

where Ω_{i,k} is the collection of enabling sets for settlement i and class k, ŝ(k) maps a destination class to a sector (health classes to *health*, school classes to *education*, MARKET to *agriculture*, ADMIN_CENTRE to *admin*, ALL_SEASON_ROAD to *rural_access*), and ŝ(n) is the sector of the node that binds the link's required standard. Tags are labels, not shares. A project tagged with four sectors counts in full under each of them, so sector totals cannot be added.

22.2.3 Sector-focused scenarios

The two sector scenarios rank projects by their own benefit per unit of cost. Scenario E credits the unserved population of health and school classes that a project helps to serve, shared among the unfunded members of each enabling set:

$$b_j^E = \sum_{g\ :\ k(g) \in \{HEALTH,\, SCHOOL\}} \max_{\omega \ni j} \frac{P_g}{|\omega \setminus F|} \tag{22.4}$$

where g is a settlement-class group not yet served, P_g its population in the class's target group, and F the set of projects already funded. Scenario D gives credit only to projects tagged with agriculture, minerals or industry, ordered by their net present value:

$$b_j^D = \mathbf{1}[\text{tags}(j) \cap \{ag, min, ind\} \neq \emptyset] \cdot (\max(NPV_j, 0) + 1) \tag{22.5}$$

Both scenarios then fund projects in descending order of b_j divided by central capital cost, until the pools eligible to fund them in each year are exhausted (Chapter 11 and Figure 11.1 describe the pools).

22.2.4 Gains in each sector's own unit

The gain of scenario A for destination class k at the horizon year is the number of people in the class's target group who meet the standard in that scenario but not on the do-nothing path:

$$\Delta S_k(A) = \sum_i P_{i,H}^{g(k)} \left(A_{i,k,H}^{A} - A_{i,k,H}^{0} \right) \tag{22.6}$$

where P^{g(k)} is the population of the group the standard applies to (all persons for primary health care, women aged 15–49 for EmONC, youth aged 13–18 for secondary school, households for markets) and A is 1 if the settlement meets the standard. Gains for different classes are reported side by side. They are never summed.

22.3 The state of the sector feeds

Table 22.2 sets out what each sector engine actually supplied to the reference run. It is the most important table in this chapter, because it determines how far each sector result can be trusted.

Table 22.2: What each sector feed supplied to the reference run

Sector	Design source	What the run used	Quality
Health	HealthModellingx facility register and standards	3,169 geocoded facilities from the MoH National Health Facility Register via HealthModellingx (1,492 HC II, 1,358 HC III, 202 HC IV, 102 hospitals, 15 regional referral hospitals). Standards are TIPRAS draft norms; the EmONC 2-hour standard is *normed*, the primary-care 60-minute standard *assumed*	Nodes observed; 3,631 register facilities have no coordinates. HealthModellingx keys facilities to 122 districts, not 135
Education	EducationModellingx school register	16,880 primary and 3,594 secondary schools from the MoES shapefiles in SDI, the only source with enrolment; vintage about 2018	Nodes observed; EducationModellingx has no transport channel yet
Agriculture	AgricultureModelling production surface	MapSPAM 2010 production cells clustered into 10 production zones with 9.76 Mt of marketed surplus; 814 markets from AID points of interest and OpenStreetMap	Production 2010, not current; marketed shares *assumed*; AgricultureModelling publishes production for only 4 regions
Minerals	MineralModelling flows and mine sites	309 mine points from AID points of interest, without output	MineralModelling's flows are labelled synthetic demonstration data and are not used
Industry	IndustryModelling industrial nodes	23 industrial parks and SEZs from AID points of interest, without output	IndustryModelling has no transport exchange
Trade	TradeModelling gateway flows	22 border posts from sparse OpenStreetMap tags (14 of them in Busia) and 6 airports from OurAirports	TradeModelling's transport exchange is synthetic and for The Gambia, not Uganda; not used

Three consequences follow.

- **Only health and education enter with real registers.** Both are observed nodes from government sources, with known gaps (missing coordinates, an old school vintage).
- **No sector supplies output trajectories.** Freight from sector nodes is therefore grown at a flagged fallback rate of 5 % a year (run fallback TI-03). This is one of the eight fallbacks of the run.
- **Minerals and industry enter as points without magnitudes.** The connection standards for mines are banded by annual output, starting at 100,000 t. With no output recorded, no mine reaches a band. The 309 mines therefore carry no connection standard, generate no freight, form no G6 gap and tag no project. The minerals linkage exists in code but is empty in this run.

22.4 Results: sector nodes without their required connection

TI-05 tested 178 sector nodes against their connection standard: 117 hospitals (102 general hospitals and 15 regional referral hospitals), 23 industrial parks and SEZs, 28 trade gateways and 10 agricultural production zones. The results are in Table 8.5 and Figure 8.18 of Chapter 8. Each sector is reported in its own unit.

- **Health: 103 of 117 hospitals (88.0 %) are not connected to standard.** A general hospital needs a gravel route, rated for 8.2 t axles, to the nearest road of district class or above; a regional referral hospital needs a sealed (surface-dressed) route.
- **Industry: 21 of 23 parks and SEZs (91.3 %) are not connected to standard.** An industrial node needs a sealed route rated for 11.5 t axles to the national primary network.
- **Trade: 16 of 28 gateways (57.1 %) are not connected to standard.** A gateway needs a sealed route rated for 11.5 t axles to the national trunk network.

- **Agriculture: all 10 production zones (100 %) are not connected to standard, with 9,764,387 t a year of marketed surplus at stake.** A zone with at least 5,000 t of surplus needs a gravel route to its nearest market; a zone with at least 50,000 t needs a sealed one.

These shares are high, and part of the reason is in the data rather than on the ground. The test treats a link of unknown surface as earth, the lowest rank, because the engine never assumes a surface it has not observed. Since 68.3 % of network km has unknown surface, many routes fail because one of their links is untagged. No link in the snapshot has a recorded axle-load limit either, so in this run only the surface part of the test can fail; the axle part is never tested. The G6 counts are therefore an upper bound on the true shortfall. They are also a precise survey list: the failing links on these 178 routes are the first links whose surface should be recorded.

The agricultural result is dominated by one zone. Of the 9.76 Mt of surplus, 8.47 Mt (86.7 %) lies in a single zone centred on Iganga, with sugarcane, banana and maize as its main crops. The next largest is in Yumbe, at 1.07 Mt. The other eight zones together hold about 0.23 Mt: two on Kalangala, and others in Buliisa, Kotido, Kaabong, Napak, Nabilatuk and Bukwo. The zones come from 2010 production data and assumed marketed shares, so they show where the method places surplus, not where today's surplus is.

22.5 Results: the candidate programme by sector

Figure 22.1 and Table 22.3 show how much of the candidate programme serves each sector. The programme has 41,890 candidate projects with a central capital cost of 92.8 tn UGX (range 56.6 to 125.3 tn) and induced routine maintenance of 205.2 bn UGX a year once all are complete. A project can serve several sectors, so the rows of the table overlap and must not be added.

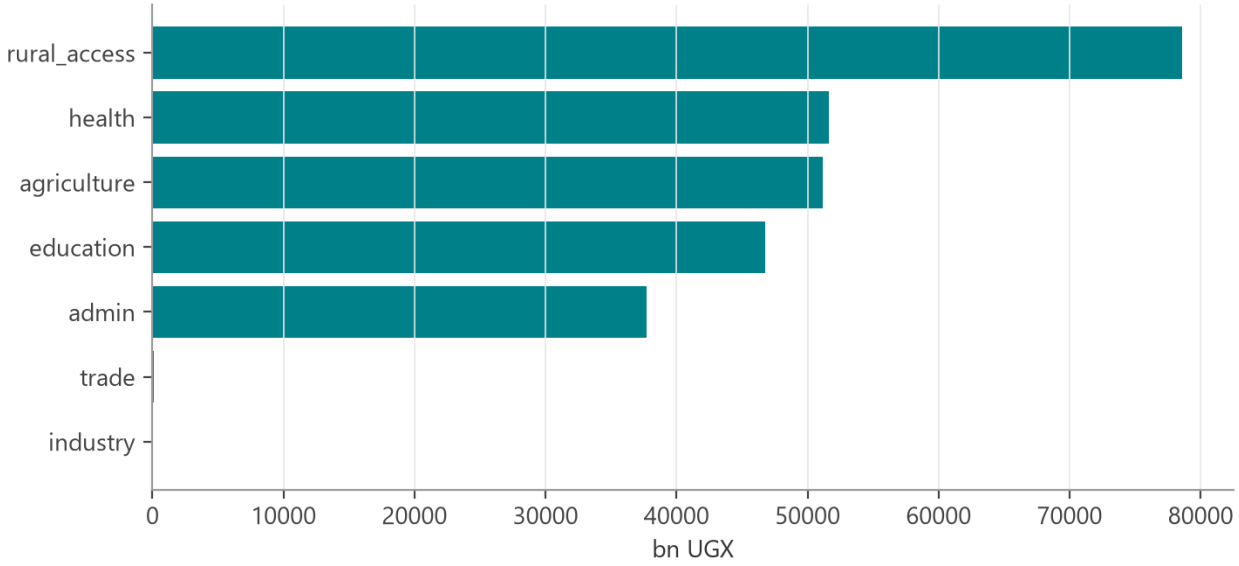

Figure 22.1: The candidate programme by sector served (a project can serve several sectors)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 22.3: Candidate programme by sector served

Sector tag	Projects	Km	Capital (bn UGX)
admin	1,847	13,593	37,700.9
agriculture	3,772	22,992	51,130.8
education	3,080	21,954	46,775.7
health	3,283	22,544	51,611.2
industry	13	8	40.8
rural_access	35,429	69,176	78,615.7
trade	7	14	75.1

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_programme_sector.csv).

The figure has three tiers.

- **Rural access dominates.** 35,429 projects, 69,176 km and 78.6 tn UGX of capital (range 47.8 to 106.8 tn), with induced routine maintenance of 194.2 bn UGX a year, help to bring rural people within 2 km of an all-season road. Almost every project with any access benefit carries this tag.
- **Four service sectors are close together.** Health (3,283 projects, 22,544 km, 51.6 tn; induced maintenance 99.7 bn a year), agriculture (3,772 projects, 22,992 km, 51.1 tn; 96.1 bn a year), education (3,080 projects, 21,954 km, 46.8 tn; 92.9 bn a year) and administrative access (1,847 projects, 13,593 km, 37.7 tn; 71.0 bn a year). The agriculture tag comes almost entirely from market access (G1 MARKET), not from the ten production zones.
- **Industry and trade are almost absent.** Industry has 13 projects, 8 km and 40.8 bn UGX of capital (induced maintenance 0.1 bn a year). Trade has 7 projects, 14 km and 75.1 bn UGX (0.2 bn a year). These are short connector upgrades on routes that bind on an industrial or gateway node standard. Two of the trade projects lie in Kanungu, near the Savannah airstrip: one 7.8 km gravel-to-sealed upgrade costing 41.1 bn UGX, also tagged for administrative, agricultural and education access, and a 1.3 km upgrade costing 7.0 bn UGX.

Every service-sector tag is far smaller than the rural-access tag. That is because the access standards for services are measured in travel time and allow motorised travel, while the RAI is measured as walking distance to an all-season road. A settlement can be within an hour of a health centre on an earth road and still be counted as without rural access.

22.5.1 Projects without a sector tag

6,130 projects (9.8 tn UGX of capital, 6.8 bn UGX a year of induced maintenance) carry no sector tag. Most are upgrades where the required standard is set by the road's function or its traffic rather than by access to a destination. There are also 72 rehabilitation and 38 periodic maintenance condition programmes, priced in tranches of 1,000 km from agency statistics. These projects serve the network, not a sector.

22.5.2 Multi-sector candidates

The design asks the interface to list candidates that serve three or more sectors, with each sector's gain shown in its own unit. Counting only the four service sectors (health, education, agriculture through markets, administrative access), the programme contains:

Table 22.4: Candidate projects by the number of service sectors they serve

Service sectors served	Projects	Km	Capital, central (bn UGX)	Induced routine maintenance (bn UGX/yr)	Projects with NPV > 0
None (rural access or network only)	35,684	146,440	25,496.9	76.8	269
One	2,843	10,343	12,371.0	22.0	85
Two	1,570	7,941	12,850.0	24.3	72
Three	1,131	8,525	19,028.1	38.7	85
All four	662	7,352	23,015.8	43.2	63

Table 22.4 shows that 1,793 projects serve three or four service sectors at once. They cover 15,877 km and cost 42.0 tn UGX, with 81.9 bn UGX a year of induced maintenance. The 662 projects that serve all four average 11.1 km, against a programme average of 3.9 km, and cost 3.1 bn UGX per km against 0.5 bn UGX per km for the programme as a whole. They are long upgrades to a sealed standard on routes that connect whole sub-counties to a health centre, a secondary school, a market and the district headquarters. These are the natural candidates

for joint financing across sector ministries. They are listed in the interface with each sector's gain beside them. The gains are not added into a score.

22.6 Results: funding by sector in each scenario

Figure 22.2 and Table 22.5 show how much capital each scenario funds in projects tagged for each sector over the MTEF window. Rows are not additive across sectors, because a project can carry several tags. Scenarios are compared here; none is preferred.

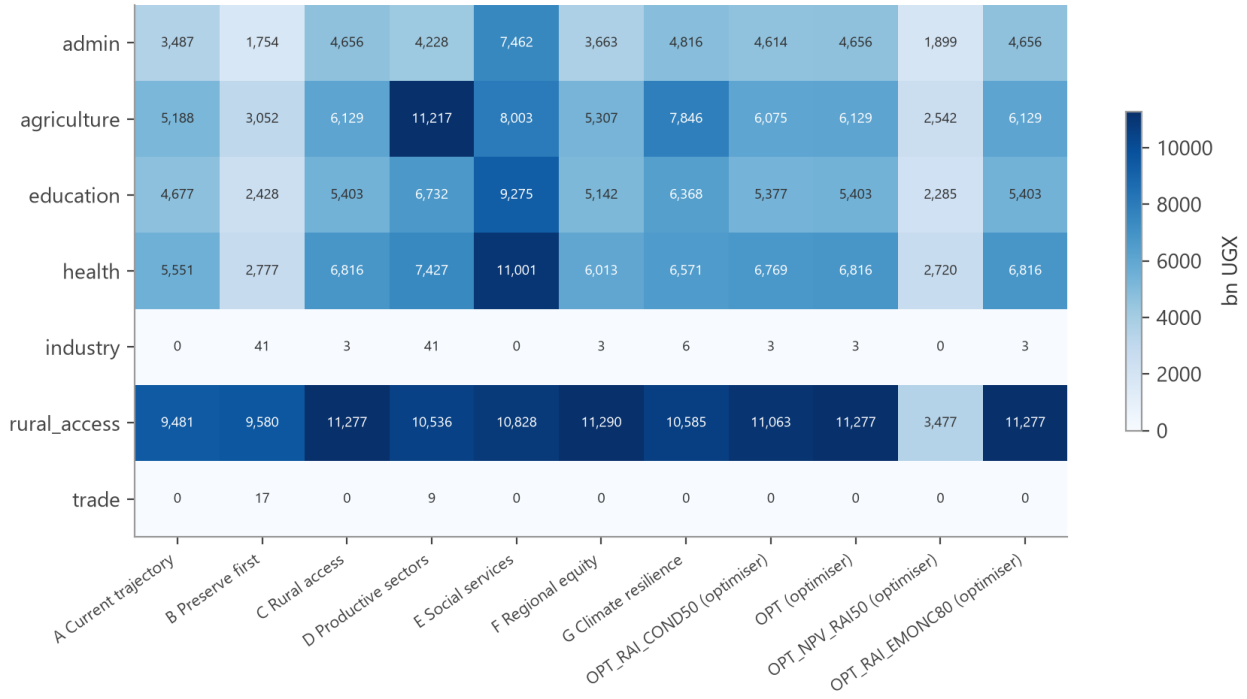

Figure 22.2: Capital funded per sector served, by scenario (bn UGX; rows are not additive)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Table 22.5: Capital funded per sector served, by scenario (bn UGX; a project can serve several sectors, so rows are not additive)

Sector tag	A	B	C	D	E	F	G	OPT_RAI_COND50	OPT	OPT_NPV_RAI50	OPT_RAI_EMONC80
admin	3,486.9	1,754.4	4,656.0	4,228.4	7,461.8	3,663.1	4,815.8	4,614.3	4,656.0	1,899.3	4,656.0
agriculture	5,187.8	3,052.2	6,129.3	11,217.3	8,002.9	5,307.5	7,845.8	6,074.9	6,129.3	2,542.4	6,129.3
education	4,676.8	2,427.7	5,403.4	6,732.1	9,274.6	5,142.4	6,368.1	5,376.9	5,403.4	2,285.0	5,403.4
health	5,550.6	2,776.5	6,816.4	7,426.9	11,000.6	6,012.6	6,571.4	6,769.4	6,816.4	2,719.8	6,816.4
industry	0.0	40.8	2.7	40.8	0.0	2.7	5.6	2.7	2.7	0.0	2.7
rural_access	9,480.8	9,579.7	11,277.4	10,535.8	10,827.9	11,290.3	10,584.9	11,063.3	11,277.4	3,477.0	11,277.4
trade	0.0	17.0	0.0	9.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/t_alloc_sector.csv).

The heat map has a clear structure.

- **Rural access is the largest row in every scenario but one.** It ranges from 9,480.8 bn UGX in A to 11,290.3 bn UGX in F. The exception is OPT_NPV_RAI50 (maximising net present value with an RAI floor at 50 % of its attainable maximum), which funds only 3,477.0 bn UGX of rural-access projects, because it funds 3,658.9 bn UGX in all.
- **Scenario D (productive sectors) funds the most agriculture.** It allocates 11,217.3 bn UGX to projects tagged for agriculture, more than twice the 5,187.8 bn of A and 1.8 times the 6,129.3 bn of C. This is the darkest cell in the agriculture row.
- **Scenario E (social services) funds the most health and education.** It allocates 11,000.6 bn UGX to health-tagged projects and 9,274.6 bn to education-tagged projects, against 6,816.4 bn and 5,403.4 bn in C. It also funds the most administrative access (7,461.8 bn), because projects that bring people to a health centre often also bring them to the district headquarters.
- **Scenario B (preserve first) funds the least for every service sector.** It spends most of its capital on the priced rehabilitation and periodic maintenance programmes, which carry no sector tag. Its health row is 2,776.5 bn and its education row 2,427.7 bn.
- **Industry and trade rows are nearly empty in every scenario.** Industry receives at most 40.8 bn UGX (B and D, which fund all 13 industry-tagged projects). Trade receives at most 17.0 bn UGX (B). C, E, F, G and the optimiser runs fund no trade-tagged project.
- **The optimiser runs with an access objective look like C.** OPT (maximising people served within 2 km of an all-season road) gives exactly the same allocation as scenario C. OPT_RAI_EMONC80 is the same again, because its floor on EmONC access could not be met: it reached 17,789 additional women within 2 hours of EmONC, 92.3 % short of the floor of 230,623.

22.6.1 Induced maintenance by sector and scenario

Every capital figure carries its maintenance tail. Table 22.6 gives the steady-state routine maintenance induced by the projects each scenario funds, by sector tag. Like the capital rows, the sector rows overlap and are not additive. The last row is the scenario total.

Table 22.6: Induced routine maintenance of funded projects, by sector tag and scenario (bn UGX a year at steady state; rows overlap)

Sector tag	A	B	C	D	E	F	G	OPT_RAI_COND50	OPT_NPV_RAI50
Administrative access	6.6	2.8	9.8	8.6	15.4	7.4	9.9	9.7	4.0
Agriculture	10.1	4.5	13.0	22.6	15.8	10.9	15.9	12.9	5.4
Education	9.5	3.2	10.9	13.7	18.4	10.0	13.2	10.9	4.8
Health	11.4	4.0	14.1	15.1	21.9	12.1	13.7	14.1	5.8
Industry	0.0	0.1	0.0	0.1	0.0	0.0	0.0	0.0	0.0
Rural access	18.9	56.5	28.3	21.6	21.7	28.6	29.1	26.6	7.5
Trade	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
All funded projects	20.5	62.7	28.3	22.9	22.2	28.6	30.6	26.6	7.9

OPT and OPT_RAI_EMONC80 have the same values as C. Two points stand out. First, B induces the most maintenance (62.7 bn UGX a year), because its upgrades are spread over the most km. That is the figure carried into the fiscal sustainability statement of Chapter 13. Second, the sector scenarios move the maintenance tail as much as the capital: E adds 21.9 bn UGX a year to health-tagged roads, D adds 22.6 bn UGX a year to agriculture-tagged roads. A sector that asks for a road also asks the district maintenance budget for its upkeep.

22.7 Results: what each sector gains, in its own unit

Table 22.7 gives each scenario's gain at the 2030 horizon over the do-nothing path, for each destination class, in the unit of the class's target group. The classes are reported side by side and never added.

Table 22.7: Gain over the do-nothing path at the 2030 horizon, by destination class (each in its own unit)

Scenario	All-season road, 2 km (persons)	Primary health, 60 min (persons)	EmONC, 2 h (women 15–49)	Secondary school, 60 min (youth 13–18)	Market, 2 h (households)
A	4,590,914	146,440	94,863	63,183	170,440
B	21,313,562	439,773	136,849	122,565	288,666
C	22,305,059	369,584	188,989	113,563	324,248
D	7,815,069	367,233	144,091	110,644	479,809
E	8,698,487	786,116	307,525	186,082	374,423
F	21,911,188	373,599	165,332	122,134	316,958
G	6,009,699	212,191	100,102	69,036	285,071
OPT_RAI_COND50	21,999,728	355,044	185,403	110,856	318,550
OPT	22,305,059	369,584	188,989	113,563	324,248
OPT_NPV_RAI50	10,186,062	111,158	63,365	41,880	97,068

On the do-nothing path in 2030, 21.0 million people are within 2 km of an all-season road, 54.9 million within 60 minutes of primary health care, 12.6 million women within 2 hours of EmONC, 7.8 million youth within 60 minutes of a secondary school and 10.4 million households within 2 hours of a market.

The table shows that each sector has its own leading scenario, and that no scenario leads for all.

- **E gives the largest health and education gains.** It brings 786,116 more people within 60 minutes of primary health care, 307,525 more women within 2 hours of EmONC and 186,082 more youth within 60 minutes of a secondary school — more than twice C's primary-health gain and 1.6 times its EmONC gain. It gives only 8.7 million in rural access, against 22.3 million in C.
- **D gives the largest market gain.** It brings 479,809 more households within 2 hours of a market, 48 % more than C. It gives only 7.8 million in rural access.
- **C and OPT give the largest rural-access gain.** 22.3 million more people within 2 km of an all-season road, but smaller service gains than E or D.
- **B is close to C on rural access (21.3 million) and second on primary health (439,773).** It also cuts poor non-urban km from 107,837 to 11,204, which no access-focused scenario does (Chapter 12).

These are exactly the trade-offs a sector ministry needs to see. A health ministry comparing E with C sees 416,532 more people within an hour of primary care and 118,536 more women within two hours of EmONC, at the price of 13.6 million fewer people with rural access. The choice between them is a policy decision; the engine only shows its terms.

Even the largest EmONC gain is small against the gap. The EmONC gap in the base year is 1.04 million women, of whom 614,403 are beyond reach because there is no qualifying facility within 2 hours even on an upgraded network. Roads alone cannot close that part. It is a facility-siting finding, and it is passed to HealthModellingx.

The agricultural production zones show no gain in any scenario. All ten remain not connected to standard in the outbound file, because the zone routes bind on unknown-surface links that no scenario upgrades. Mineral, industrial and trade gains are not reported, because no magnitudes or flows were supplied for these nodes.

22.8 Sector by sector

22.8.1 Agriculture

Figure 22.3 shows the 814 crop markets as points over the district market-access gap (G1 MARKET, in households beyond 2 hours of a market). Darker districts have larger gaps.

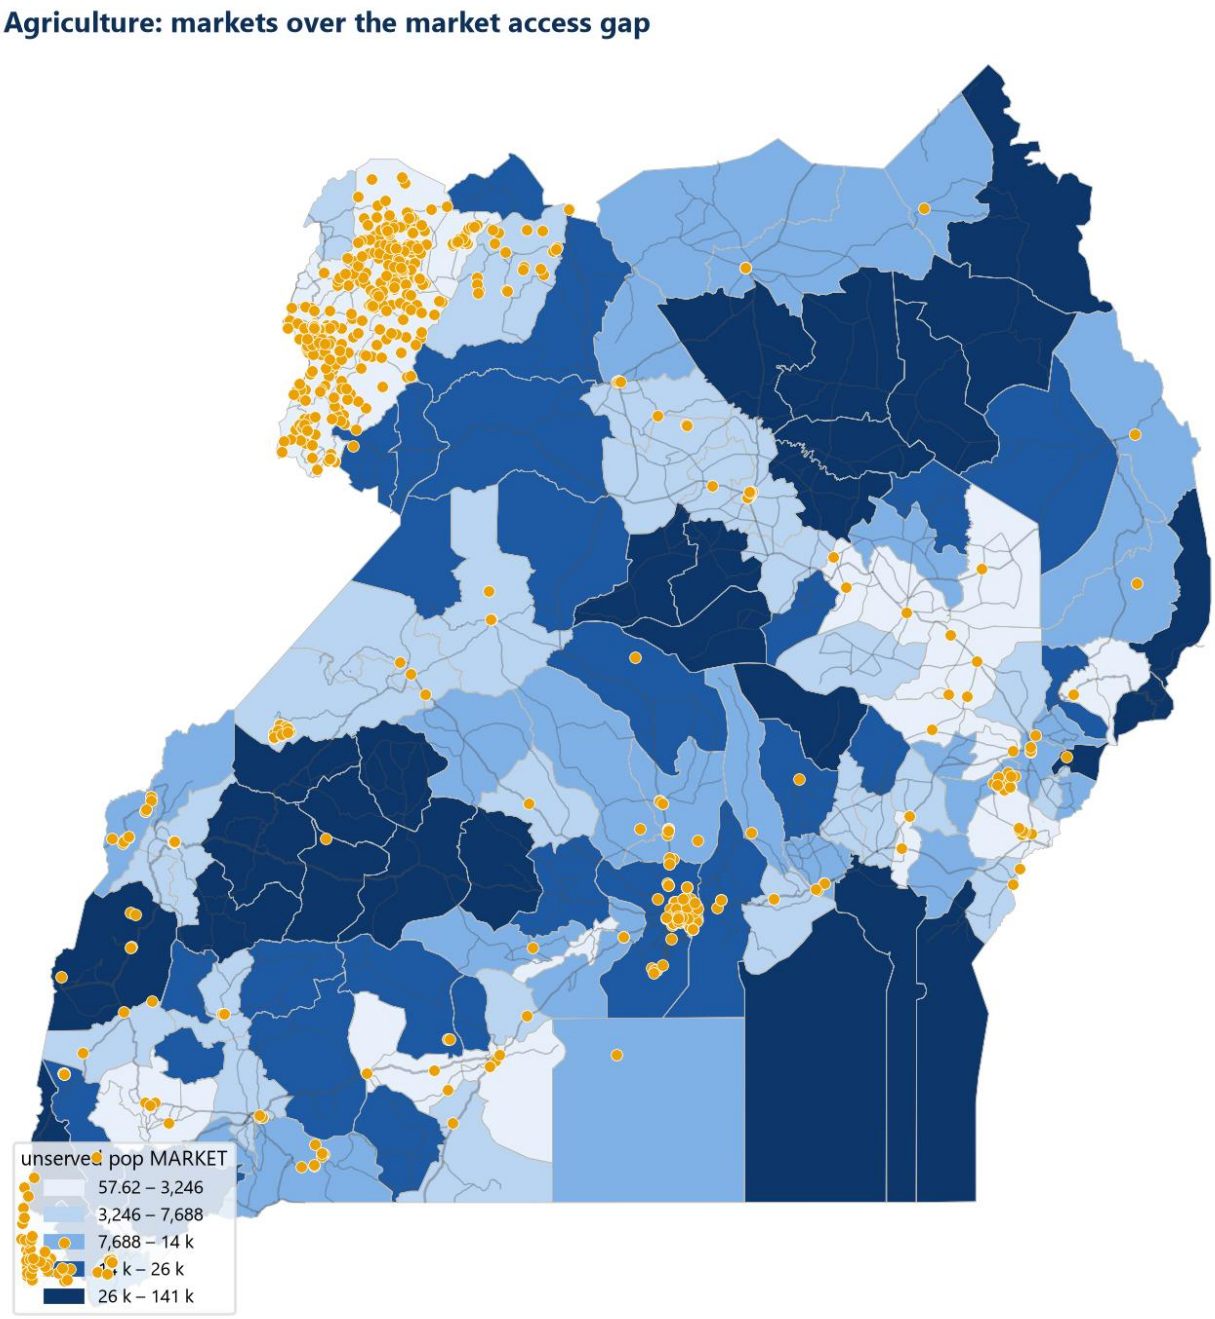

Figure 22.3: Agriculture: markets over the market access gap (each sector's gap in its own unit)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Two patterns are visible. First, the market register is very uneven. West Nile is thick with points: Arua has 149 markets and Yumbe 145, against 115 in Kampala. Much of the mid-west and north-east has only a handful. This is a mapping effect in the source data (AID points of interest and OpenStreetMap marketplaces), not a survey of trading activity. Second, the darkest districts are in the mid-west and the north-east. Mubende (141,370) and Kakumiro (140,790) have the largest gaps, followed by Agago (67,794), Apac (61,026), Kibaale (56,286) and Kagadi (55,486). The mid-west block is the HH cluster of the market gap found in the GIS analysis of Chapter 21. Of the national market gap of 2.33 million, 1.52 million is *no_destination*: no market is within 2 hours even on an upgraded network. Part of that may be markets missing from the register rather than missing on the ground.

The agricultural chain is built: production zones from gridded production, a logit split of each zone's surplus across markets within reach, and a gravity distribution from markets to consumption centres and gateways. In the reference run, 9.76 Mt of surplus reached 9 markets in the first stage, and the gravity model balanced 9.69 Mt

across consumption centres in 500 iterations. The gravity parameter (β = 0.2 per hour) and the market logit parameter (θ = 0.5 per hour) are assumed, not calibrated.

22.8.2 Minerals

Figure 22.4 shows the 309 mine points over G4 km below the required surface. The mines are in 48 districts, with large concentrations in Wakiso (100 points), Nebbi (24), Kasese (23), Pader (19), Buliisa (16) and Mukono (11).

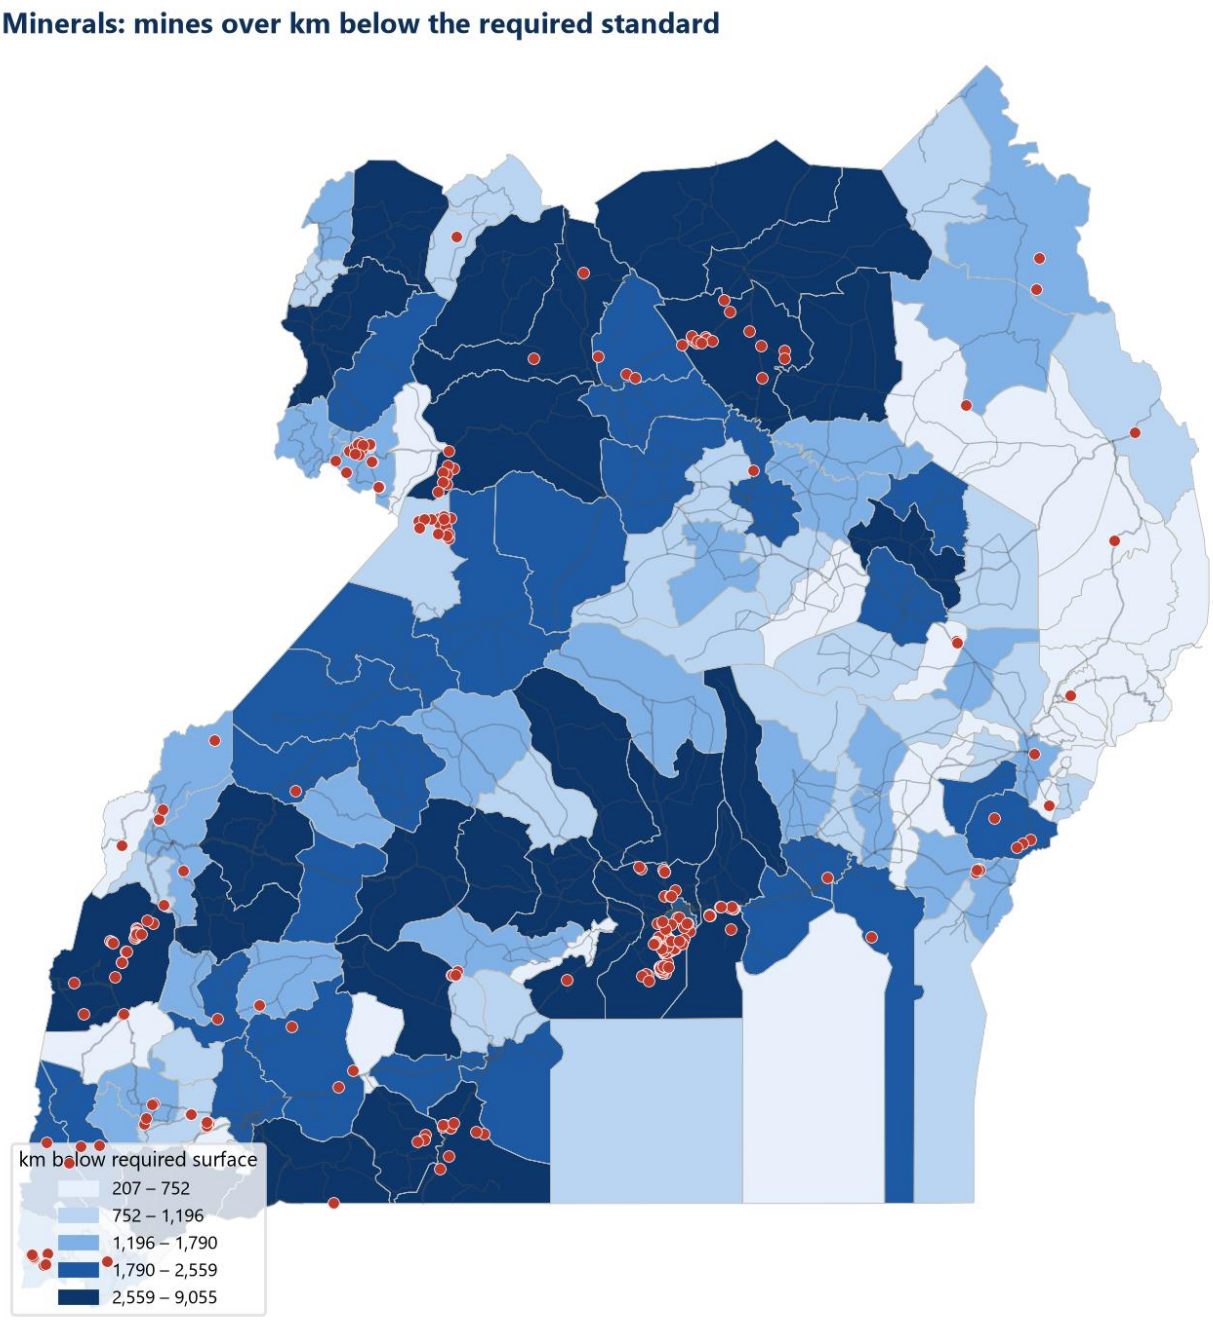

Figure 22.4: Minerals: mines over km below the required standard (each sector's gap in its own unit)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The map can show where mines sit, but not what they need. The points come from the AID points of interest without output, ownership or status, so each is a site, not a production node. The points carry no mineral type, so a quarry cannot be told from a mine. No connection standard can be assigned without output, and no mineral freight is generated. The underlying layer is a G4 count of km, which is largest in the populous central districts

(Wakiso, 9,055 km; Mukono, 5,463 km) and in the north and mid-west. It is not a measure of the mines' route standard.

What is needed is a register of mines with output from the Directorate of Geological Survey and Mines, or real flows from MineralModelling. MineralModelling's `minerals_to_transport_flows.csv` exists but is labelled synthetic demonstration data, so it is not loaded as a Uganda result. When real flows arrive, TI-03 will take them as given, skip the gravity step for minerals and test each mine's export route against the axle and surface standard for its output band.

22.8.3 Health

Figure 22.5 shows the 3,169 geocoded health facilities over the primary-health access gap (G1 HEALTH_PRIMARY, persons beyond 60 minutes of a facility at HC II level or above in the wet season).

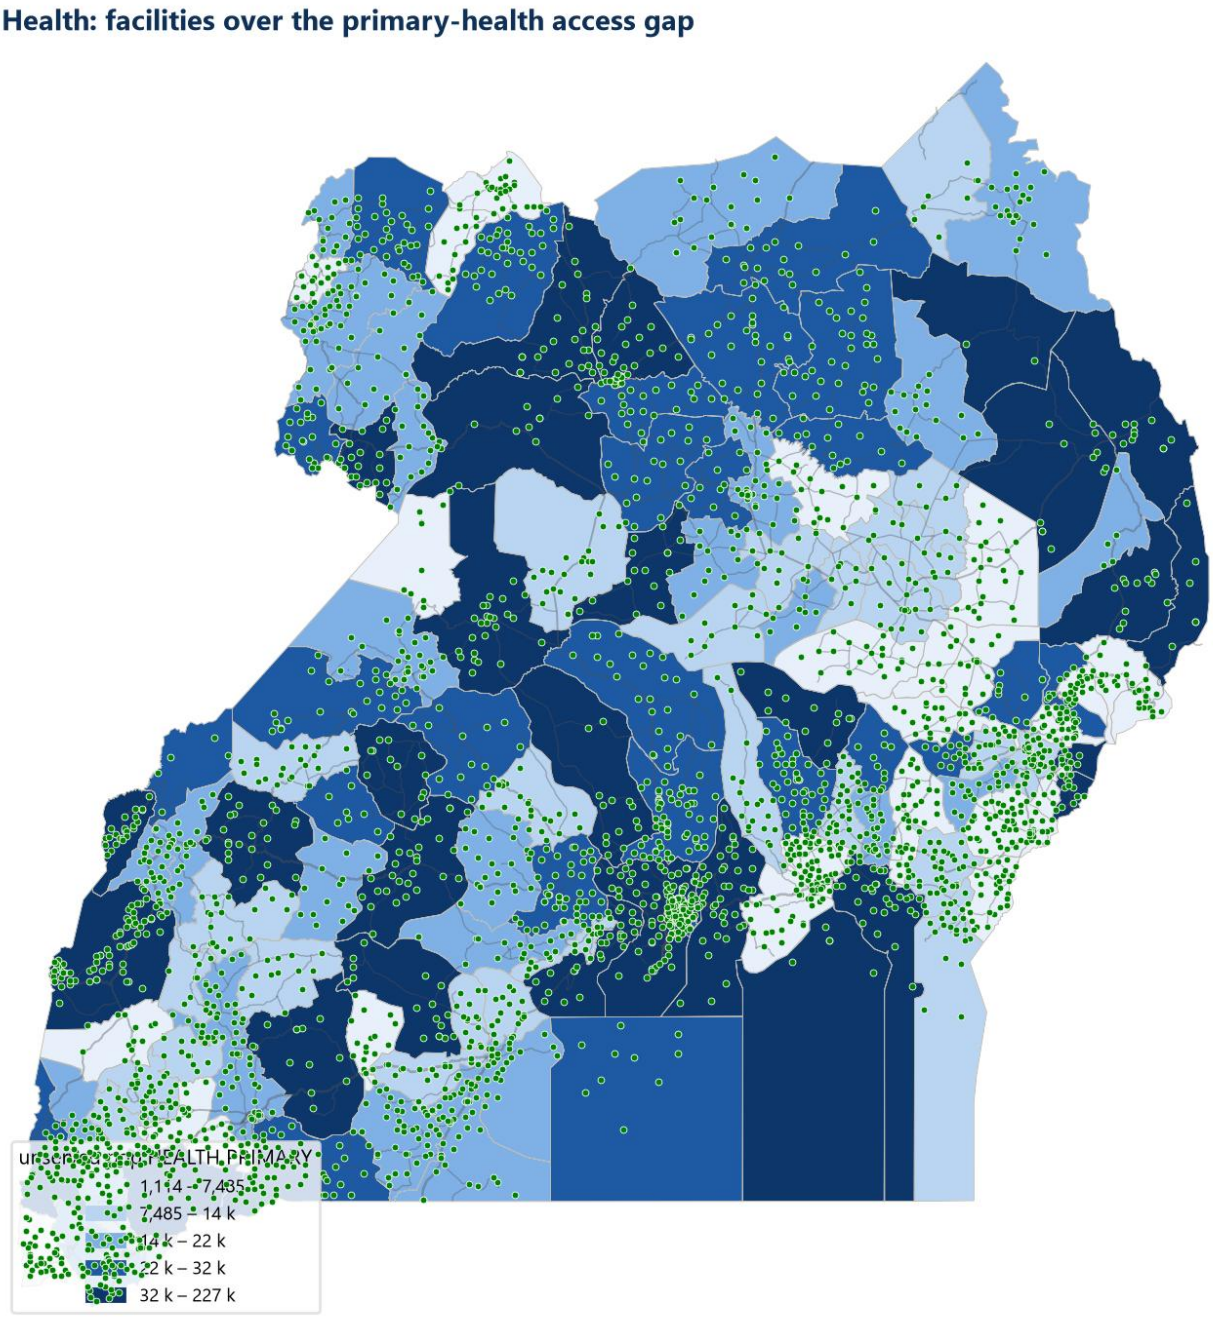

Figure 22.5: Health: facilities over the primary-health access gap (each sector's gap in its own unit)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

Facilities are dense in the south-west, around Kampala and on the Mount Elgon slopes, and sparse in Karamoja and across parts of the north and mid-west. The darkest districts — the largest gaps — include Kotido (227,078 persons), Kasese (105,665), Buvuma (77,175), Amudat (74,471), Bududa (71,074) and Mubende (68,700). Kasese and Bududa show that a large gap can occur where facilities are not scarce: mountain terrain and the wet season stretch travel times beyond the standard. Kotido's gap is mostly a lack of facilities.

The cause split matters for the health sector. Of the 3.21 million people beyond the primary-health standard, 2.36 million are *no_destination*: no qualifying facility is within 60 minutes even on an upgraded network. Only 0.85 million could be brought within the standard by roads. For EmONC, the split is 614,403 *no_destination* and 420,690 *network_substandard*. HealthModellingx receives this split as a siting finding: for most of the health gap, a road will not fix it; a facility or outreach will.

Before any health figure enters a plan, the design asks for a check of the facility register against SDI's health accessibility layer (gate GIS-6). The reference run passed the register through TIPRAS's own checks, but the SDI comparison is not yet reported. Health access is the best-supported sector result in this chapter. TI-04 found that 93.8 % of the population requiring primary health care meets the standard in the wet season, and 91.5 % of women 15–49 meet the EmONC standard.

22.8.4 Education

Figure 22.6 shows the 20,474 schools over the secondary-school access gap (G1 SCHOOL_SECONDARY, youth aged 13–18 beyond 60 minutes of a secondary school).

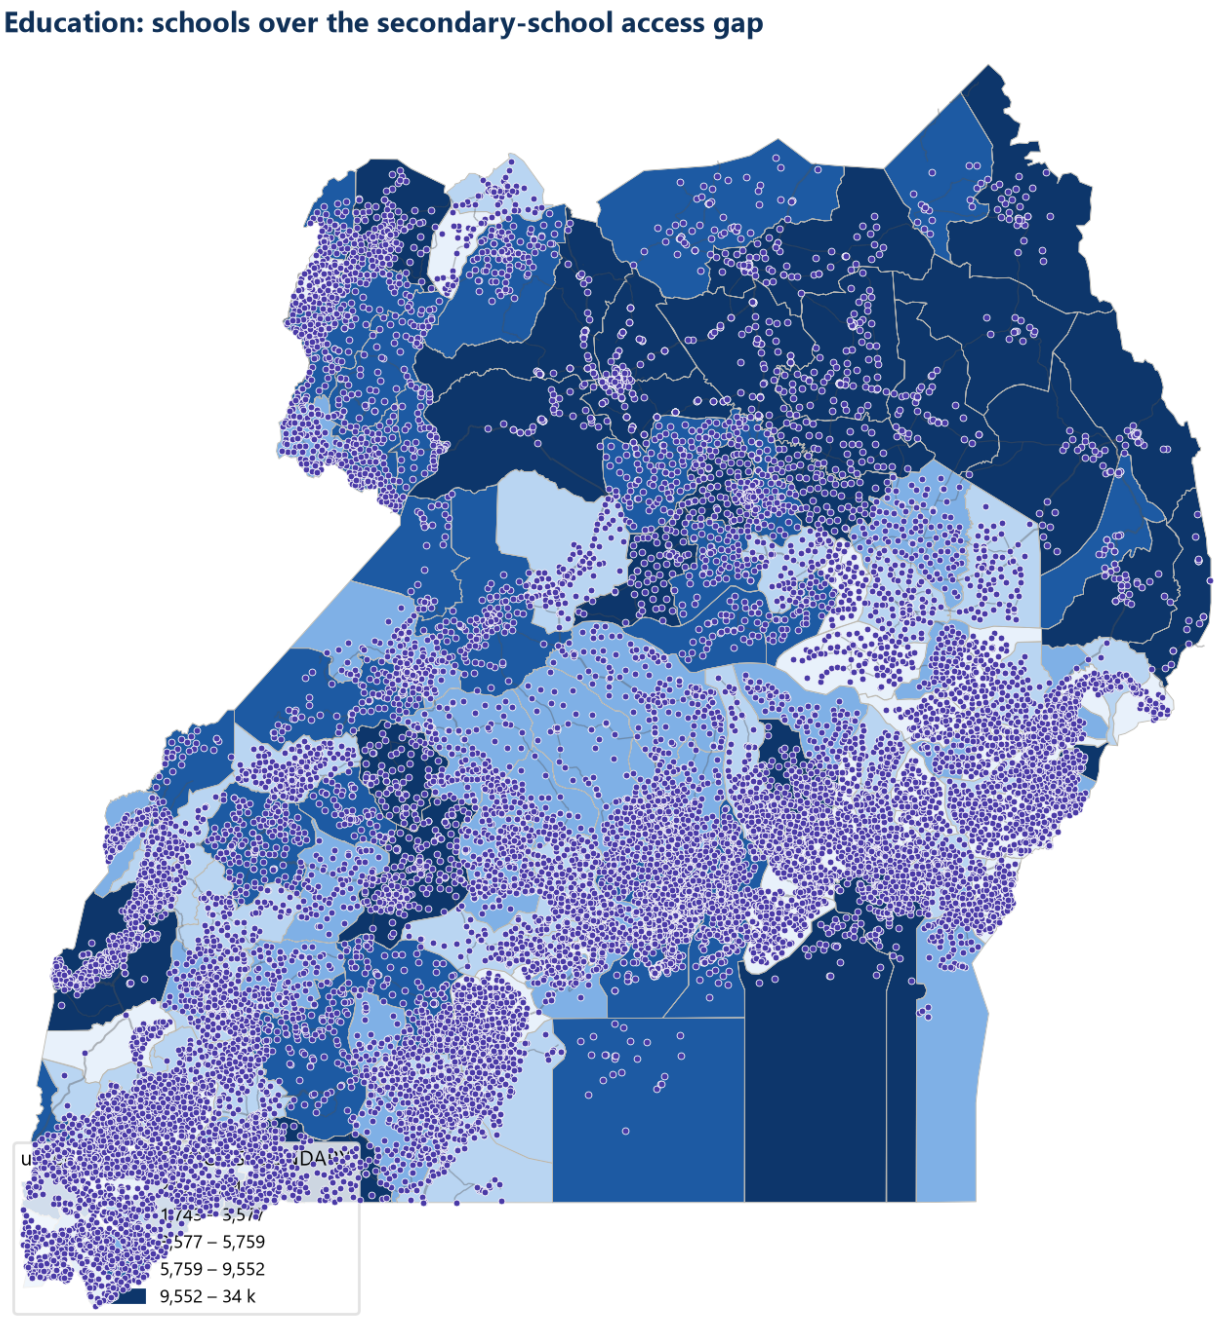

Figure 22.6: Education: schools over the secondary-school access gap (each sector's gap in its own unit)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The map is dominated by the density of the school register. Primary schools cover the south, the centre and the east almost continuously. The darkest districts lie across the north and north-east: Alebtong (33,976 youth), Amudat (25,798), Otuke (25,600), Agago (23,163), Kotido (23,161) and Napak (20,644). These form the largest LISA hot spot in the report (14 districts; Chapter 21). The national secondary-school gap is 860,840 youth, of whom 591,642 are *no_destination*.

The primary-school gap is different in kind. The standard is a 3 km walk in the wet season, and 2.86 million of the 2.94 million children beyond it are *no_destination*. A road cannot shorten a walk to a school that is too far away. The fixes that do belong to transport are footbridges and vented fords on seasonal crossings, which the decision rule places before upgrades. 1,958 river crossings on the network have no recorded structure. They are listed for survey in Chapter 8 (Table 8.4); they are not assumed to be passable.

22.8.5 Industry

Figure 22.7 shows the 23 industrial parks and SEZs over G4 km below the required surface.

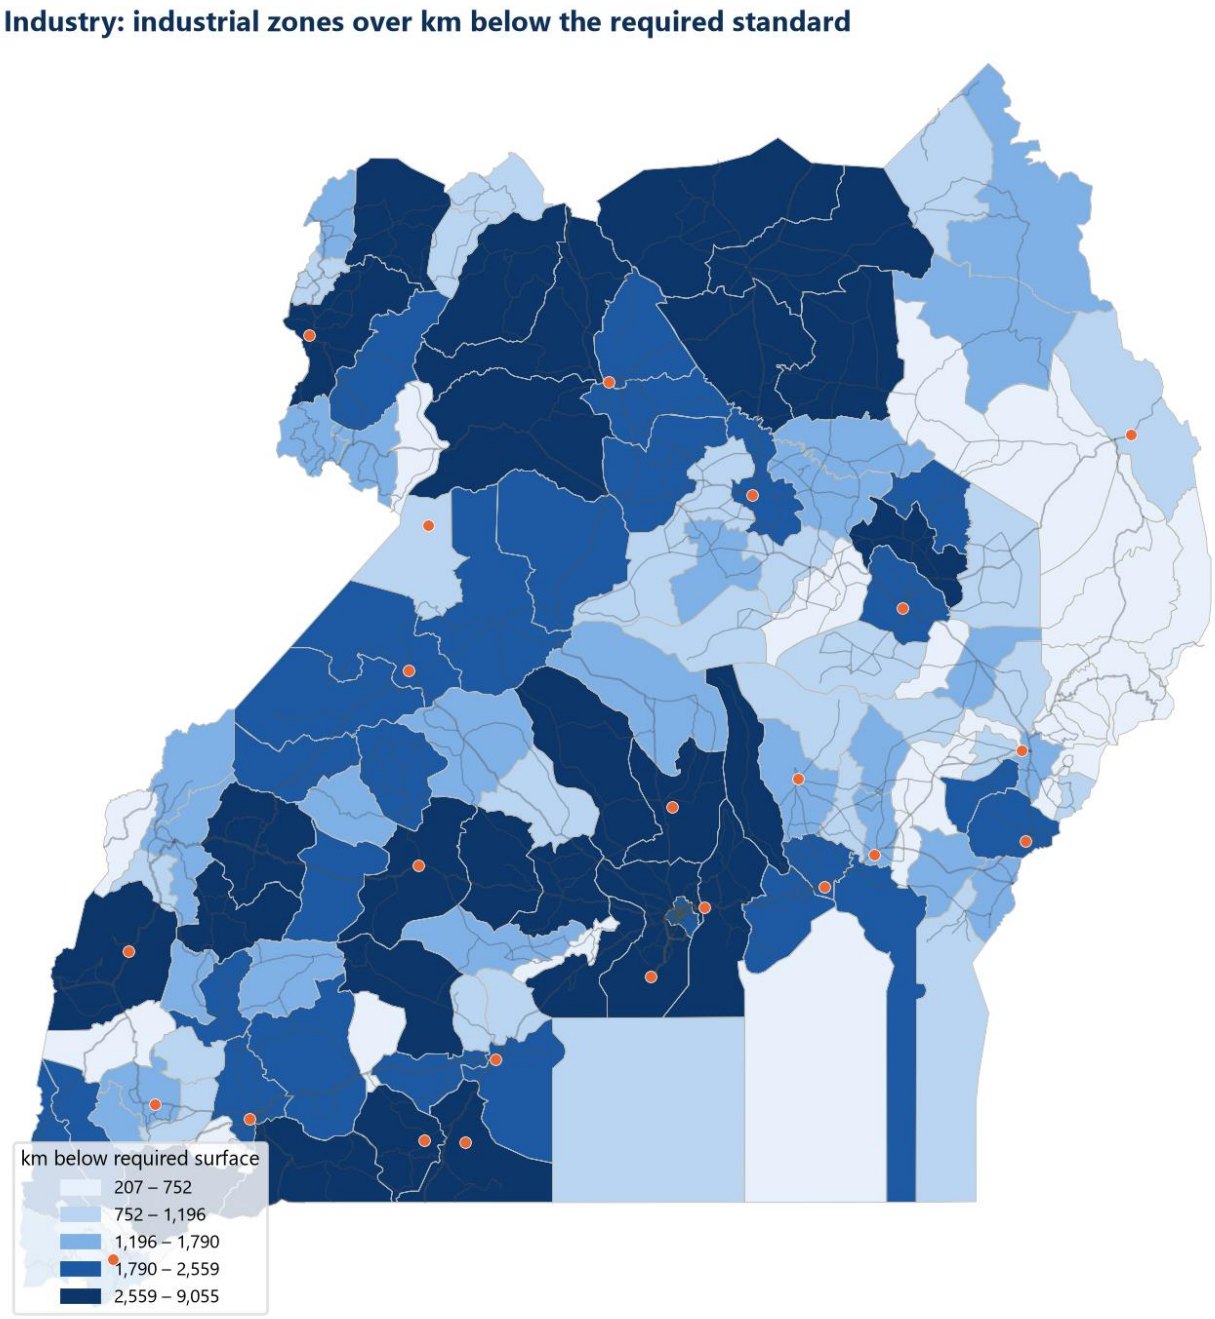

Figure 22.7: Industry: industrial zones over km below the required standard (each sector's gap in its own unit)

Source: AICTEM TIPRAS engine, reference run 014c9cd7

The parks are spread across 23 districts, one each, from Arua and Gulu in the north to Kabale, Mbarara and Rakai in the south-west and Mbale and Tororo in the east. The register comes from AID points of interest and needs checking: the point named "Kabaale Industrial Park" falls in Kabale district, while the industrial park of that name associated with the oil refinery is in Hoima district. Of the 23 parks, 21 lack a sealed, 11.5 t route to the national primary network. The programme contains 13 short upgrades (8 km in all, 40.8 bn UGX, with 0.1 bn UGX a year of induced maintenance) whose standard is set by an industrial node. The park connections are the cheapest sector connections in the programme. Their benefit cannot be valued until IndustryModelling supplies output and freight for each park.

22.8.6 Trade

The 28 trade gateways are 22 border posts and 6 airports: Entebbe International, Arua, Gulu, Soroti, Kisoro and the Savannah airstrip in Kanungu. The border register is sparse and uneven: 14 of the 22 border points lie in Busia district, and the remaining 8 are spread over six districts (Kasese and Tororo with two each; Kabale, Mbale, Koboko and Amuru with one each). Sixteen of the 28 gateways lack a sealed, 11.5 t route to the national trunk network. The programme contains 7 trade-tagged projects (14 km, 75.1 bn UGX, with 0.2 bn UGX a year of induced maintenance). Without TradeModelling's gateway flows for Uganda, no corridor or border capacity gap (G3 at gateways) can be computed.

22.9 What TIPRAS sends back

TI-10 writes the outbound exchange files to `data/exchange/outbound/`. Each row carries the country, run id, scenario, schema version, the producer (TIPRAS), an *exploratory* flag and a basis. Table 22.8 lists them.

Table 22.8: Outbound files written for the sector engines

Receiver	File	Content	Rows
HealthModellingx (HPRAS)	`transport_access_UGA_<date>.csv`	By district: population and population within 30, 60 and 120 minutes of a health facility, mixed mode, by year and scenario, in the schema HPRAS reads for its access module	135 districts
EducationModellingx (EPRAS)	`transport_access_SCHOOL_PRIMARY_…` and `…_SCHOOL_SECONDARY_…`	By settlement: wet-season travel time to the nearest school and whether it meets the standard	175,367 settlements per file
AgricultureModelling	`transport_cost_UGA_<date>.csv`	By production zone: surplus, crop mix, travel time to market on the required network, connected at standard	10 zones
FINEX	`transport_requirement_UGA_<date>.csv`	Capital requirement by fiscal year, economic code and financing pool	18 rows
MTEF	`transport_mtef_UGA_<date>.csv`	The plan in MTEF format	—

Three planned channels are not yet written. MineralModelling's route capacity file, IndustryModelling's logistics cost file and TradeModelling's gateway cost file all need inbound magnitudes or flows first. The agricultural file does not yet carry a cost per tonne, only travel time and the connection flag, because vehicle operating costs are uncalibrated (Chapter 23). The outbound files in the folder are written per run. The latest set is from run 0bad432d, a later run than the reference run 014c9cd7; the agricultural zone list is the same in both.

None of these files is a sector outcome. HPRAS receives population within a travel-time band, not lives saved. EPRAS receives minutes and a yes/no standard flag, not enrolment. AgricultureModelling receives travel times, not prices or yields. The response belongs to the sector engine.

22.10 Data quality and limitations

- **Two sectors are real, four are placeholders.** Health and education enter with government registers. Agriculture enters with 2010 production data and assumed marketed shares. Minerals, industry and trade enter as points without magnitudes. Every result for the last four is a demonstration of the method, not a Uganda finding.
- **Unknown surface drives the G6 counts.** A route with one untagged link fails its standard. The shares of 88 %, 91 %, 57 % and 100 % not connected are upper bounds until the routes are surveyed.

- **Access standards are mostly draft norms.** The EmONC 2-hour standard is *normed* (WHO/Lancet). The 60-minute primary-health standard, the 3 km primary-school walk, the 60-minute secondary-school standard and the 2-hour market standard are *assumed* and await confirmation by the sector ministries. The connection standards for every node family are illustrative.
- **Group populations use national shares.** Women 15–49 (23.5 % of the population), children 6–12 (21 %), youth 13–18 (15 %) and households (one per 4.5 persons) are applied uniformly to every district. District age structure from UBOS would change every sector's gap.
- **Registers carry mapping bias.** Markets are over-represented in West Nile, borders in Busia. Schools date from about 2018. 3,631 health facilities lack coordinates. Some G1 *no_destination* gaps may be register gaps.
- **No sector switch-off run yet.** The design asks for runs with each sector disabled, so that each linkage's contribution is visible as the difference between the enabled and disabled runs. These runs have not been made.
- **Sector tags are labels, not shares.** A project tagged for four sectors counts in full under each. Sector rows of capital and maintenance cannot be added, and no sector "owns" a share of a multi-sector project.

22.11 Implications for planning

- **Show each sector its own gains, side by side.** E gives the largest health and education gains, D the largest market gain, C and the optimiser the largest rural-access gain. A sector ministry should see the scenario comparison in its own unit, without a combined score.
- **Use the multi-sector candidates as the agenda for joint financing.** 1,793 projects serve three or four service sectors at once, at 42.0 tn UGX of capital and 81.9 bn UGX a year of induced maintenance. They are where health, education, agriculture and local government have a shared interest in the same road.
- **Separate what roads can fix from what they cannot.** Most of the primary-health, EmONC and primary-school gaps are *no_destination*. The siting findings should go to HealthModellingx and EducationModellingx, and the transport plan should not be judged on gaps it cannot close.
- **Survey the 178 sector routes first.** The G6 shortfall is inflated by unknown surface. Recording surface and structure on the routes from hospitals, parks, gateways and production zones to their targets would turn an upper bound into a measured gap at low cost.
- **Request magnitudes from the productive sectors.** Mines need output from the Directorate of Geological Survey and Mines, parks need output and freight from IndustryModelling or the Uganda Investment Authority, and gateways need flows from TradeModelling for Uganda. Until they arrive, the minerals, industry and trade linkages remain empty.
- **Budget the maintenance tail with the sector.** E adds 21.9 bn UGX a year of routine maintenance on health-tagged roads and D 22.6 bn UGX a year on agriculture-tagged roads. A sector that asks for roads should see what it commits district budgets to maintain.
- **Confirm the sector access standards.** Only the EmONC standard is normed. The sector ministries' confirmation of the other standards is a precondition for treating any sector gap in this chapter as more than exploratory.

CHAPTER 23

23 TransportModelling reuse

Which parts of the existing TransportModelling engine does TIPRAS reuse, which values did it take from that engine's configuration, and how far can those values be trusted for Uganda? This chapter answers those questions. It implements design chapter 22 (`docs/22_TRANSPORTMODELLING_REUSE.md`), which assessed TransportModelling component by component in September 2026 and gave each one a verdict: port as it is, port and adapt, consume its results through an exchange, take its parameters or method only, or do not copy. The chapter compares those verdicts with what is now in `src/`, and with the norm rows seeded from TransportModelling's configuration. The short answer is that TIPRAS took the *methods* of about a dozen components and the *parameter values* of six configuration blocks. Every one of those values is marked UNCALIBRATED by its own source file, and several are explicitly "not Ugandan". They carry the quality tier *assumed* in every run, and the vehicle operating cost values in particular make every economic appraisal in this report exploratory.

23.1 What TransportModelling is

TransportModelling (`D:\EDIP\models\TransportModelling\AICTEM_Transport_Infrastructure`) is a Julia package for transport engineering and economic appraisal. When design chapter 22 assessed it, it had 26 build phases and 2,701 `@test` lines expanding to about 3,536 assertions. It had been run on Uganda's OpenStreetMap network of about 825,000 links. It is strict about provenance, and it refuses missing or illustrative inputs unless the caller opts in.

It does engineering and appraisal well. It does not do what TIPRAS does: the chain from norms-based needs to gaps, costs and allocation. The two engines therefore overlap in their building blocks (graphs, travel time, deterioration, appraisal, optimisation) but not in their purpose. Design chapter 22 used five verdicts to decide what to take, and Table 23.1 restates them.

Table 23.1: The five reuse verdicts of docs/22 §22.1

Verdict	Meaning
Port as-is	copy the code with its tests and add an origin header
Port and adapt	copy, then change it to fit the norms tier, `basis` tags, factor traces, stable identifiers and the no-composite rule
Consume	leave it in TransportModelling and read its results through an exchange folder
Seed or method only	take parameter values into a draft norm version, or take the method into the design and validation
Do not copy	incompatible with a TIPRAS principle, or superseded

The adaptations matter more than the copying. TransportModelling hard-codes some defaults in its source, for example a fallback speed of 25 km/h, a default of two lanes and a maximum snapping distance of 5 km. TIPRAS forbids literal coefficients in `src/`. Every such value had to become a row in the norm tables, with a source and a quality tier, and the code had to raise an error if the row is missing.

23.2 What was ported, component by component

Design chapter 22 listed 23 components. Table 23.2 gives each one's verdict and what the TIPRAS source tree holds today. "Built" means the method is in `src/` and runs on Uganda; "partly" means some of the named functions are there; "not built" means nothing equivalent was found in the code.

Table 23.2: TransportModelling components: verdict and state in TIPRAS

#	TransportModelling component	Verdict	State in TIPRAS
1	`NetworkGraph.jl` (graph, shortest routes, components, criticality)	port and adapt	**built** as `src/GIS/Network.jl`: CSR adjacency forward and reverse, stable `link_id` from the snapshot. The degree-2 simplification was not built
2	`Topology.jl` (noding checks, crossing triage)	method only	**not built**: there is no `geo validate` command
3	`SpeedProfile.jl` (speeds by class × surface)	port and adapt	**built**: speeds are read from `speed_profile` norms in `src/Engine.jl`; the literal 25 km/h fallback is gone
4	`Accessibility.jl` (snapping, multi-source Dijkstra, time bands)	port and adapt	**built** as `src/GIS/TravelTime.jl` and `Snap.jl`: snap to links, Dijkstra on the reversed graph
5	`Crossings.jl` (cut edges, criticality components, crossing lifecycle)	port and adapt the cut edges and components; do not copy the ranking	**partly**: Tarjan cut edges and separate criticality components in `src/GIS/Criticality.jl`; the crossing lifecycle comparison is not built
6	`Deterioration.jl` (HDM-4 IRI projection, Markov step)	port and adapt	**partly**: the IRI projection is in `TI04_Supply.jl` with norm parameters; unknown condition is not projected; the Markov form is not built
7	`Maintenance.jl` (intervention catalogue, lifecycle optimisation, backlog)	port and adapt	**not built** as such: TI-06 uses norm cycles, and the backlog comes from agency poor km by network
8	`VehicleOperatingCost.jl` (linear VOC in IRI)	port and adapt	**built** in `TI07_Investment.jl`, coefficients from norms
9	`Appraisal.jl` (NPV, IRR checks, residual value, sensitivity)	port as-is	**re-implemented**: NPV, EIRR and BCR in `TI07_Investment.jl`; residual value, switching values and sensitivity are not there
10	`UnitCost.jl` (one-door UCSN reader with refusals)	port and adapt	**built** as `src/Data/Prices.jl`, with an origin header dated 2026-09-23
11	`Exchange.jl` (exchange header and payload structures)	port as-is	**partly**: TI-10 writes header columns on its outbound files; there is no exchange module
12	`Fiscal.jl` (MTEF reader with protected and contestable ceilings)	port and adapt	**not ported**: TI-08 reads the MTEF ceiling file directly; protected and contestable are not separated
13	`Portfolio.jl` (multi-period MILP, ε-constraint frontier)	port and adapt; replace the weighted objective	**re-implemented** as `src/Optimise/Portfolio.jl`: one objective, soft floors, frontier; dependencies, exclusions, concurrency and regional shares not carried over
14	`Poverty.jl` (concentration index)	port as-is	**built** in `src/GIS/SpatialStats.jl`, with sign tests
15	`Sectors.jl` (source-tagged logistics flows)	port and adapt	**not built**: the sector flows on disk are synthetic and are not read
16	`Units.jl`, `Provenance.jl` (`MoneyUnit`, data classes)	port and adapt	**built**: `MoneyUnit` in `src/Core/Units.jl`; the four TIPRAS bases in `src/Core/Types.jl`

#	TransportModelling component	Verdict	State in TIPRAS
17	`Assignment.jl` (link capacity, BPR, all-or-nothing, equilibrium)	port and adapt capacity and all-or-nothing; consume equilibrium	**built** for all-or-nothing loading (`TI03_Demand.jl`) and capacity from lane norms (`TI04_Supply.jl`); BPR parameters are seeded but not used
18	`Demand.jl`, `Freight.jl`, `Balancing.jl` (gravity, Furness)	consume results; port Furness	**built**: Furness gravity in `TI03_Demand.jl`, with a unit test
19	`Alignment.jl` (raster least-cost paths)	consume	**not available**: TransportModelling writes no alignment output
20	`ClimateRisk.jl` (expected annual damage)	method only	**not used**: G7 reports flood-exposed critical links; damage is not estimated
21	rail, ports, airports, waterways, urban	consume	**not available**: no exchange output; nodes only
22	data ingest (AID, UNPS, reference, ATDGR)	seed or method only	**used as data**: the AID and OSM extracts were read by the geography snapshot builder
23	`Inventory.jl` (typed validation finding codes)	port and adapt the catalogue	**not built**

Counted across the 23 rows, 9 components are built (wholly or in their adapted form), 2 are re-implemented in TIPRAS's own form, 3 are partly built, 1 (the data ingest) was used only for its data, and 8 are not built or not available. The pattern is clear. The components that the planning chain needs on every run (graph, speeds, travel time, cut edges, prices, appraisal, optimiser, gravity and loading) are all present. The components that would refine a result once better data exist (lifecycle optimisation, Markov deterioration, crossing lifecycle cost, damage estimation) are not. The integration components (exchange module, MTEF reader, sector intake) were bypassed by simpler direct reads.

23.2.1 The adapted components in more detail

Graph and travel time. The most important change from TransportModelling is the direction of the shortest-path search. TransportModelling computed travel time from each facility to every node. On one-way links that gives the time *from* the facility, not *to* it, which is the wrong direction for access. TIPRAS runs its search on the reversed graph from a single "super-source" joined to all destinations of a class, so each vertex's label is the time from that vertex to its nearest destination. Origins and destinations are snapped to the nearest point on a link rather than to the nearest vertex. TransportModelling's literal 5 km maximum snap distance was removed from the code; a `walk_leg` norm row holds a maximum of 10 km, taken from TransportModelling's own configuration.

Prices. `src/Data/Prices.jl` is the only file whose header names its TransportModelling origin. It keeps TransportModelling's refusals and tightens them. A price is resolved through the `unit_cost_map` norm, never by guessing. Illustrative (E7) and evidence-candidate (E5) prices are accepted only in exploratory runs. A duplicate code in a price vintage raises an error, where TransportModelling took the first match. A missing price raises `PriceNotFound` and the project is listed as unpriced, never costed at zero. In run 8 that rule left 23 spot improvements unpriced, and they were not funded.

Deterioration. The IRI projection keeps the HDM-4 form: roughness grows each year by a structural term proportional to the equivalent standard axles carried, and an environmental term proportional to current roughness, up to a maximum. The coefficients are norms. TransportModelling's literal maximum IRI of 20 and base year of 2024 were removed. Unknown condition stays unknown: the function skips a link with no IRI rather than inventing a starting value. Because no Uganda link has a surveyed condition, the projection is idle on every Uganda run. It is coded and waiting for data.

Optimiser. TransportModelling's portfolio model maximises a weighted sum of several objectives. That is a composite score, which TIPRAS forbids. The TIPRAS optimiser keeps the structure (a project starts at most once in the window, its capital is phased over its build years, it must finish by the horizon, money comes from pools it is

eligible for) and replaces the objective. Each specification names one measure to maximise, in its own unit, and turns every other measure into a floor with a reported shortfall. TransportModelling's dependency, exclusion, concurrency and regional-share constraints were not carried over.

23.2.2 What was not ported, and why

Some components were left out on purpose, as design chapter 22 required. Table 23.3 lists the items marked "do not copy" and confirms their absence from `src/`.

Table 23.3: Items marked "do not copy" and their state

Item	Why it is excluded	State in `src/`
Portfolio weighted objective and its weights	a composite score	absent; one declared objective per specification
`rank_crossings` and its weights (0.35, 0.30, 0.20, 0.15, with a ×3.0 cut-edge multiplier)	a composite score	absent; cut-edge components reported separately
Feeder-road priority weights	a composite score	absent
The manuals' priority index and multi-criteria weights	composite scores	absent
`PolicyIntelligence.jl` and the advisor modules	they mark options "recommended"	absent; TI-09 states that no scenario is marked recommended
`unit_cost_per_lane_km` in `asset_management.yml`	no currency and about four orders of magnitude below UCSN (2,500 against 18,000,000 UGX per km per year for paved routine maintenance)	not seeded
Literal defaults (two lanes, unknown→fair, 25 km/h, IRI 20, base year 2024, 5 km snap)	silent defaults	removed; each is a norm lookup or an explicit "unknown"
The rejected speed calibration	multipliers pinned at their bounds (250 km/h on motorways)	not used
`Topology.jl` node identifiers	sequential, so not stable between runs	not used; identifiers come from SDI

A search of `src/` finds no weighted objective, no ranking function and no use of the word "recommended" except in TI-09's statement that no scenario is marked so. This is the most important result of the reuse exercise. The engine took TransportModelling's methods without taking its composite scores.

Other components were left out for practical reasons, not principle. The exchange module, the MTEF reader and the sector intake were replaced by direct file reads, because TransportModelling and most sector models publish no exchange files for Uganda. The outbound header that TI-10 writes has six columns (country, run, scenario, schema version, producer and the exploratory flag). Design chapter 22 asked for three more: the geography version, norm version and price vintage. They are in each run's manifest but not on the outbound files.

23.3 Seed norms from TransportModelling configuration

Design chapter 22 §22.3 directed that TransportModelling's configuration files be loaded into a **draft** norm version only, each row with quality *assumed* and a source naming the file and date. `scripts/seed_norms.jl` does this, reading `network.yml`, `asset_management.yml`, `appraisal.yml` and `traffic.yml` as they stood on 2026-09-23. In the current draft norm set (`UGA-2026.09-draft-f`, 646 rows), 172 rows (27 %) cite TransportModelling as their source. All 172 are of quality *assumed*. Table 23.4 shows where they sit.

Table 23.4: Norm rows seeded from TransportModelling configuration (draft-f)

Norm table	Rows from TM	Source file	What they set
speed_profile	96 of 98	network.yml	88 motorised speeds by class and surface; 8 walking speeds by surface
deterioration	29	asset_management.yml	structural and environmental rates and maximum IRI by surface; the drainage multiplier; IRI by condition class
vehicle	29	asset_management.yml	for 7 vehicle types: passenger-car units, equivalent standard axles, base operating cost and its sensitivity to roughness; the reference IRI
vehicle_mix	7	asset_management.yml	the default traffic composition
maintenance_norm	6	asset_management.yml	periodic maintenance cycles; IRI after works
road_capacity	2	traffic.yml	BPR α and β
walk_leg	2	network.yml	access-leg walking speed and maximum snap distance
gravity_beta	1	network.yml	the freight gravity decay

The two speed_profile rows not from TransportModelling are the assumed ferry speed of 8 km/h, including waiting, for motorised traffic and foot passengers.

23.3.1 The values and their status

Every one of these values is uncalibrated for Uganda. The speed table is the most influential, because it drives every travel time and therefore every access result. Its motorised speeds run from 80 km/h (trunk road, paved) down to 12 km/h (track, earth). A community road is assumed to carry 18 km/h on earth, 20 km/h on engineered earth and 22 km/h on gravel. A district road is assumed to carry 22, 25 and 30 km/h on the same surfaces. The file itself calls these generic sub-Saharan defaults and "NOT Ugandan". TransportModelling measured that they apply to 96.1 % of Uganda's network km. In TIPRAS the only observed adjustment is OSM maxspeed, which is used as an upper cap. Surface "unknown" has its own speed in the table (for example 60 km/h on trunk roads and 22 km/h on community roads), while for the access standard an unknown surface is ranked as earth.

The deterioration values are HDM-4 in form and illustrative in value: a structural rate of 0.28 per million standard axles and an environmental rate of 0.035 a year on paved roads, against 0.90 and 0.070 on unpaved roads, with maximum IRI of 16 and 20 m/km. IRI by condition class is 3.5 (good), 5.0 (fair), 7.0 (poor) and 11.0 m/km (bad). These values are idle in Uganda runs today, because no link has a condition to project from. They become active the day a condition survey arrives, so they will need review then.

The maintenance cycles are 8 years for resealing and 5 years for regravelling, with roughness reset to 2.0 m/km after rehabilitation, 2.8 after periodic maintenance and 1.8 after upgrading. The cycles drive the periodic maintenance requirement; the reset values drive the appraisal benefits.

23.3.2 Vehicle operating costs: the most fragile values

The appraisal in TI-07 values a project's benefits as savings in vehicle operating cost (VOC) and travel time. The VOC part uses TransportModelling's linear form, in which the cost per vehicle-km rises in a straight line with road roughness above a reference level:

$$\mathrm{VOC}(\mathrm{IRI}) = \mathrm{VOC}_0 \cdot (1 + s \cdot \max(\mathrm{IRI} - \mathrm{IRI}_{ref}, 0)) \tag{23.1}$$

where

- IRI is the road's roughness in m/km;
- IRI_ref is the reference roughness, 2.0 m/km;
- VOC_0 is the mix-weighted base operating cost per vehicle-km on a smooth road, in UGX;

- s is the mix-weighted increase in cost per unit of roughness above the reference.

The two mix-weighted coefficients come from the vehicle table and the default mix:

$$\mathrm{VOC}_0 = f_x \sum_h m_h c_h, \qquad s = \sum_h m_h s_h \tag{23.2}$$

where

- h is a vehicle type (car, minibus, bus, motorcycle, light, medium and heavy truck);
- m_h is its share of traffic in the default mix (motorcycles 0.30, cars 0.24, minibuses 0.18, light trucks 0.13, medium trucks 0.08, heavy trucks 0.04, buses 0.03);
- c_h is its base operating cost per km (from 0.10 for a motorcycle to 1.45 for a heavy truck);
- s_h is its sensitivity to roughness (from 0.030 to 0.075);
- f_x is the exchange rate norm, 3,700 UGX per USD.

Computed from the draft-f norms, the mix-weighted base cost is 0.424 per vehicle-km, or about 1,570 UGX at the assumed rate, and the mix-weighted sensitivity is 0.047. A road at the "poor" roughness of 7.0 m/km therefore costs about 24 % more per vehicle-km to drive than one at the reference, and a "bad" road at 11.0 m/km about 42 % more.

Three things make these values fragile, and the run log says so in its fallback list: "VOC coefficients are TransportModelling UNCALIBRATED values in assumed USD converted at an assumed FX; appraisal is exploratory".

1. **The currency is assumed.** TransportModelling's file gives the base costs with no currency. TIPRAS records them as USD per vehicle-km and labels the unit "currency assumed; TM file states none".
2. **The exchange rate is assumed.** The appraisal converts at 3,700 UGX per USD, a norm labelled "illustrative 2026 exchange rate; replace with FINEX path". The evidence prices, by contrast, are converted with World Bank WDI rates, whose 2026 value is extrapolated to 3,580 UGX per USD. The two parts of the same appraisal, costs and benefits, therefore use different exchange rates, about 3 % apart.
3. **The values are not Ugandan.** They are HDM-4-form illustrations. Design chapter 22 directs that they be replaced by a RED or HDM-4 form once a Uganda VOC study exists.

The other appraisal parameters are also assumed: a 12 % discount rate pending the MoFPED public investment guidance (TransportModelling's own file leaves it empty by design), a 20-year period, a value of time of 3,000 UGX per person-hour, and an average occupancy of 2.2. Time savings on links that are closed in the wet season are valued at three times the improved time, a placeholder flagged in the run log until a closure-cost norm exists.

These assumptions explain a result reported in the investment chapter. In run 8, 574 of the 41,757 appraised projects (1.4 %) have a positive NPV at 12 %. That share depends on the VOC coefficients, the value of time and the discount rate, all three assumed. It should be read as a test of the method, not as a finding about the economics of Uganda's roads. The access results, which do not use the VOC values, do not carry this weakness.

23.4 Tests and benchmarks

Design chapter 22 asked that each component's TransportModelling tests be ported with it, so that every adaptation would be a deliberate and visible change. That was not done. TIPRAS has its own suite of 2,265 assertions in nine testsets (see the build status chapter). It tests the ported methods against the manual's worked examples, not against TransportModelling's tests: Furness on a 2×2 case, Dijkstra against brute force, cut edges on a small graph with a pendant link, and the concentration index on cases of known sign.

The golden country `TST` was built afresh rather than grown from TransportModelling's `ZZZ` triangle. It meets the purpose that design chapter 22 §22.4 set out: it has districts, a cut edge and a seasonal drift. The Braess case from

TransportModelling's assignment tests was not carried over, and TransportModelling's health access benchmark (88.1 % of population within 30 minutes, motorised) has not been added as a regression check.

23.5 Requests back to TransportModelling

Design chapter 22 §22.10 listed five requests back to TransportModelling. None of them blocked the build, because TIPRAS runs on norm-tier fallbacks, and none has been met. Table 23.5 gives their state.

Table 23.5: Requests back to TransportModelling and their state

Request	State	Effect on TIPRAS today
Write exchange files for origin–destination matrices, link volumes, deterioration parameters and alignment candidates	not met	TI-03 uses its own gravity fallback; G1 cannot separate "network missing" from "no destination"
Fix the one-way direction in accessibility	not needed by TIPRAS (fixed in its own port)	none
Delete the duplicate `road_pbf.csv`	not known	none
Persist the noded network with stable identifiers, or adopt SDI's	not met	TIPRAS takes geometry and identifiers from SDI
Receive TIPRAS's funded programme for corridor appraisal	not met	no outbound channel to TransportModelling

There is a new and more basic issue. At the time of writing (October 2026), the TransportModelling folder is not present at its estate path. The current norm set and geography snapshot were built while it was present, and they are stored with TIPRAS, so runs are not affected. But `scripts/seed_norms.jl` reads the four configuration files from that path, and `scripts/build_geo_snapshot.jl` reads the OSM tag table, the AID points of interest and the elevation model from it. Neither script can be re-run until the folder is restored or its contents are copied into TIPRAS's own data folders.

23.6 Data quality and limitations

- **Every seeded value is* assumed.* All 172 TransportModelling rows carry that tier, and none may enter an approved norm version without national review.
- **Speeds dominate the access results.** They are generic defaults on most of the network. The rural access index does not depend on them (it is a distance measure), but every time-based access class does.
- **The appraisal is the weakest link.** VOC base costs in an assumed currency, at an assumed exchange rate, with a mix that is not Ugandan, make every NPV, EIRR and BCR exploratory. The NPV-maximising optimiser specification inherits this.
- **Ported code has no ported tests.** The methods are tested on worked examples and on `TST`, not against TransportModelling's own cases.
- **Some idle components wait for data.** Deterioration and the rehabilitation trigger are coded but do nothing without link condition.
- **The source is currently missing.** The TransportModelling folder is absent, so its configuration and data extracts cannot be re-read.

23.7 Implications for planning

- The access results in this report rest on TransportModelling's speeds but not on its cost or appraisal values. They can be discussed with planners as exploratory, with the speed caveat stated.

- The economic results (NPV, EIRR, BCR and the NPV-maximising optimiser run) should not be quoted outside the engine's own development until a Uganda VOC study, a national value of time and the MoFPED discount rate replace the assumed values.
- A single exchange-rate source should be used across costs and benefits. Today the evidence prices use 3,580 UGX per USD and the appraisal uses 3,700.
- The TransportModelling folder should be restored to its estate path, or its Uganda inputs copied into TIPRAS, so that the draft norms and the geography can be rebuilt and audited.
- The deterioration and lifecycle parameters will become active as soon as condition data arrive. They should be reviewed at the same time, with the road agency, rather than left at their illustrative values.
- Porting TransportModelling's tests for the adapted components, and adding its health access benchmark as a regression check, would make each adaptation visible and testable, as design chapter 22 intended.
- The exclusion of every composite score and every "recommended" label held through the port. Any future reuse of TransportModelling code should be checked against the same list.

CHAPTER 24

24 Mathematical models

This chapter is the reference for every mathematical model the engine uses, in one notation. The module chapters described *when* and *why* each model runs; this chapter states *exactly what* it computes, so that a planner, a reviewer or an auditor can check any number in the report against its definition. It implements design chapter 23 (`docs/23_MATHEMATICAL_MODELS.md`). For each model it gives the equations, explains each variable, and says how the Uganda implementation differs from the design where it does. Two rules run through the whole chapter. First, every parameter written as a Greek letter or a named coefficient is a *norm*: a row in the norm set `UGA-2026.09-draft-f` with a source, a year and a quality tier. No parameter has a default value in code. Second, the seven gap dimensions are defined separately and are never combined into a single index, score or ranking.

24.1 What the chapter covers

The models follow the chain of the ten modules, from population to budget, plus the spatial models that serve all of them. Table 24.1 lists the main symbols.

Table 24.1: Main sets and quantities used in this chapter

Symbol	Meaning	Unit
i	settlement (a populated 1 km cell)	—
d	district (135 in Uganda)	—
t, t_0, H	plan year, base year (2026), horizon (2030)	year
k	destination class (all-season road, primary health, EmONC, primary school, secondary school, market, administrative centre)	—
n	sector node (facility, school, market, mine, park, gateway, production zone)	—
l = (u, v)	network link from vertex u to vertex v	—
h	vehicle class	—
s	season (wet, dry)	—
w	works type (RM, PM, SI, RH, UP and others)	—
j	candidate project	—
π	financing pool	—
P_{i,t}	population of settlement i in year t	persons
len_l	length of link l, in the projected coordinate system	km
tt_{i,k,t}	travel time from i to the nearest destination of class k	minutes
τ_k	access threshold for class k	minutes or km
x_l	load on link l	vehicles or PCU per day
r	discount rate	per year

The indicator $1[\cdot]$ equals 1 when the condition inside is true and 0 otherwise. The positive part $[a]^+$ equals max(0, a).

24.2 Method: population and the planning grid (TI-01)

24.2.1 Raking to official totals

The gridded population (WorldPop 2020, aggregated to 1 km cells) does not add up to the official district totals. TI-01 scales it so that it does, separately for each district, urban or rural class and year:

$$\rho_{d,\,u,\,t} = \frac{P^{NSO}_{d,\,u,\,t}}{\sum_{i \,\in\, I_{d,\,u}} P^{grid}_{i,\,t_0}} \tag{24.1}$$

$$P_{i,\,t} = \rho_{d(i),\,u(i),\,t} \cdot P^{grid}_{i,\,t_0} \tag{24.2}$$

where P^{NSO} is the official projection for district d and class u (here UBOS 2022 totals carried forward at an assumed 3.0 % a year), P^{grid} the gridded population of cell i in the base year, and I_{d,u} the cells of district d in class u. By construction the raked cells add up exactly to the official total. A factor outside the range 0.85 to 1.15 is a warning that the grid and the official figures disagree. In the reference run, 1,235 district-class-years were outside that range.

24.2.2 Group populations

Most access standards apply to a population group rather than to everyone: women aged 15–49 for EmONC, children aged 6–12 for primary school. The group population is:

$$P^{g}_{i,\,t} = P_{i,\,t} \cdot a_{g,\,d(i),\,u(i),\,t} \tag{24.3}$$

where a_g is the share of group g in the population. In the draft norms these shares are national constants (for example 0.235 for women 15–49 and 0.21 for children 6–12), quality *assumed*, pending district age structure from UBOS or HouseholdModelling.

24.2.3 Agricultural production zones

Cells whose production density exceeds a threshold θ_ag (20 t per km² in the draft norms) are joined by contiguity into production zones z. The marketed surplus of crop c in zone z is:

$$Y^{surplus}_{c,\,z,\,t} = \mu_c \sum_{i \,\in\, z} Y_{c,\,i,\,t} \tag{24.4}$$

where Y_{c,i,t} is the production of crop c in cell i and μ_c the marketed share (for example 0.45 for maize and 1.0 for sugarcane and coffee, all *assumed*). A zone with less than Y_min = 5,000 t of surplus is not a freight origin. The reference run found 10 zones holding 9.76 Mt.

24.3 Method: network, link time and seasons (TI-04)

24.3.1 Seasonal availability

The network is a directed graph: a two-way road is two arcs. A link is available in season s only if its surface, its condition and every structure on it allow passage:

$$a_{l,\,s,\,t} = \mathbf{1}[surf_{ok}(l,\,s)] \cdot \mathbf{1}[cond_{ok}(l,\,t)] \cdot \prod_{b \,\in\, B(l)} \mathbf{1}[passable(b,\,s,\,t)] \tag{24.5}$$

where B(l) is the set of structures (bridges, drifts, fords) on link l. The three tests come from the season-rule table. For example, unengineered earth in a flood-exposed zone fails in the wet season, and a ford on a seasonal river is impassable in the wet season. In the reference run, 5,013 links (3,299 km) close in the wet season.

24.3.2 Link travel time

The free-flow time of vehicle class h on link l is:

$$t^{0}_{l,\,h} = \frac{60\,len_l}{v_{l,\,h}}, \qquad v_{l,\,h} = V(mode_l, surf_l, cond_l, terrain_l, s, h) \tag{24.6}$$

where v is the speed read from the speed profile and the speed factors for condition, terrain and wet season. A missing speed raises an error; there is no fallback speed. Where traffic approaches capacity, the time rises by the Bureau of Public Roads volume–delay function:

$$t_{l,\,h}(x_l) = t^{0}_{l,\,h}\left[1 + \alpha\left(\frac{x_l}{Q_l}\right)^{\beta}\right] \tag{24.7}$$

where x_l and Q_l are the design-hour flow and capacity in PCU per hour, and α = 0.15 and β = 4 in the draft norms (*assumed*).

24.3.3 The walking leg

A settlement is joined to the network at its nearest link. The off-network leg takes:

$$w_i = \frac{60\,\delta_i\,\varphi_i}{v_{walk}} \tag{24.8}$$

where δ_i is the straight-line distance from the settlement to its snap point in km, ϕ_i a terrain factor (1.0 flat, 1.15 rolling, 1.4 mountainous) and v_walk the walking speed (5 km/h). All three are *assumed* norms. In the reference run the median snap distance was 0.24 km, and 1,039 settlements lay more than 5 km from any mapped road.

24.4 Method: accessibility (TI-02, TI-04, TI-09)

24.4.1 Shortest paths to the nearest destination

The cost of a path R is the sum of its link times and node penalties, and the shortest-path distance d(u, v) is the minimum over all paths. Rather than search from every settlement to every destination, the engine joins all destinations of class k to one artificial super-source and runs one Dijkstra search on the reversed graph. The label at each vertex u is then:

$$D_k(u) = \min_{j \in J_k\,(t)} \left[d(u, snap(j)) + w_j\right] \tag{24.9}$$

where J_k(t) is the set of functional destinations of class k in year t, snap(j) the point where destination j joins the network and w_j its own egress walk. The travel time of settlement i to class k is:

$$tt_{i,\,k,\,t} = w_i + D_k(snap(i)) \tag{24.10}$$

An unreachable settlement has an infinite travel time, stored as missing and never as zero. Searching the reversed graph gives the right answer on one-way links.

24.4.2 Meeting the standard and the Rural Access Index

A settlement meets a time standard if:

$$A_{i,\,k,\,t} = \mathbf{1}\left[\, tt_{i,\,k,\,t} \le \tau_{k,\,u(i)} \,\right] \tag{24.11}$$

where τ_{k,u} is the threshold for class k and settlement type u (urban or rural), for example 60 minutes for primary health care in rural areas and 120 minutes for EmONC. The all-season road class uses a distance standard instead: 2 km in a straight line from the settlement to the nearest link that is all-season and open in the wet season. The Rural Access Index of district d (SDG indicator 9.1.1) is the rural population share that meets it:

$$RAI_{d,\,t} = \frac{\sum_{i \in I_d \cap I^{R}} P_{i,\,t}\, A_{i,\,ARS,\,t}}{\sum_{i \in I_d \cap I^{R}} P_{i,\,t}} \tag{24.12}$$

where I^R is the set of rural settlements and ARS the all-season road class. The national RAI is the same ratio over all rural settlements. In the reference run it is 28.7 % when unknown surface is treated as earth and 93.5 % when it is treated as gravel.

24.4.3 Potential accessibility

For agriculture and industry the engine also computes a descriptive index of market potential:

$$H_{i,\,t} = \sum_{j \in J_k(t)} M_j \exp\!\left(-\beta_H\, tt_{i,\,j,\,t}\right) \tag{24.13}$$

where M_j is the size of market j and β_H a decay parameter. It is emitted to the sector engines for information. It is never used as a gap or a weight.

24.5 Method: movement generation and distribution (TI-02, TI-03)

24.5.1 Freight and passenger generation

Freight leaving production node n and arriving at consumption centre m, per commodity c, is:

$$O^{c}_{n,\,t} = \mu_c\, Y_{c,\,n,\,t}, \qquad D^{c}_{m,\,t} = q_c\, P_{m,\,t}\,(1 - self_{c,\,m}) \tag{24.14}$$

where q_c is consumption per head and self_{c,m} the self-supplied share. In the reference run each district's consumption centre is its most populous settlement, and consumption is 0.35 t per person a year (*assumed*). Total consumption demand is scaled to total supply before distribution, and the scaling factor is logged. Passenger trips by purpose p are:

$$O^{p}_{i,\,t} = \rho_{p,\,u(i)}\, P_{i,\,t} \left(\frac{\bar{y}_t}{\bar{y}_{t_0}}\right)^{\epsilon_p} \tag{24.15}$$

where ρ_p is the trip rate, ȳ the GDP per head from FINEX or CDCGE and ε_p the income elasticity. Growth comes from the sector models; a fixed growth rate is allowed only as a flagged fallback. In the reference run no sector supplied output trajectories, so freight grows at the fallback rate of 5 % a year.

24.5.2 Generalised cost

The cost of moving one unit of commodity c (a tonne or a trip) from origin o to destination d on the cheapest path combines time and money:

$$gc^{c}_{od} = \min_{R} \sum_{l \in R} \left[\frac{VoT_c\, t_{l,\,h(c)}}{60} + \frac{VOC_{h(c)}(IRI_l, v_l)\, len_l}{load_{h(c)}} \right] \tag{24.16}$$

where VoT_c is the value of time per hour per unit, VOC the vehicle operating cost per vehicle-km, IRI_l the link's roughness and load_h the average load per vehicle. Node penalties (border time and charges) are added for each node on the path.

24.5.3 The doubly constrained gravity model

Flows between origins and destinations follow Wilson's doubly constrained gravity model:

$$T_{od} = A_o \, B_d \, O_o \, D_d \, f(gc_{od}) \tag{24.17}$$

where O_o is the total leaving origin o, D_d the total arriving at destination d, and f a deterrence function. The balancing factors make the row and column totals match:

$$A_o = \frac{1}{\sum_d B_d \, D_d \, f(gc_{od})}, \qquad B_d = \frac{1}{\sum_o A_o \, O_o \, f(gc_{od})} \tag{24.18}$$

The deterrence function is exponential:

$$f(gc) = \exp(-\beta \, gc) \tag{24.19}$$

with β = 0.2 per hour for freight in the draft norms (*assumed*, not calibrated). The model is solved by Furness iteration: start from B = 1, update A and B alternately, and stop when every row and column total is within 0.1 % of its target, or after 500 iterations. The reference run logged 500 iterations for the agricultural distribution, with 9.69 Mt distributed.

Where observed trip lengths exist, β is calibrated by Hyman's method, which adjusts β until the modelled mean cost equals the observed mean c̄^{obs}:

$$\beta_{n+1} = \beta_n \, \frac{\bar{c}(\beta_n)}{\bar{c}^{obs}} \tag{24.20}$$

No observed trip-length data exist for Uganda in the estate, so this step has not been run.

24.5.4 The two-stage agricultural chain

The surplus of each production zone is first split across the markets within the market standard by a logit model:

$$share_{z,\,m} = \frac{\exp(-\theta \, gc_{zm})}{\sum_{m' \in M_z} \exp(-\theta \, gc_{zm'})} \tag{24.21}$$

where M_z is the set of markets within reach of zone z and θ = 0.5 per hour (*assumed*). If no market is within reach, the whole surplus goes to the nearest market and the zone registers a G1 market gap. In the second stage, each market's throughput is distributed to consumption centres and gateways by the gravity model above. This is what makes feeder roads carry visible freight.

24.6 Method: traffic loading, AADT and axle loads (TI-03)

24.6.1 Loading flows onto links

With δ^{od}_l equal to 1 if link l is on the chosen path from o to d, the load of commodity c on link l is:

$$x_l^c = \sum_{od} T_{od}^c \, \delta_l^{od} \tag{24.22}$$

The reference run uses all-or-nothing assignment on the shortest generalised-cost path. Incremental assignment, which loads the flow in four slices (40 %, 30 %, 20 %, 10 %) and updates link times by the BPR function after each,

is specified but not yet implemented. The user-equilibrium assignment belongs to TransportModelling. In the reference run 288,725 of 817,382 links carry modelled flow, and freight amounts to 2.722 bn tonne-km.

24.6.2 Vehicles and AADT

Tonnes and trips become vehicles by class h:

$$veh_{h,\,l} = \sum_{c} \frac{x_l^{\,c}\,\eta_{c,\,h}}{payload_h\, lf_h}(1 + e_h) + \sum_{p} \frac{x_l^{\,p}\,\eta_{p,\,h}}{occ_h} \tag{24.23}$$

where η is the share of a commodity or purpose carried by class h, payload_h and lf_h the payload and load factor, e_h the empty-return factor and occ_h the occupancy. Annual average daily traffic and passenger car units follow:

$$AADT_l = \tfrac{1}{365}\sum_h veh_{h,\,l}, \qquad PCU_l = \tfrac{1}{365}\sum_h pcu_h\, veh_{h,\,l} \tag{24.24}$$

24.6.3 Equivalent standard axles

Pavement damage grows with the fourth power of axle load. The equivalence factor of a vehicle with axle groups a carrying loads P_a is:

$$EF_h = \sum_a \left(\frac{P_a}{P_{std}(a)}\right)^4 \tag{24.25}$$

where P_std = 8.16 t for a single axle. Overloading raises the effective factor:

$$EF_h^{\,eff} = EF_h\left[1 + \omega_h(\Omega_h - 1)\right] \tag{24.26}$$

where ω_h is the share of vehicles overloaded and Ω_h the damage multiplier when overloaded. Annual equivalent standard axles (ESA) on link l are:

$$ESA_{l,\,t} = 365 \sum_h veh_{h,\,l,\,t}\, EF_h^{\,eff} \tag{24.27}$$

As implemented, the engine reads one ESA per vehicle for each class from the vehicle norm table, in place of computing it from axle loads; those values come from TransportModelling and are *assumed*. Design traffic over a design life of n years at growth rate g is:

$$N_{des,\,l} = ESA_{l,\,t_0}\,\frac{(1+g)^n - 1}{g}\, f_{dir}\, f_{lane} \tag{24.28}$$

where f_dir is the directional split and f_lane the lane distribution factor. In the reference run n = 15 years, g = 5 % (the fallback growth rate) and f_dir = 0.5.

24.7 Method: the required standard (TI-03)

A standard is a record of surface, width, lanes, all-season status, axle load, speed and structural capacity. Standards are ordered field by field: S ≼ S' when every field of S is at most the matching field of S'. Four tests each set a minimum standard for a link: the traffic it carries, its functional role, the sector nodes whose route it lies on, and the access paths it serves. The required standard is the join, the field-wise maximum, of the four:

$$S^{*}_{l,\,t} = S^{\,traffic}_{l,\,t} \sqcup S^{\,function}_{l} \sqcup S^{\,node}_{l,\,t} \sqcup S^{\,access}_{l,\,t} \tag{24.29}$$

The *binding test* is the one that sets the surface field. Ties are recorded. If the design traffic exceeds the structural capacity of the required surface, the surface is promoted to the next rank that can carry it. In the reference run, 656,479 links bind on function, 160,189 on traffic, 686 on a sector-node standard and 28 on mode. The standard gap of link l in field ϕ is:

$$\Delta_{\varphi}(l,t) = rank(S_{\varphi}^{*}) - rank(S_{\varphi}) \tag{24.30}$$

A positive value is a shortfall. Unknown surface is treated as earth for this comparison, which is why 242,068 km are below their required surface and 169,343 km of gap links with unknown surface are listed for survey rather than for works.

24.8 Method: pavement deterioration (TI-04, TI-06)

TransportModelling estimates the deterioration parameters; TIPRAS applies them. For paved roads, roughness (IRI, in m/km) grows each year by a structural term driven by axle loads and an environmental term:

$$IRI_{l,\,t+1} = \min\left(IRI_{l,\,t} + \left(k_s\, MESA_{l,\,t} + k_e\, IRI_{l,\,t}\right)\psi_{drain}\,\psi_{rain},\ \ IRI_{max}\right) \tag{24.31}$$

where MESA is millions of ESA a year, k_s and k_e the structural and environmental coefficients (0.28 and 0.035 for asphalt and surface dressing in the draft norms), ψ multipliers for drainage and rainfall, and IRI_max a ceiling (16 m/km). For gravel roads, the gravel layer wears away:

$$GT_{l,\,t+1} = \left[GT_{l,\,t} - \left(g_0 + g_1 AADT_{l,\,t}\right)\psi_{rain}\,\psi_{grad}\right]^{+} \tag{24.32}$$

where GT is the gravel thickness in mm and g_0 and g_1 the loss coefficients. When GT falls below a minimum, the link becomes earth and re-enters the season rule as earth. For networks that report condition only by class (good, fair, poor, bad), a Markov alternative moves the distribution of a group of links between classes:

$$\pi_{g,\,t+1} = \pi_{g,\,t}\,\mathbf{M}(surf, band) \tag{24.33}$$

where π_g is the row vector of class shares and **M** a row-stochastic transition matrix by surface and traffic band.

The Uganda run cannot apply these models link by link. No link in the snapshot has a recorded roughness or condition, so the do-nothing projection leaves unknown condition unknown and invents no IRI. Network condition enters only as agency statistics of poor km by network class (139,691 km in poor condition), and the condition programmes are priced in tranches from those totals. The coefficients are TransportModelling values flagged *UNCALIBRATED, HDM-4 form, not Ugandan*.

24.9 Method: capacity and queues (TI-04, TI-05)

Road capacity in PCU per hour comes from the norm table by lanes, adjusted for width, heavy vehicles and terrain:

$$Q_l = Q_{base}(lanes_l)\; f_w(width_l)\; f_{hv}\; f_{terrain}(terrain_l) \tag{24.34}$$

The draft norms give 900 PCU/h for one lane, 2,000 for two and 5,600 for four, all *assumed*. Converting to a daily capacity with the design-hour factor K = 0.1, the volume-to-capacity ratio is:

$$vc_{l,\,t} = \frac{PCU_{l,\,t}\,K}{Q_l} \tag{24.35}$$

A link is over capacity when vc exceeds the target of 0.8. Lanes are unknown on 793,861 of the 817,382 links; on the 276,431 loaded links among them, capacity is reported as unknown, never assumed adequate. For border posts, ferries and port gates whose waiting time is not observed, the design specifies a queue model. With one server it is Kingman's formula:

$$W_q \approx \frac{\rho}{1-\rho} \cdot \frac{C_a^2 + C_s^2}{2} \cdot \frac{1}{\mu} \tag{24.36}$$

where ρ = λ/μ is utilisation (arrival rate over service rate), and C_a and C_s the coefficients of variation of arrival and service times. At ρ ≥ 1 the queue is unstable and the node is reported as a capacity gap. The queue model is not yet implemented, because no gateway flows were supplied.

24.10 Method: the seven gap dimensions (TI-05)

The gaps are defined separately. Each has its own unit and none is ever added to another.

G1 access. The unserved population of class k in district d, and the depth of each settlement's shortfall:

$$U_{k,d,t} = \sum_{i \in I_d} P_{i,t}^{g(k)} \left(1 - A_{i,k,t}\right) \tag{24.37}$$

$$depth_{i,k,t} = \left[tt_{i,k,t} - \tau_{k,u(i)}\right]^{+} \tag{24.38}$$

Each unserved settlement is given one cause: *no_destination* if it stays unserved even on the required-standard network with candidate links; *network_missing* if the required path needs a link that does not exist; *network_substandard* if a link on the path is below standard or closes in the wet season; *network_deteriorated* otherwise. Without candidate alignments from SDI or TransportModelling, *network_missing* cannot be separated from *no_destination* in this run.

G2 condition. Km in worse condition than the target, by link group g (network class × district):

$$CG_{g,t} = \sum_{l \in g} len_l \, \mathbf{1}[cond_{l,t} \succ target_g] \tag{24.39}$$

In this run G2 is computed from agency statistics by network, not from links.

G3 capacity. The excess load over the target volume-to-capacity ratio:

$$CapG_{l,t} = x_{l,t} - \overline{vc}\, Q_l^{day} \tag{24.40}$$

G4 standard. The standard gap by field, as defined in Equation (24.30), and the all-season gap: a link that must be all-season but is not available in the wet season.

G5 redundancy. The criticality measures of the network section below: population isolated by the loss of one link, and cut edges.

G6 sector connection. The connection test of Chapter 22 (Equation (22.1)): a node fails if any link on its route to its target falls below its connection standard.

G7 climate exposure. A link is a resilience gap when it is exposed, critical and without a reasonable detour:

$$RG_l = \mathbf{1}[exposure_l \geq \underline{e}] \cdot \mathbf{1}[crit_l \geq \underline{c}] \cdot \mathbf{1}[detour_l > \bar{\rho}] \tag{24.41}$$

where the three thresholds are norms. A link with no hazard data is recorded as *unassessed*, never as safe.

24.11 Method: costing (TI-06)

24.11.1 Works cost

The capital cost of project j, at each of the low, central and high price points, is the sum over its links of quantity times unit cost, plus overheads:

$$X_j^{cap} = \left(\sum_{l \in L_j} q_w(l)\, \kappa_w(S_l \to S_l^{*}, terrain_l) \right) (1 + \theta_w^{DS} + \theta_w^{cont}) \tag{24.42}$$

where q_w(l) is the quantity (km, lane-km or units), κ_w the unit cost read from UCSN through the unit-cost map, θ^{DS} the design and supervision rate and θ^{cont} the contingency rate. A missing price is never zero: the project is marked unpriced and cannot be funded. In the reference run 41,867 of 41,890 projects are priced, at 92.8 tn UGX central (56.6 to 125.3 tn), with 205.2 bn UGX a year of induced routine maintenance.

24.11.2 Cash phasing

Capital is spread over the construction period n_j by an S-curve, the regularised incomplete beta function I_ξ(a, b):

$$cash_{j,k} = X_j^{cap} \left[I_{k/n_j}(a_w, b_w) - I_{(k-1)/n_j}(a_w, b_w) \right], \quad k = 1, \ldots, n_j \tag{24.43}$$

The draft norms use a = b = 2, the symmetric S, with 3 years for an upgrade and 2 for rehabilitation (*assumed*).

24.11.3 Present value

Every money stream y_t is discounted to the base year:

$$PV(y) = \sum_{t=t_0}^{t_0 + \Theta} \frac{y_t}{(1+r)^{t-t_0}} \tag{24.44}$$

with r = 12 % and an analysis period Θ = 20 years in the draft norms (*assumed*). The design adds a residual value for layers whose life outlasts the analysis period; it is not yet implemented.

24.11.4 Maintenance requirement and induced maintenance

The routine maintenance requirement of the network is the sum over maintainable links of length times the routine maintenance rate for the link's surface:

$$M_t^{RM} = \sum_{l} len_l\, rm(surf_{l,t}) \tag{24.45}$$

In the reference run this is 954.7 bn UGX a year. Every capital project changes the network it leaves behind, so it changes the maintenance requirement. The induced routine maintenance of project j after completion is:

$$IM_j = \sum_{l \in L_j} len_l \left[rm(S_l^{*}) - rm(S_l) \right] \tag{24.46}$$

It may be negative: replacing gravel by a sealed surface can remove regravelling. This is the term that makes every capital figure in the report carry its maintenance. The condition backlog is the cost of overdue periodic maintenance and rehabilitation:

$$BL_{t_0} = \sum_l len_l \left(\mathbf{1}[IRI_{l,\,t_0} \geq IRI^{trig}]\, pm(\cdot) + \mathbf{1}[cond_l \in \{poor, bad\}]\, rh(\cdot) \right) \tag{24.47}$$

In the reference run it is computed from agency poor-km totals by network: 7.25 tn UGX central.

24.11.5 Lifecycle optimisation

For a link with a known condition, the best maintenance policy over the analysis period is found by backward dynamic programming over condition states c:

$$V_t(c) = \min_{w \in W(c)} \left[\kappa_w(c) + U(c, w) + \frac{V_{t+1}(T(c, w))}{1 + r} \right] \tag{24.48}$$

where κ_w(c) is the cost of works w in state c, U the road user cost, T the state transition after works and deterioration, and V_Θ = −RV the terminal value. The policy gives the timing of periodic maintenance. It is specified for porting from TransportModelling but not used in the Uganda run, which has no link condition.

24.12 Method: road user costs and appraisal (TI-07)

24.12.1 Vehicle operating cost

The design form is a polynomial in roughness and speed:

$$VOC_h(IRI, v) = a_{0h} + a_{1h} IRI + a_{2h} IRI^2 + \frac{a_{3h}}{v} + a_{4h} v^2 \tag{24.49}$$

As implemented, the engine uses TransportModelling's linear special case, weighted by the vehicle mix:

$$VOC(IRI) = c_0\, X \left(1 + \sigma \left[IRI - IRI_{ref}\right]^+\right) \tag{24.50}$$

where c_0 is the mix-weighted base cost per vehicle-km in US dollars, X the exchange rate (3,700 UGX per dollar, *assumed*), σ the mix-weighted roughness sensitivity and IRI_ref = 2 m/km. The coefficients are TransportModelling values marked uncalibrated. Because no link has an observed roughness, the VOC term contributes nothing to the benefits of the reference run.

24.12.2 User benefits

Where a project lowers the generalised cost of travel from gc^0 to gc^1, and flow rises from T^0 to T^1, the user benefit follows the rule of a half:

$$\Delta CS = \sum_{od,\,c} \tfrac{1}{2} (gc^0_{od} - gc^1_{od})(T^0_{od} + T^1_{od}) \tag{24.51}$$

With demand held fixed, as in the reference run, this is the cost saving times the flow. As implemented, the annual benefit of project j is the time saved on its links, valued per person:

$$B_j = \sum_{l \,\in\, L_j} 365\; AADT_l\; \frac{t_l^{cur} - t_l^{after}}{60}\; occ\; VoT \tag{24.52}$$

with an occupancy of 2.2 persons per vehicle and a value of time of 3,000 UGX per person-hour (both *assumed*). A link closed in the wet season has no current time; its saving is valued at three times the improved time. That is a flagged placeholder until a closure-cost norm exists.

24.12.3 Net present value, benefit–cost ratio and EIRR

Benefits start after construction. Costs are the capital cash flow and, after completion, the induced maintenance:

$$NPV_j = \sum_{t=1}^{\Theta} \frac{B_{j,t} - C_{j,t}}{(1+r)^t} \tag{24.53}$$

$$BCR_j = \frac{PV(B_j)}{PV(C_j)} \tag{24.54}$$

The economic internal rate of return is the discount rate r* that sets NPV to zero, found by Brent's method on the bracket from −50 % to 500 %:

$$\sum_{t=1}^{\Theta} \frac{B_{j,t} - C_{j,t}}{(1+r^*)^t} = 0 \tag{24.55}$$

If the cash flow changes sign more than once, the root may not be unique; the design then reports NPV only. In the reference run 41,757 projects were appraised and 574 have a positive NPV at 12 %. Effects that are not monetised — access, sector connection, resilience — are reported in their own units as cost-effectiveness ratios:

$$CE_{j,k} = \frac{PV(C_j)}{\Delta U_k(j)} \tag{24.56}$$

the present cost per person brought within the standard for class k. They are never converted into money, and a cost-effectiveness ratio is never used as a ranking on its own.

24.13 Method: criticality and resilience (TI-05)

Removing link l from the network and recomputing access gives three separate measures, never summed:

$$iso_l = \sum_{i,k} P_i \, \mathbf{1}\left[tt_{i,k}^{-l} = \infty \wedge tt_{i,k} < \infty\right] \tag{24.57}$$

$$\Delta tt_l = \sum_{i,k} P_i \left(tt_{i,k}^{-l} - tt_{i,k}\right) \mathbf{1}\left[tt_{i,k}^{-l} < \infty\right] \tag{24.58}$$

$$detour_l = \frac{d^{-l}(u,v)}{d(u,v)} \tag{24.59}$$

where the superscript −l denotes a value computed without link l. The first counts the people cut off; the second the person-minutes lost by those who can still travel; the third the ratio of the detour to the original distance.

Links whose loss disconnects the network — cut edges, or bridges in graph theory — are found by Tarjan's algorithm in time proportional to the size of the graph. A depth-first search gives each vertex a discovery time disc(v) and a low-link value low(v), the earliest discovery time reachable from v's subtree through one back edge. A tree edge (u, v) is a cut edge exactly when:

$$low(v) > disc(u) \tag{24.60}$$

The reference run found 65,548 cut edges with population behind them. The largest isolates 716,686 persons, in Sironko. For adaptation appraisal (a later phase), the expected annual damage over return periods RP_1 < … < RP_K, with exceedance probabilities p_k = 1/RP_k and damages D_k, is:

$$EAD_l = \sum_{k=1}^{K-1} \tfrac{1}{2}\left(D_k + D_{k+1}\right)\left(p_k - p_{k+1}\right) \tag{24.61}$$

24.14 Method: spatial statistics (GIS)

Chapter 21 gives the spatial statistics in full. For reference, with z_d = x_d − x̄ and row-standardised queen weights w_de, the global and local Moran statistics are:

$$I = \frac{n}{S_0} \cdot \frac{\sum_d \sum_e w_{de} z_d z_e}{\sum_d z_d^2} \tag{24.62}$$

$$I_d = \frac{z_d}{m_2} \sum_e w_{de} z_e \tag{24.63}$$

Inference is by 999 permutations with a fixed seed. The population-weighted Gini coefficient of travel time, sorted in O(n log n), is:

$$Gini = 1 - \sum_i p_i \left(S_i + S_{i-1} \right) \tag{24.64}$$

where settlements are sorted by travel time, p_i is settlement i's population share and S_i the cumulative share of total population-weighted travel time up to and including i. This is the computational form of Equation (21.6).

24.15 Method: resource allocation (TI-08)

24.15.1 Rule-based scenarios

Each rule-based scenario is a greedy, multi-year knapsack. Maintenance is funded first. Then the scenario's benefit b_j is computed in its own unit for each priced project, and projects are funded in order of benefit per unit of central capital cost:

$$\rho_j = \frac{b_j}{X_j^{cap}} \tag{24.65}$$

A project is funded if its cash profile fits the remaining eligible pools in every year from its earliest start; otherwise later starts are tried. After each funding, the benefits of projects that share settlements with it are recomputed, because a settlement served once cannot be served again. The engine uses a lazy heap so that only the top of the ranking is recomputed. Scenario C credits the rural-access population, E the health and school populations, D the net present value of projects tagged for agriculture, minerals or industry, G the flood-exposed critical links, and B a fixed works order (rehabilitation and periodic maintenance before spot improvements before upgrades). Scenario A, which should allocate by historical shares, uses a fallback ordering because no historical allocation data are available.

Scenario F first sets district envelopes in proportion to unserved need, weighted towards districts with low coverage:

$$share_d = \frac{U_d \left(U_d / R_d \right)^{\alpha}}{\sum_{d'} U_{d'} \left(U_{d'} / R_{d'} \right)^{\alpha}} \tag{24.66}$$

where U_d is the rural population without all-season access in district d, R_d the district's rural population and α = 1. Scenario C's rule is then applied within each district envelope.

24.15.2 The optimiser

The optimiser is a mixed-integer linear programme solved by HiGHS. Its decision variables are x_{j,t} ∈ {0, 1} (project j starts in year t), f_{j,π,t} ≥ 0 (cash to j from pool π in year t), z_ω ∈ [0, 1] (enabling set ω complete) and s_g ∈ [0, 1] (settlement-class group g newly served). The constraints built for the Uganda run are:

$$\sum_{t} x_{j,\,t} \le 1 \tag{24.67}$$

$$\sum_{\pi \in E_j} f_{j,\,\pi,\,t} = \sum_{\tau} x_{j,\,\tau}\, cash_{j,\,t-\tau+1} \tag{24.68}$$

$$f_{j,\,\pi,\,t} = 0 \quad \forall \pi \notin E_j \tag{24.69}$$

$$\sum_{j} f_{j,\,\pi,\,t} \le B_{\pi,\,t} - h_{\pi,\,t} \tag{24.70}$$

$$z_\omega \le y_j \quad \forall j \in \omega, \qquad s_g \le \sum_{\omega \in \Omega_g} z_\omega \tag{24.71}$$

The first allows each project to start at most once. The second ties the cash drawn from the pools to the S-curve of a started project. The third enforces pool eligibility and earmarks: E_j is the set of pools allowed to fund project j's works type. The fourth is the envelope: spending from each pool in each year cannot exceed the pool B_{π,t} after routine maintenance h_{π,t} is deducted. The fifth links projects to people: an enabling set is complete only when all its projects are complete (y_j is the sum over start years of x_{j,t}), and a group is served only if at least one of its enabling sets is complete. A start whose cash would run past the horizon is forbidden.

These coverage sets are what let the optimiser value network effects without routing inside the solver. For each baseline-unserved settlement and class, TI-07 records the minimal sets of projects whose joint completion brings it within the standard. In the reference run there are 566,392 enabling sets for 144,746 settlement-class groups.

24.15.3 Measures, floors and the frontier

Three measures are available, each in its own unit: people served for a class k, net present value, and km of condition programmes:

$$served_k = \sum_{g} pop_g\, s_g, \quad npv = \sum_{j} NPV_j\, y_j, \quad cond = \sum_{j \in PRG} km_j\, y_j \tag{24.72}$$

The user declares one objective and may add floors on the others. As implemented, floors are soft and stated against each measure's attainable maximum f^max, the optimum of that measure alone. The problem solved is:

$$\max \ \tfrac{f_1}{f_1^{max}} - \Lambda \sum_{q} r_q \quad \text{s. t.} \quad f_q \ge \varepsilon_q (1 - r_q),\ 0 \le r_q \le 1 \tag{24.73}$$

where ε_q = ϕ_q f_q^max is a relative floor and Λ = 10. The objective term lies between 0 and 1, so the penalty dominates it and a floor is given up only when no allocation within the envelope can meet it. The relative shortfall r_q is reported for every floor, and any r_q > 0 is logged as a finding. In the reference run, OPT_RAI_EMONC80 (maximising people served by all-season roads with an EmONC floor at 80 % of its attainable maximum) fell 92.3 % short of its EmONC floor. The frontier sweeps ϕ over 0, 0.25, 0.5, 0.75 and 1.

The design's fuller formulation adds dependencies between projects, exclusions, contracting capacity, monotonic access over time, district access floors and class weights. It also specifies the augmented ε-constraint method, which adds a small slack reward to avoid weakly dominated frontier points:

$$\max \; f_1 + \delta \sum_q \frac{sl_q}{r_q} \quad \text{s. t.} \quad f_q - sl_q = \varepsilon_q, \; sl_q \geq 0 \tag{24.74}$$

where δ is between 10^{-6} and 10^{-3} and r_q here is the range of f_q. These elements are not yet built. Three properties of the formulation are tested: raising the envelope never worsens the objective, tightening a floor never improves it, and a zero envelope reproduces the do-nothing path.

24.16 Method: outcomes, mode choice and budget (TI-09, TI-07, TI-10)

24.16.1 Outcome simulation

The network state — standard, roughness, gravel thickness and seasonal availability of every link — evolves year by year:

$$\Sigma_{t+1} = \Phi(\Sigma_t, works_t, m_t) \tag{24.75}$$

Links in completed works are reset to their new standard, maintained links follow the maintained path and unmaintained links the do-nothing path. Access is recomputed on the new state. Because the gains of projects interact — two projects may serve the same settlement, or one may be useless without the other — the engine reports an interaction adjustment for each outcome family ϕ:

$$IA_\varphi = \frac{\sum_j B^{iso}_{j,\varphi} - B^{net}_\varphi}{\sum_j B^{iso}_{j,\varphi}} \tag{24.76}$$

where B^{iso}_j is project j's benefit evaluated alone and B^{net} the benefit of the funded set evaluated together on the network. A warning is raised when the absolute value exceeds 0.15. In the reference run the adjustment for the RAI ranges from 0.21 (OPT_NPV_RAI50) to 0.87 (B). The sum of isolated benefits therefore overstates the network gain by a wide margin in every scenario, which is why outcomes are always taken from the network simulation.

24.16.2 Road or rail for bulk flows

For a bulk flow F (t a year) over distance ℓ (km), the annualised cost of each mode uses the capital recovery factor:

$$CRF(r, n) = \frac{r(1+r)^n}{(1+r)^n - 1} \tag{24.77}$$

$$AC_{road}(F) = CRF\, X^{cap}_{road} + M_{road} + F\ell\, \frac{voc_{truck}}{payload} + F\ell\, \frac{EF^{eff}_{truck}}{payload}\, mdc \tag{24.78}$$

where mdc is the marginal pavement damage cost per ESA-km. The rail cost has the same form, with rail operating cost and transhipment at both ends. The break-even tonnage F* solves AC_road(F*) = AC_rail(F*). This comparison needs mineral flows, which were not supplied, so it was not run.

24.16.3 Fiscal years and the fiscal loop

Plan-year cash is converted to fiscal years (July to June in Uganda) with monthly phasing weights ϖ:

$$y_{\phi}^{FY} = \sum_{t} y_t \sum_{m \in t \cap \phi} \varpi_{t,m} \tag{24.79}$$

The conversion preserves totals: the sum over fiscal years equals the sum over plan years, and the engine asserts it. The three budget views (programme, economic and administrative) must reconcile; in the reference run each totals 20,894.9 bn UGX for scenario B. The design specifies a bounded fixed-point loop with FINEX, in which TIPRAS proposes a requirement, FINEX returns a revised envelope, and the loop stops when the envelope changes by less than 0.5 %:

$$\frac{\|B^{(\nu+1)} - B^{(\nu)}\|_1}{\|B^{(\nu)}\|_1} < 0.005 \tag{24.80}$$

or after five iterations. The loop is not yet built: TIPRAS writes the requirement file for FINEX, but no FINEX reader returns a revised envelope.

24.16.4 Uncertainty

Three bounded sources are propagated by re-running with their low and high values: unit cost (the UCSN or evidence triple), driver growth and the gravity parameter (±25 % when uncalibrated). Results are reported as ranges over the corner combinations, not as probabilities. In the reference run only the cost range is propagated, giving the capital range of 56.6 to 125.3 tn UGX.

24.17 Implementation map

Table 24.2 summarises which models are implemented as specified, which in a simplified form, and which not at all in the reference run.

Table 24.2: Implementation status of the mathematical models in the Uganda reference run

Model	Status	Note
Population raking, group shares, production zones	Implemented	Group shares are national constants
Season rule, free-flow time, walking leg	Implemented	Condition unknown on every link
BPR volume–delay	Specified	Loading is all-or-nothing
Multi-source Dijkstra, access test, RAI	Implemented	RAI reported as bounds for unknown surface
Potential accessibility	Specified	Descriptive only
Freight generation, Furness gravity, logit chain	Implemented	β and θ assumed; no Hyman calibration
Passenger trip generation	Partial	No FINEX income path
AON loading, AADT, PCU	Implemented	Incremental assignment not built
ESA and design traffic	Simplified	ESA per vehicle from norms; design life and direction fixed in code
Required standard (join, binding test, structural promotion)	Implemented	Unknown surface ranked as earth
Deterioration (paved, gravel, Markov)	Implemented, not applied	No link condition in the estate
Capacity and V/C	Implemented	Lanes unknown on most links
Queue model	Not built	No gateway flows
Gap definitions G1–G7	Implemented	G2 from agency statistics; G1 cause cannot separate missing links
Works cost, overheads, S-curve	Implemented	E5 evidence and E7 UCSN illustrative prices
Residual value, lifecycle DP	Not built	—
Maintenance requirement, induced maintenance, backlog	Implemented	Backlog from agency poor km
VOC, time benefits, NPV, BCR, EIRR	Simplified	Linear VOC unused (no IRI); wet closure placeholder

Model	Status	Note
Safety, emissions, switching values	Not built	—
Criticality, Tarjan cut edges	Implemented	EAD is a later phase
Moran's I, LISA, concentration index, Gini	Implemented	Shared with HPRAS
Greedy scenarios A–G	Implemented	A uses a fallback ordering
MILP optimiser, soft floors, frontier	Partial	C1, C4–C6, C13–C14 built; dependencies, capacity and AUGMECON not
Network simulation, interaction adjustment	Implemented	—
Bulk mode comparison	Not run	No mineral flows
Fiscal-year conversion, three views	Implemented	Fiscal loop with FINEX not built
Uncertainty ranges	Partial	Cost range only

24.18 Data quality and limitations

- **The models are only as good as their norms.** Of the 646 rows of the draft norm set, 68 are *normed* and 578 *assumed*. Every Greek letter in this chapter is a row in that set, and most are illustrative. The run is exploratory.
- **Missing network attributes switch models off.** No link has a recorded condition, roughness or axle limit, and 68.3 % of km has unknown surface. The deterioration models, the VOC term of the appraisal and the axle part of the connection test therefore have no effect in this run. The engine never fills these attributes by inference.
- **Several parameters are uncalibrated.** The gravity and logit parameters, the BPR parameters, the vehicle operating costs and the deterioration coefficients are TransportModelling or illustrative values, not estimated for Uganda. The ±25 % gravity range in the uncertainty design reflects that.
- **Two coefficients are fixed in code.** The design life of 15 years and the directional factor of 0.5 used for design traffic are set in the TI-03 source rather than read from a norm. They should move to the norm table, as the engine's rules require.
- **Demand is fixed.** The appraisal holds demand constant ($T^1 = T^o$), so generated traffic and the rule of a half have no effect. Benefits are time savings only, valued with an assumed value of time.
- **The optimiser is partial.** Without dependencies, exclusions or contracting capacity, the optimiser can fund combinations that could not be built in sequence or by the available contractors. Its results should be read beside the rule-based scenarios, not instead of them.

24.19 Implications for planning

- **Every number in the report can be traced to an equation in this chapter.** A planner who questions a figure can find its definition here, its parameters in the norm set and its value in the run tables.
- **Calibrate the parameters that drive the money.** The value of time, the discount rate, the vehicle operating costs and the gravity parameter decide which projects have a positive NPV and how freight is spread. A national VOC study and observed trip lengths would replace the most influential assumptions.
- **Supply the link attributes the models need.** Condition, roughness and surface by link would switch on the deterioration models, the VOC benefits and a true G2 condition gap. Without them, a large part of the specified mathematics is inactive.
- **Read the access gains from the network simulation, never from sums of project gains.** The interaction adjustment of 0.21 to 0.87 shows how far isolated project benefits overstate what a funded set achieves together.
- **Complete the optimiser before relying on it for sequencing.** Dependencies, exclusions and contracting capacity are needed before an optimiser schedule can be handed to an implementing agency.

- **Keep the seven gaps separate in every model.** The formulation never adds persons to km or tonnes to minutes. Trade-offs between gaps are shown as frontiers and floors, and the choice among them remains a policy decision.

ANNEX A

A Annex: district summary

Table A.1 gives one row for each of the 135 districts of the reference run 014c9cd7, in alphabetical order. The columns are kept separate on purpose: each gap dimension is shown in its own unit, and no column combines them.

- **RAI today** is the Rural Access Index in the base year 2026, the share of the rural population within 2 km of an all-season road, computed with unknown surface treated as earth (the lower bound). The national value is 28.7 %; district values run from 0 % (Buvuma and Kalangala) to 83.0 % (Manafwa).
- **RAI 2030 (C)** is the same index at the 2030 horizon under scenario C (rural access). It is shown for comparison only. Scenario C is one of the compared scenarios, not a recommended one.
- **G1 beyond 2 km** is the population without an all-season road within the standard walking distance (2 km for rural settlements, 0.5 km for urban ones), and **G1 EmONC** the number of women aged 15–49 more than 2 hours from emergency obstetric care (persons).
- **G4 km not all-season** is the km of links that must be all-season but close in the wet season.
- **G5 isolated pop** is the largest population that the loss of a single link would cut off (persons).
- **G7 exposed pop** is the population that depends on flood-exposed critical links without a reasonable detour (persons). Only Ntoroko, Nakaseke and Napak have a non-zero value.
- **Capital (bn UGX)** is the central capital cost of the candidate projects located in the district, at evidence (E5) and illustrative UCSN (E7) prices, including design, supervision and contingency. The district rows add up to 85.5 tn UGX, with 205.2 bn UGX a year of induced routine maintenance once all projects are complete. The 7.25 tn UGX of national condition programmes, priced from agency statistics rather than by link, is not assigned to districts.

Every value is exploratory: 68.3 % of network km has unknown surface, condition is known only as agency statistics by network, and the norms are a draft. A district's capital figure is the cost of closing its gaps to the draft standards, not an allocation; the allocations by scenario are in Chapter 11.

Table A.1: District summary: access, gaps and the candidate programme

District	RAI today	RAI 2030 (C)	G1 beyond 2 km	G1 EmONC	G4 km not all-season	G5 isolated pop	G7 exposed pop	Capital (bn UGX)
Abim	12.5%	58.5%	176,725	17,917	4	26,431	0	1,024.0
Adjumani	12.8%	67.2%	236,644	7,618	88	7,591	0	947.1
Agago	2.7%	69.0%	297,516	43,166	40	3,250	0	1,596.0
Alebtong	20.4%	53.7%	262,492	1,512	0	5,225	0	682.0
Amolatar	22.5%	54.5%	175,438	2,228	0	6,583	0	504.8
Amudat	3.0%	40.7%	167,036	25,682	21	24,584	0	987.8
Amuria	17.2%	33.3%	232,530	3,493	4	6,855	0	657.8
Amuru	27.9%	67.1%	187,055	17,941	59	6,763	0	1,284.6
Apac	12.7%	54.3%	245,561	33,411	15	2,920	0	908.1
Arua	29.7%	77.7%	578,375	3,543	86	19,495	0	1,140.0
Budaka	45.1%	88.7%	175,165	2,031	2	5,126	0	522.6
Bududa	79.5%	90.5%	143,265	17,465	2	13,480	0	350.8
Bugiri	28.2%	62.5%	363,254	2,624	8	9,040	0	621.9
Bugweri	39.0%	77.4%	131,317	874	0	13,496	0	261.2
Buhweju	2.9%	27.9%	174,591	1,238	0	9,800	0	93.6
Buikwe	29.8%	45.2%	336,805	1,209	2	107,012	0	422.7
Bukedea	29.1%	79.3%	233,020	4,144	0	9,664	0	511.4
Bukomansimbi	13.5%	48.4%	161,179	2,498	1	4,598	0	96.1
Bukwo	12.0%	71.2%	134,444	16,744	27	153,287	0	269.1
Bulambuli	26.8%	74.2%	236,611	8,218	1	711,139	0	208.7
Buliisa	44.2%	48.4%	119,537	82	17	11,292	0	275.4
Bundibugyo	43.3%	99.1%	172,000	7,513	41	295,449	0	401.7
Bunyangabu	31.6%	89.5%	162,357	4,438	66	5,474	0	488.7
Bushenyi	44.3%	58.1%	161,290	3,424	11	4,690	0	422.1
Busia	35.6%	90.3%	266,520	703	14	7,908	0	534.4
Butaleja	16.1%	60.6%	314,776	4,882	30	8,735	0	429.4
Butambala	50.2%	69.4%	60,098	1,960	0	11,716	0	140.3
Butebo	38.2%	94.3%	102,635	469	0	3,197	0	219.6
Buvuma	0.0%	6.6%	171,109	14,596	0	107,001	0	0.2
Buyende	7.8%	34.4%	490,376	10,104	4	39,395	0	279.6
Dokolo	19.7%	59.6%	214,956	2,746	4	11,575	0	537.6
Gomba	23.4%	54.3%	156,928	3,008	6	4,978	0	558.0
Gulu	21.9%	77.3%	234,536	7,939	26	7,091	0	742.3
Hoima	43.0%	74.7%	251,056	2,826	2	5,651	0	332.7
Ibanda	29.7%	89.2%	226,514	5,340	1	2,632	0	617.9
Iganga	58.8%	77.6%	174,409	3,710	6	12,468	0	170.5
Isingiro	17.3%	76.6%	619,309	6,758	20	6,463	0	1,795.9

District	RAI today	RAI 2030 (C)	G1 beyond 2 km	G1 EmONC	G4 km not all-season	G5 isolated pop	G7 exposed pop	Capital (bn UGX)
Jinja	30.8%	85.6%	348,648	952	1	11,094	0	712.2
Kaabong	5.1%	74.0%	144,606	3,364	60	12,263	0	473.1
Kabale	35.5%	81.2%	175,773	11,863	11	12,455	0	603.2
Kabarole	40.2%	91.6%	221,692	4,729	36	357,373	0	711.9
Kaberamaido	14.3%	60.9%	144,027	10,106	0	9,485	0	312.1
Kagadi	17.1%	98.9%	443,669	5,661	3	9,277	0	1,262.5
Kakumiro	6.0%	98.9%	616,800	4,645	2	8,682	0	1,706.0
Kalaki	20.2%	70.8%	139,731	332	0	19,484	0	564.8
Kalangala	0.0%	39.2%	84,690	8,766	0	9,436	0	4.3
Kaliro	24.8%	56.9%	268,456	5,581	8	11,426	0	377.8
Kalungu	19.9%	53.9%	171,223	5,409	8	4,220	0	108.0
Kampala	82.6%	100.0%	197,586	6,721	5	17,213	0	637.0
Kamuli	34.3%	51.8%	421,161	13,147	10	11,702	0	598.3
Kamwenge	15.7%	95.2%	376,805	2,340	12	3,191	0	1,148.2
Kanungu	51.6%	96.9%	159,711	5,678	115	48,808	0	1,353.5
Kapchorwa	76.4%	84.0%	33,876	666	11	5,911	0	24.3
Kapelebyong	2.1%	57.0%	126,218	16,755	33	10,045	0	592.1
Karenga	0.1%	39.6%	83,347	8,797	120	9,304	0	472.0
Kasese	52.7%	94.1%	401,785	21,509	225	13,104	0	2,065.5
Kassanda	17.4%	92.6%	299,935	2,984	5	18,192	0	1,646.5
Katakwi	34.9%	38.7%	153,848	1,812	36	15,128	0	604.8
Kayunga	31.5%	36.6%	316,721	3,780	33	16,867	0	513.9
Kazo	16.5%	66.8%	233,159	8,956	12	5,421	0	490.1
Kibaale	8.1%	97.7%	240,448	3,615	2	7,888	0	1,097.3
Kiboga	33.4%	72.9%	132,158	2,163	1	18,192	0	86.1
Kibuku	42.5%	54.2%	180,033	7,055	2	18,814	0	151.1
Kikuube	26.7%	79.8%	353,202	3,768	23	11,826	0	724.2
Kiruhura	19.8%	63.7%	184,214	7,068	0	2,674	0	655.9
Kiryandongo	37.6%	53.9%	228,319	1,104	21	6,145	0	653.7
Kisoro	33.4%	93.9%	248,895	4,404	41	367,166	0	804.4
Kitagwenda	31.4%	93.7%	139,394	5,217	7	3,774	0	701.1
Kitgum	25.8%	68.1%	179,433	9,864	49	9,021	0	769.1
Koboko	17.1%	70.6%	246,861	5,028	18	40,825	0	223.6
Kole	8.2%	96.9%	325,179	5,539	2	6,385	0	1,317.7
Kotido	4.6%	62.4%	218,494	54,928	145	10,130	0	1,045.7
Kumi	33.9%	74.7%	229,442	856	0	11,571	0	459.7
Kwania	3.7%	64.1%	258,451	45,262	0	5,904	0	906.9
Kween	48.6%	66.5%	69,735	2,877	18	173,795	0	235.0
Kyankwanzi	17.9%	74.8%	301,499	4,180	5	6,991	0	1,011.1

District	RAI today	RAI 2030 (C)	G1 beyond 2 km	G1 EmONC	G4 km not all-season	G5 isolated pop	G7 exposed pop	Capital (bn UGX)
Kyegegwa	12.4%	68.2%	526,602	2,885	33	7,799	0	524.8
Kyenjojo	19.3%	91.8%	533,226	8,640	8	6,198	0	1,195.6
Kyotera	38.2%	85.3%	177,326	3,911	33	22,217	0	663.5
Lamwo	25.5%	65.0%	127,901	3,599	52	13,167	0	565.5
Lira	23.0%	72.5%	433,536	4,134	13	8,712	0	849.8
Luuka	31.0%	54.5%	217,824	1,979	0	7,330	0	562.7
Luwero	27.8%	67.9%	416,757	5,689	0	8,382	0	596.6
Lwengo	28.8%	51.2%	229,413	1,885	0	9,187	0	97.1
Lyantonde	10.7%	31.4%	112,566	2,101	0	3,726	0	59.3
Madi Okollo	17.5%	72.5%	174,754	5,894	33	7,710	0	723.2
Manafwa	83.0%	99.6%	37,771	4,697	12	3,305	0	265.6
Maracha	51.1%	80.4%	124,547	1,189	7	3,943	0	187.9
Masaka	38.2%	73.3%	216,743	5,138	3	9,610	0	291.1
Masindi	41.9%	68.7%	231,742	2,700	14	14,193	0	451.5
Mayuge	16.7%	40.1%	570,505	13,783	7	8,748	0	243.7
Mbale	81.5%	95.4%	178,991	2,075	5	33,340	0	754.2
Mbarara	38.1%	72.9%	243,723	3,226	1	16,142	0	611.8
Mitooma	7.7%	64.8%	214,809	1,814	35	5,319	0	324.0
Mityana	33.4%	76.7%	266,609	3,330	3	12,089	0	831.2
Moroto	15.2%	35.0%	113,668	14,222	51	8,285	0	405.7
Moyo	40.8%	81.4%	62,623	2,681	33	79,923	0	329.3
Mpigi	40.2%	68.3%	188,418	13,268	4	10,158	0	374.7
Mubende	21.3%	86.5%	546,880	29,669	33	8,982	0	1,258.8
Mukono	46.6%	65.9%	410,667	14,263	4	23,123	0	594.9
Nabilatuk	14.4%	28.1%	100,407	3,143	17	19,115	0	365.4
Nakapiripirit	11.1%	28.4%	125,984	6,055	12	36,235	0	356.2
Nakaseke	25.0%	64.8%	208,717	4,973	8	12,134	12,089	850.6
Nakasongola	30.9%	60.5%	177,847	18,287	9	4,129	0	411.9
Namayingo	7.0%	70.9%	294,299	9,573	2	16,289	0	371.6
Namisindwa	52.2%	85.8%	150,546	8,165	33	11,416	0	262.3
Namutumba	27.2%	50.9%	269,892	1,752	5	18,641	0	323.7
Napak	27.5%	32.8%	138,621	12,919	11	24,308	6,745	577.7
Nebbi	40.8%	97.8%	202,591	19,802	166	8,747	0	965.1
Ngora	22.0%	69.0%	160,760	2,933	3	31,577	0	288.6
Ntoroko	39.4%	92.1%	58,936	12,590	463	356,701	298,306	662.3
Ntungamo	30.0%	57.4%	458,958	2,607	0	9,985	0	975.2
Nwoya	27.5%	62.1%	237,582	11,439	30	7,094	0	720.3
Obongi	19.5%	59.5%	25,342	2,600	53	16,342	0	179.6
Omoro	26.7%	63.7%	180,109	5,738	11	20,698	0	496.5

District	RAI today	RAI 2030 (C)	G1 beyond 2 km	G1 EmONC	G4 km not all-season	G5 isolated pop	G7 exposed pop	Capital (bn UGX)
Otuke	7.9%	30.2%	157,104	9,601	8	25,473	0	425.9
Oyam	24.6%	76.9%	423,515	4,393	12	4,803	0	1,467.1
Pader	15.1%	90.4%	201,789	18,084	29	5,316	0	1,036.2
Pakwach	14.2%	74.7%	206,711	14,377	16	20,512	0	456.8
Pallisa	25.1%	84.5%	306,894	650	0	14,531	0	574.8
Rakai	10.4%	82.7%	350,328	7,738	28	5,806	0	1,116.6
Rubanda	33.3%	97.0%	173,499	13,692	13	368,252	0	584.7
Rubirizi	35.5%	99.3%	110,130	3,478	34	6,116	0	507.6
Rukiga	15.6%	63.3%	103,772	1,531	4	2,729	0	466.0
Rukungiri	40.3%	93.2%	235,166	4,173	93	16,609	0	1,432.8
Rwampara	22.1%	64.3%	133,966	279	0	26,038	0	176.9
Serere	55.1%	60.8%	204,439	736	8	27,574	0	844.3
Sheema	27.4%	42.7%	187,978	1,150	1	3,960	0	410.1
Sironko	66.3%	87.3%	119,524	1,079	5	716,686	0	267.2
Soroti	33.9%	82.6%	275,970	2,150	1	4,596	0	928.1
Ssembabule	8.1%	45.7%	331,672	20,912	8	8,898	0	945.1
Tororo	34.7%	88.0%	461,081	1,426	7	15,144	0	980.4
Wakiso	47.8%	76.6%	1,709,126	13,904	10	20,798	0	1,320.0
Yumbe	18.9%	68.2%	703,292	24,182	39	30,875	0	1,024.1
Zombo	59.4%	96.8%	129,232	4,672	87	10,042	0	1,113.8

Source: AICTEM TIPRAS engine, reference run 014c9cd7 (report/build/data/a_districts.csv).

ANNEX B

B Annex: glossary and abbreviations

This annex explains the abbreviations (Table B.1) and the technical terms (Table B.2) used in the report. Where a term has a precise definition in the engine, the definition given here is the one the engine uses, and the chapter that states it in full is named.

Table B.1: Abbreviations

Abbreviation	Meaning
AADT	Annual average daily traffic: vehicles per day, averaged over the year
AC	Asphalt concrete, a sealed paved surface
AID	Africa Infrastructure Data extracts held by TransportModelling (points of interest, wealth indices)
AON	All-or-nothing assignment: each flow is loaded entirely onto its cheapest path
ATDGR	Africa Transport Data and GIS Repository: the estate's read-only transport inventory
AUGMECON	Augmented ε-constraint method for tracing a Pareto frontier
BCR	Benefit–cost ratio: present value of benefits divided by present value of costs
BPR	Bureau of Public Roads volume–delay function
CDCGE	The estate's computable general equilibrium model
COD-AB	Common Operational Dataset of administrative boundaries
CODPS	Common Operational Dataset of population statistics (UBOS 2022 district totals)
CRS	Coordinate reference system
DBST	Double bituminous surface treatment (surface dressing), a sealed surface
DEM	Digital elevation model
DGSM	Directorate of Geological Survey and Mines
DUCAR	District, urban and community access roads
E5, E7	Price evidence tiers: E5 an evidence candidate price from contract and budget documents; E7 an illustrative UCSN price
EAD	Expected annual damage
EIRR	Economic internal rate of return
EmONC	Emergency obstetric and newborn care
EPRAS	Education Planning and Resource Allocation System (EducationModellingx)
ESA	Equivalent standard axle: the pavement damage of one 8.16 t single axle
FINEX	The estate's public finance and expenditure model
FX	Foreign exchange rate
FY	Fiscal year (July to June; for example FY2024/25)
FYRR	First-year rate of return
GIS	Geographic information system
GRIP4	Global Roads Inventory Project, version 4
HC II, HC III, HC IV	Health centre levels II, III and IV in the Ugandan health system
HDM-4	Highway Development and Management model, version 4
HeiGIT	Heidelberg Institute for Geoinformation Technology (2024 road surface estimates)
HPRAS	Health Planning and Resource Allocation System (HealthModellingx)
IRI	International Roughness Index, in m/km
ITIS	Integrated Transport Infrastructure and Services, the transport programme of Uganda's budget
IWI	International Wealth Index

Abbreviation	Meaning
JRC	Joint Research Centre of the European Commission (flood hazard maps)
KCCA	Kampala Capital City Authority
LISA	Local indicator of spatial association (local Moran)
LVSR	Low-volume sealed road
MESA	Million equivalent standard axles
MILP	Mixed-integer linear programme
MoES	Ministry of Education and Sports
MoH	Ministry of Health
MoWT	Ministry of Works and Transport
MTEF	Medium-term expenditure framework
NHFR	National Health Facility Register
NPV	Net present value
NSO	National statistics office (UBOS in Uganda)
OD	Origin–destination
OSM	OpenStreetMap
PCU	Passenger car unit
PM	Periodic maintenance (regravelling, resealing)
PV	Present value
RAI	Rural Access Index (SDG indicator 9.1.1)
RH	Rehabilitation
RM	Routine maintenance
RRH	Regional referral hospital
RWI	Relative Wealth Index
SDG	Sustainable Development Goal
SDI	The estate's spatial data infrastructure (geography, population, noded network, hazards)
SEZ	Special economic zone
SI	Spot improvement (footbridges, vented fords, drainage, short sections)
TIPRAS	Transport Infrastructure Planning and Resource Allocation System, the engine described in this report
TI-01 … TI-10	The ten modules of TIPRAS, from the spatial base to the budget
UBOS	Uganda Bureau of Statistics
UCSN	Unit Cost, Service Norms and Budget Conversion Database: the estate's governed parameter service for unit costs (transport codes TRN-*)
UGX	Uganda shilling (bn = 10^9, tn = 10^{12})
UNRA	Uganda National Roads Authority
UP	Upgrading: raising a link to a higher surface standard
URF	Uganda Road Fund
UTM	Universal Transverse Mercator projection (Uganda: zone 36N, EPSG:32636)
V/C	Volume-to-capacity ratio
VOC	Vehicle operating cost
VoT	Value of time
WorldPop	Gridded population estimates (2020, constrained)

Table B.2: Glossary of technical terms

Term	Meaning in this report
Access standard	The maximum travel time or distance to a destination class that a settlement should meet, for example 2 hours to EmONC. A norm row with a source and a quality tier

Term	Meaning in this report
Agricultural production zone	A contiguous group of cells whose crop production density exceeds a threshold; a freight origin if its marketed surplus is at least 5,000 t a year
All-season road	A road that is open to motorised traffic throughout the year, including the wet season
Assumed	Quality tier of a value set by the analyst without a sourced Ugandan standard; to be replaced
Basis	The field on every value that records how it was obtained (observed, estimated, modelled, assumed)
Binding test	The test (function, traffic, sector node, access or mode) that sets a link's required surface
Candidate project	A package of works on adjacent links proposed by the decision rule of TI-07, with its cost, induced maintenance and appraisal
Cause of an access gap	Why a settlement is unserved: no destination within reach, a missing link, a substandard network or a deteriorated network
Concentration index	A measure of whether a quantity is concentrated where need is highest; positive means pro-need
Connection standard	The road class, surface and axle load that the route from a sector node to its target must meet
Coverage set (enabling set)	A minimal set of candidate projects whose joint completion brings a settlement within an access standard
Cut edge	A link whose removal disconnects part of the network from the rest
Do-nothing path	The state of the network at the horizon if no capital works are funded
Envelope	The money available for transport in each year and financing pool
Exploratory	Status of a run that uses draft norms, candidate prices or data below the quality gates; its results are not for decisions without review
Factor trace	The stored record of every coefficient applied to a result, with its source, so that any value can be explained
Fallback	A flagged substitute used when a required input is missing, recorded in the run log and on every page that uses it
Financing pool	A source of funds with its own eligible works types, such as government development, donor projects, the Road Fund or district grants
Frontier	The set of best attainable values of one objective for each level of a floor on another
Gap dimension	One of the seven separate measures of shortfall (G1 access, G2 condition, G3 capacity, G4 standard, G5 redundancy, G6 sector connection, G7 climate exposure). They are never combined
Generalised cost	The cost of a trip or a tonne moved, combining time valued in money and vehicle operating cost
Geography snapshot	An immutable, versioned copy of the boundaries, network, settlements and nodes used by a run (here UGA-2026.09-c)
Gravity model	A model that distributes flows between origins and destinations in proportion to their sizes and inversely to the cost between them
Hot spot, cold spot	A district with a high (HH) or low (LL) value among neighbours with similar values, significant in the LISA test
Induced maintenance	The change in annual routine maintenance requirement caused by a completed project; every capital figure carries it
Interaction adjustment	The share by which the sum of projects' isolated benefits overstates the benefit of the funded set on the network
Measured	Quality tier of a value obtained by a survey or count
Modelled	Quality tier of a value computed by the engine from other values
Moran's I	A statistic that tests whether similar values are found next to each other on a map
Norm	A standard, threshold, coefficient or unit-cost mapping held as data in the norm set, never in code
Normed	Quality tier of a value taken from an official or published standard

Term	Meaning in this report
Observed	Quality tier of a value recorded directly in a source dataset
Planning graph	The network of links and vertices on which travel times and flows are computed
Quality tier	The label that records how reliable a value is: observed, measured, normed, imputed, modelled or assumed
Raking	Scaling gridded population so that it adds up to official district totals
Required standard	The minimum standard a link must meet, the field-wise maximum of the traffic, function, sector-node and access tests
Rule of a half	The convention that new traffic attracted by a cheaper route gains, on average, half the cost saving
S-curve	The phasing of a project's cash over its construction years
Scenario	A rule for allocating the envelope (A to G), or an optimiser run named by its objective and floors. Scenarios are compared; none is recommended
Sector node	A point that a sector needs connected: a health facility, school, market, mine, industrial park, gateway or production zone
Sector tag	A label on a candidate project naming each sector whose gap it helps to close; a project can carry several
Settlement	A populated 1 km² cell of the planning grid
Snapping	Joining a settlement or node to the nearest point on the network
Spatial lag	The average value of a district's neighbours
Survey requirement	Links whose surface or structures are unknown and must be surveyed before works can be decided
Unit cost	The cost per km, per lane-km or per unit of a works type, read from UCSN or from the evidence price file
Unpriced	A project with no unit cost for one of its works; it cannot be funded until priced and is never costed at zero
Unserved population	Persons (in the standard's target group) who do not meet an access standard
Works type	The kind of intervention: routine and periodic maintenance, spot improvement, rehabilitation, upgrading, structures, new construction and others