EDIS
Economic Development Intelligence System
Data • Models • Better Decisions

Supporting evidence-based economic policy for a prosperous, resilient and sustainable Africa

FINEX

Forecasting Model of Internal and External Balance

An Open-Source Julia Implementation for Macroeconomic Forecasting, Monetary Policy and Fiscal Policy Analysis

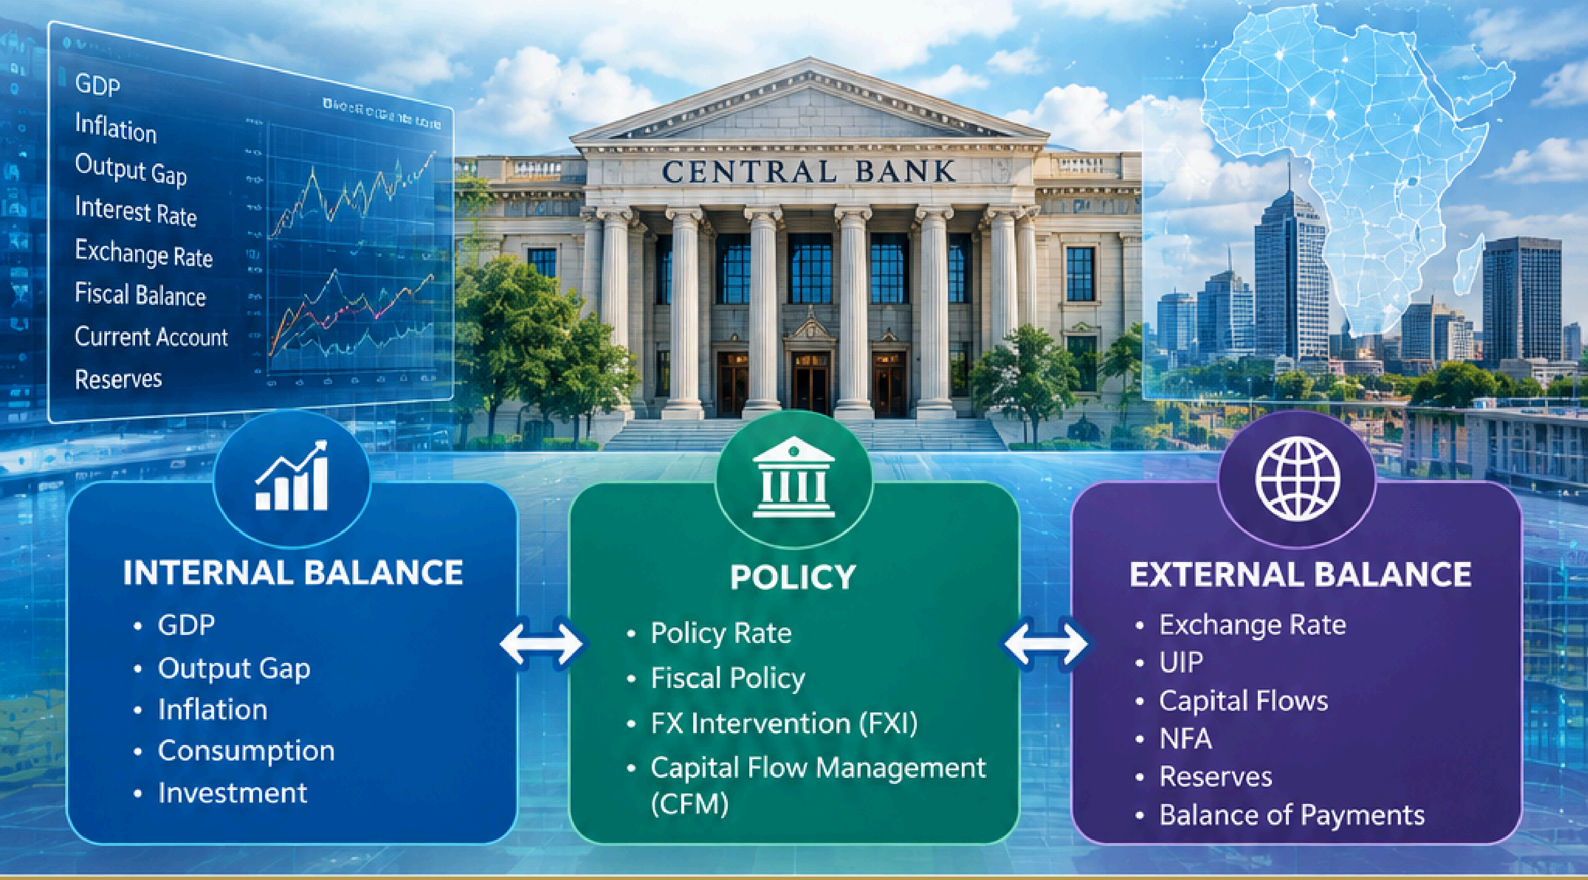

A Three-Country Application: Uganda, South Africa and Kenya

Structure, Estimation, Forecasts and Policy Simulation

Estimated by Simulated Method of Moments on 2001–2024

John-Mary Matovu

jmatovu@gmail.com • edisgov.org

Economic Development Intelligence System (EDIS)

Abstract

We apply FINEX — the IMF's semi-structural model of internal and external balance — to Uganda, South Africa and Kenya, estimating 16 behavioural coefficients and 40 shock variances for each by Simulated Method of Moments on annual data for 2001–2024, using identical code and an identical moment set throughout. All three project to convergent steady states, but they do not fit equally well: Uganda and South Africa reach data–model moment correlations of 0.783 and 0.774 while Kenya reaches 0.191, a weakness traceable entirely to the world oil block, which carries 57 per cent of Kenya's total moment misfit and in which the model produces essentially no variance against a data variance of 549. Re-estimating with that block held out raises Kenya's fit to 0.532 over the remaining 155 moments, but almost none of that is a better estimate: the full-sample parameters already scored 0.528 on those same moments, so the apparent gain is the arithmetic of dropping the thirteen worst conditions rather than evidence of improved fit. We therefore report the full-moment-set estimate throughout, and read 0.191 as a statement about the oil block rather than about the model's grip on Kenya's domestic economy. The estimated structures differ sharply and coherently. Uganda's transmission runs through a strongly interest-sensitive consumption block and a heavily smoothed policy rule; South Africa's through inertial inflation, a steep Phillips curve and an aggressive inflation response; Kenya combines South African inflation persistence with a Ugandan interest channel. Fiscal multipliers are near 0.6. We report baseline projections to 2035, impulse responses to all 83 structural shocks, a dated policy scenario with explicit attribution, and a decomposition of each country's history since 2006. We are explicit about what the exercise cannot establish: roughly half the shock variances are unidentified from a single country's 24 observations. The third case earns its place by isolating an outlier — South Africa's near-zero interest elasticity of consumption, which two countries left ambiguous and three identify as a probable identification failure.

Keywords: semi-structural modelling; forecasting; monetary policy; fiscal policy; Simulated Method of Moments; Uganda; South Africa; Kenya.

Contents

1. Introduction

The macroeconomic questions that matter most in emerging and developing economies are joint questions. Whether a monetary tightening is affordable depends on what it does to the debt ratio as well as to inflation. Whether a fiscal consolidation is self-defeating depends on what it does to growth, and therefore to the denominator of the ratio it is meant to improve. Whether a commodity price shock is a fiscal problem or an external one is not a choice the analyst gets to make; it is both at once, and the two channels interact through the exchange rate and the currency composition of public debt. A model that treats the external accounts as a residual — as the standard quarterly projection model largely does — cannot answer any of them, because the accounting that links them is precisely what it leaves out.

FINEX, the Forecasting Model of Internal and External Balance (Berg et al., 2023), is built to answer them. It is semi-structural: its behavioural equations are written directly in terms of output gaps, inflation rates and policy rules, with coefficients that are calibrated or estimated rather than derived from an explicit optimisation problem. In exchange for that theoretical looseness it buys three things. The balance of payments and the government budget constraint are complete, binding identities rather than residuals, and the stocks they generate — public debt by currency, net foreign assets, reserves — accumulate and feed back. Aggregate demand is disaggregated finely enough that fiscal instruments act on identifiable components, so multipliers differ by instrument as a matter of structure rather than assumption. And potential output depends on private and public capital stocks, so investment — including public investment — shifts the path the economy converges to rather than merely the cycle around it. Policy has level effects, not only stabilisation effects, and a consolidation that falls on public investment is a different object from one that falls on transfers.

A model of that ambition is only useful for a particular country once its parameters have been fitted to that country's data, and this is where the practical difficulty lies. FINEX carries 436 variables and 83 structural shocks. The annual data available for a typical African economy runs to roughly two dozen usable observations. The ratio rules out the two standard approaches. Maximum likelihood requires a well-conditioned likelihood surface, and with 83 shocks and 24 observations that surface is flat in most directions, so the optimiser wanders and the result is whichever local mode it happened to stop at. Bayesian estimation is feasible but converts the problem rather than solving it: with this little data the posterior would largely reproduce the priors, and the honest description of the output would be "the priors, slightly updated" — which is a legitimate exercise, but not one that can be said to have learned much from the country's own experience.

This report therefore estimates FINEX by **Simulated Method of Moments**, and the method is worth setting out at the start because it shapes everything that follows. The estimator matches a chosen set of second moments — the variance and first autocovariance of every observable, plus a set of cross-covariances spanning the demand, price, policy, external and fiscal channels, 166 conditions in all for Uganda and 168 for South Africa. Means are handled separately by the steady-state calibration, so every series is demeaned before the moments are formed: the steady state is identified by sample means and stated policy anchors, and the dynamics by second moments. The virtue of this division of labour is that it is explicit about what the estimate is being asked to reproduce, and therefore explicit about what it fails to reproduce. When the estimation breaks down — as it does here for the world commodity block — the failure appears as a named list of missed moments rather than as an unremarkable likelihood value.

What makes the estimation tractable at all is a structural property of the model rather than a numerical trick. The model's autocovariances turn out to be *exactly linear* in the shock variances, with the coefficients of that linear relationship depending only on the behavioural parameters. Section 5.3 derives this. Its consequence is that the parameter vector splits cleanly: for any trial value of the 16 behavioural coefficients, the 40 shock variances solve a non-negative least squares problem exactly, by an active-set algorithm (Lawson and Hanson, 1974), rather than being searched. Only the 16 coefficients need numerical optimisation, by Nelder–Mead (Nelder and Mead, 1965) over a logistic reparameterisation of their bounds. A 56-dimensional search becomes a 16-dimensional one, each evaluation costing a single model solve and a single least-squares problem, and 272 evaluations suffice. The non-negativity constraint is not a convenience but a requirement: a variance must be non-negative, and imposing it as a constraint rather than by reparameterisation is what allows an uninformative shock to sit cleanly at zero and announce that it is uninformative, rather than being pushed to some small positive number that looks like an estimate.

One restriction on that search is load-bearing and easily missed. Three of the model's gap equations combine a lag and a lead, and stability requires the two coefficients to sum to less than one; at or above one the equation has a root on or outside the unit circle and the entire model is explosive. The SMM objective will not catch this, because a near-explosive impulse response either still yields finite second moments over a truncated horizon or fails outright and is silently skipped. An unrestricted first pass on Uganda scored a *better* objective than the restricted estimate and produced a projection that diverged; the same failure elsewhere produced 252% GDP growth and -135% of GDP public debt by 2035. Goodness of fit and dynamic stability are independent properties of an estimate, and both must be tested. Every result reported here has passed both.

We apply this to three countries. Uganda is low-income, fast-growing, remittance- and aid-supported, with a small government and largely concessional external debt. South Africa is upper-middle-income and slow-growing, with a large government, deep domestic financial markets and market-priced debt. Kenya sits between them on almost every structural ratio, growing at close to Uganda's rate but carrying nearly twice its interest burden. All three were estimated with identical code, an identical moment set and identical methodology, differing only in data and in a short country-specific calibration block — which makes the comparison unusually clean, since differences in the results are differences in the economies rather than in their treatment.

All three project to convergent, well-defined steady states, but they do not fit equally well. Uganda and South Africa reach data–model moment correlations of 0.783 and 0.774; Kenya reaches only 0.191. That the framework transfers at all across economies this different is a result, and not a foregone one — a good moment fit is no guarantee of a usable projection, and the two properties must be established separately. That one of the three fits far less well is equally informative, because the weakness turns out to be diagnosable: every one of Kenya's six largest moment misses is an oil moment, and three of them are exact zeros.

The estimates differ sharply across the three countries, and coherently. Uganda's consumption is strongly interest-sensitive and its central bank smooths heavily, so a modest policy shock produces a large and persistent real effect. South African inflation is markedly more inertial and its Phillips curve markedly steeper, and the Reserve Bank smooths less and responds to inflation roughly twice as strongly — an estimated rule close to a textbook Taylor rule, against Uganda's heavily smoothed and weakly responsive one. Both patterns are consistent with what is independently known about those regimes. Kenya combines South African inflation persistence with a Ugandan interest channel, and it is Kenya that resolves the report's sharpest anomaly: its interest elasticity of 0.59 sits all but on top of Uganda's 0.58 and twenty-eight times above South Africa's 0.02, which turns a puzzle into a diagnosis. Fiscal multipliers, however, come out strikingly similar — an impact effect on growth of 0.62 in Uganda, 0.66 in Kenya and 0.57 in South Africa per standard deviation of the shock. **The third country turns that agreement from a coincidence into a result and removes its usual explanation.** Read on two countries it was natural to attribute the small multiplier to demand leaking into imports in Uganda and to monetary offset in South Africa. Kenya's import persistence is 0.77 against Uganda's 0.004 — a factor of 207 — and its multiplier is the same, so leakage cannot be what makes the multiplier small; and no country's policy rate moves by more than 0.14 percentage points in response to the fiscal impulse, so offset cannot be either. What all three share is the reversal: the first year's growth impulse is given back in the second, leaving a cumulative effect on the level of output of +0.05 %, +0.02 and -0.02 (Section 14). The sharpest contrast is arithmetic rather than behavioural: South Africa raises roughly twice Uganda's revenue as a share of GDP and spends nearly twice as much, yet runs a similar deficit, because the difference is absorbed by an interest bill of 6.5% of GDP against Uganda's 2.9% — on a debt ratio 25 points higher and a growth rate a quarter as fast. In the baseline projection Uganda's debt falls to 48.8% of GDP by 2030 and settles at its 50% anchor, while South Africa's rises to 77.7% before edging back toward its 75% anchor.

We are equally explicit about what the exercise does not establish. With 24 annual observations and 40 shock variances, the non-negative least squares assigns the covariance it can explain to a subset of shocks and drives the collinear remainder to exactly zero — between 15 and 21 of 40 across the three. It bites hardest on the world commodity block: the oil price gap ends with a model variance of exactly zero against a data variance of 549 in Uganda, which is not a poor fit but a refusal to fit. Since the world-price processes are common to every country, the same moments are missed in both countries alike. The remedy is structural and is set out in Section 17: these variances should be estimated once, directly off the foreign series, and held fixed while the domestic shocks are allocated. Three further limitations bound the results — a level mismatch between the national accounts and the balance of payments in the external block, a Ugandan policy-rate series that ends in 2019, and the absence of any off-model information, notably Ugandan oil production. And one specific estimate, South Africa's near-zero interest-rate elasticity of consumption, we report and use consistently while identifying it in

Section 18 as a probable identification problem rather than a finding about the South African economy.

Three conventions run through what follows and are the report's main defence against over-claiming. Every shock scale is labelled as estimated or calibrated, since a response to a calibrated shock has an estimated shape and an assumed magnitude. Cross-country comparisons are made per unit of standard deviation wherever the two countries' shock scales differ, and the text says when it is doing so. And results that are not credible are flagged rather than smoothed over.

One further difference from the published model is practical rather than economic, and it bears on who can use this work. FINEX as published is written for the IRIS Toolbox, which runs on MATLAB under a commercial licence. The implementation used throughout this report is a port to Julia that depends on no proprietary software at any stage — model definition, solution, estimation, simulation, charting and the interactive front end all run on open-source packages, three of them maintained and released by the Bank of Canada. Section 2 sets out what was replaced by what, and what the substitution costs as well as what it buys.

The consequence is that nothing stands between the model and a prospective user except the data. The estimation, the projections, the scenario engine and a web application serving all of them can be installed on an ordinary machine or an unattended server by any central bank or ministry of finance, with no per-seat or per-core licence to buy and no metering to work around; the same code that produced every number in this report is what the application runs. That is a lower barrier than a proprietary toolbox imposes, and for institutions that must rerun the same analysis on new data over many years it also removes a dependency that has to be renewed and budgeted for indefinitely.

It would be wrong, however, to present the licence as the binding constraint. Applying the same pipeline across the continent, we find that data availability, not software cost, is what determines whether a country can be modelled at all: of the fifty-four African economies screened for this project, nineteen have the complete set of series the model requires. Removing the licence widens who *may* run FINEX; it does not by itself widen where FINEX can be estimated.

The remainder proceeds as follows. Section 2 describes the open-source implementation. Section 3 sets out the model, equation by equation. Sections 4 and 5 describe the data and the estimation in detail. Section 6 presents the baseline forecasts to 2035 and the conditional forecasting machinery. Sections 9 to 12 report the policy simulations, including a dated scenario with explicit attribution of the interaction between instruments, and Section 14 examines fiscal policy in depth. Section 13 decomposes both countries' histories since 2006 and Section 15 reads that decomposition as a narrative. Section 18 draws the two-country comparison together, and Section 17 states what the exercise cannot do. Appendices document the observable panels, all 83 shocks, the full parameter set, the steady states, the numerical methods and the code.

2. An open-source implementation: Julia in place of MATLAB

FINEX as published is written for the IRIS Toolbox, which runs on MATLAB. The implementation used throughout this report is a port to Julia (Bezanson et al., 2017) that depends on no proprietary software at any stage: model definition, solution, estimation, simulation, charting and the web application all run on open-source packages. The modelling core rests on `ModelBaseEcon.jl`, `StateSpaceEcon.jl` and `TimeSeriesEcon.jl`, which are maintained by the Bank of Canada and released open source (Bank of Canada, 2024): the substitution is one central bank's proprietary-platform toolbox for another central bank's open-source one, not a move from institutional software to hobbyist software. This section sets out what the translation had to handle and what the substitution buys and costs — consequences, some favourable and one a real cost, that ought to be stated rather than assumed.

2.1 What the translator had to handle

Most of the difficulty is in the IRIS model file, whose preprocessor is a small language of its own. Reproducing it faithfully was the precondition for trusting anything downstream, and the constructs that had to be supported were:

- `!for ?cc = $[ext_cc_list]$ !do ...!end` — country loops, which expand the foreign block over the United States and the euro area.
- `!if / !elseif / !else / !end` — the monetary-regime switches that select between inflation targeting with a float, a managed float, and a fixed rate with partial capital controls. Each regime yields a different model, and all three are generated.
- `{-1}`, `{+1}` lead and lag markers, translated to `[t-1]` and `[t+1]`.
- `&X` steady-state references, translated to `@sstate(X)`.
- `!!` separating an equation's dynamic form from its explicit steady-state form, which become a dynamic equation and a steady-state constraint respectively.
- `diff(X)`, expanded to `(X - X[t-1])`.
- MATLAB's element-wise operators `.*`, `./`, `.^`, which have no separate meaning here and become the ordinary ones.
- Quote-aware `%` comment stripping and `...` line continuation — necessary because equation descriptions contain percent signs and the naive strip corrupts them.

The three smooth-approximation helpers deserve a note because they are where a silent numerical difference could most easily have crept in. IRIS supplies `sabs`, `smax0` and `smin0` as MATLAB functions giving differentiable approximations to $|b|$, $\max(b, 0)$ and $\min(b, 0)$ with a smoothing parameter. They matter: the effective lower bound on the policy rate and the nonlinear risk-appetite terms are written with them, so their exact form determines how sharply those nonlinearities bite. They are reimplemented from the `.m` sources rather than re-derived, and the Julia definitions are one-liners that can be checked against the originals by eye.

2.2 What the substitution buys

Three gains are concrete, and one of them was decisive for this project.

Reproducibility with a pinned environment. Julia records exact package versions in a manifest, so the environment that produced these results can be reconstructed by another person on another machine. The comparative literature identifies precisely this as MATLAB's weakness: the absence of a package management system is described as its "Achilles heel", leaving users to rely on the vendor's own toolboxes, write everything themselves, or share code by unreliable means (Coleman et al., 2021). For an institution that must rerun the same analysis on new data over a decade, that distinction matters more than any difference in language syntax.

No licence between the model and its user. The model, its estimation and its web front end can be installed on any machine, including a server that runs unattended, without a per-seat or per-core licence. That is what makes the job-queue architecture in Appendix E possible: workers can be started and stopped freely because nothing is metered.

One language from model to web page. The estimation, the projection, the scenario engine and the application serving them are the same language and share the same model object. There is no export step between the model and the tool that presents it, and therefore no opportunity for the two to disagree.

2.3 The web application, and the audience it reaches

The same estimated model is also served as a web application, and the reason to describe it here rather than only in Appendix G is that it is a consequence of the substitution rather than an accessory to it. An unlicensed stack can be installed on a server; a server can be pointed at by a browser; and a browser is the only thing a user then needs. **What changes is not what the model can do but how many people can ask it anything.**

Three audiences, three different barriers. The barrier a browser removes is not the same one for each of them.

- **The modeller** is not the beneficiary. They have the code, and the application adds nothing they could not do at a prompt — except a queue that keeps long runs off their own machine and a record of what was run.

- **The analyst who does not write Julia** is. A macroeconomist in a forecasting or debt unit can select a shock, scale it, read the response across every headline variable, compose a dated scenario and export the table to a spreadsheet, without installing anything or reading a line of the model. This is the largest group in most finance ministries and central banks, and the one a command-line pipeline serves worst.

- **The decision-maker** is reached indirectly, through a document. A report generated from the application carries an executive answer, the channels the model attributes the result to, the validation verdict and the assumptions behind it — in a fixed order, so a minister's office learns where to look once.

The question bank is the bridge from a policy question to a model variable. 930 questions drawn from the macro-fiscal work programme are served, organised by 31 topics and 10 originating desks, and each is marked by what FINEX can actually answer: 434 are supported outright and 496 are partial, meaning the question also asks for something the model does not carry — spatial detail, unit costing, expenditure ceilings, poverty microsimulation — in which case the model answers the part it can and states what it left out. A further 120 were removed from the bank because FINEX has no variable for their topic at all (Credit Growth, Food Inflation, Macro-Sector Consistency, Money Supply).

Marking the boundary is the point of that design. A model asked a thousand questions will answer all of them if nothing stops it, and the answers it should not have given are indistinguishable from the rest. Publishing the status alongside the question puts the limit in front of the user before they reach it rather than after.

Speed comes from separating what can be rescaled from what must be solved. Nothing in the browser solves the model. Responses are linear in shock size, so any query that rescales a stored impulse response — a what-if over any combination of the 83 structural shocks, a change of horizon, a different shock size — is exact and returns in under a second; anything that moves the steady state, such as an inflation target, a debt anchor, a tax rate or the policy regime, requires the model to be re-solved and history re-inverted, and is queued. Each control says which kind it is before it is touched, and each answer says whether the linear scaling was still exact at the size requested.

Reach is worth little without reproducibility, so results are addressable. Every view has a URL: a page, a tab within it, a selected question, a stored run, a filtered list and a fully specified quantitative query. Colleagues exchange links to results rather than instructions for reproducing clicks, and every generated answer carries an analysis identifier that regenerates the same numbers from the same frozen evidence package. One process serves all three estimated countries, so the comparison this report makes is available to a user without a second installation.

Finally, the application carries the paper's qualifications rather than hiding them. The identification and convergence verdicts appear on every page that reports a number; a shock whose standard deviation the estimation could not identify is labelled as such wherever it is used, so that the distinction between the shape of a response and its size survives contact with a user who has not read Section 11. **An interface that made this model easier to use without making its limits visible would be a disservice**, and the 10 pages shown in Appendix G are built on that principle.

2.4 What it costs

One cost is real and should not be minimised. Julia's ecosystem is younger than MATLAB's, and the comparative literature notes that its library coverage, while growing quickly and in places already richer than MATLAB's for economics, remains less mature (Coleman et al., 2021). In practice this showed up here not as missing functionality but as fragility around the edges: a package whose precompilation step reached the network and failed on an expired certificate, an optimiser whose interface differs from the MATLAB equivalent in ways that had to be read from source, and start-up compilation latency long enough that the web application keeps long-lived worker processes rather than starting one per request.

None of these is a modelling limitation, and none changed a number in this report. They are the ordinary friction of a younger toolchain, and they are worth stating so that an institution weighing the same substitution can budget for them rather than discover them.

A second and more important qualification is that a port is only as good as its verification. The equations here were checked against Appendix A.IV of Berg et al. (2023) across the whole range — all 353 of them, in a system that solves square — and the variable, shock and parameter lists were checked against the paper's tables. That verification, not the choice of language, is what licenses the results in this report.

3. The FINEX framework

3.1 What kind of model this is

FINEX is a *semi-structural* model. Its behavioural equations are written directly in terms of gaps, inflation rates and policy rules, with coefficients that are calibrated or estimated rather than derived from an explicit optimisation problem. In exchange for that theoretical looseness it buys two things a small gap model cannot offer: a complete and binding set of accounting identities for the balance of payments and the government budget, and a disaggregated demand side in which fiscal instruments act on identifiable components.

Three features distinguish it from the earlier generation of quarterly projection models (Berg et al., 2023):

- **External balance is a first-class object.** The balance of payments is an identity that must hold each period, the current and financial accounts are modelled component by component, and the stock of net foreign assets accumulates from the flows.
- **Fiscal policy is explicit.** Revenue is linked to specific tax bases and specific expenditure components, so multipliers differ by instrument rather than being imposed as a single number.
- **Trends are partly endogenous.** Potential output depends on private and public capital stocks, so investment — including public investment — affects the level of the balanced-growth path, not merely the cycle around it.

The third point deserves emphasis because it is unusual. In most gap models the trend is exogenous scenery against which the cycle plays out. In FINEX a sustained change in interest rates, by changing investment, changes potential output; and a fiscal consolidation that falls on public investment lowers the path the economy converges to. Policy has level effects, not only stabilisation effects.

3.2 Gaps, trends, and the steady state

Every non-stationary variable in FINEX is decomposed into a *trend* and a *gap*. For a real quantity $\mathcal{X}_t$ expressed in logarithms,

$$x_t = 100 \cdot \log \mathcal{X}_t, \qquad x_t = \overline{x}_t + \widehat{x}_t \tag{1}$$

where $\overline{x}_t$ is the trend and $\widehat{x}_t$ the gap. Trend growth is written in log-differences, $\Delta\overline{x}_t = \overline{x}_t - \overline{x}_{t-1}$, which is stationary even though the level is not.

Three distinct concepts must be kept apart, and confusing them is the most common source of misreading model output:

- the **gap** $\widehat{x}_t$ — the business cycle, which closes as stabilising policy returns the economy to trend;
- the **trend** $\overline{x}_t$ — the medium-run path, which itself moves with capital accumulation, relative prices, remittances and fiscal policy;
- the **steady state** — the balanced growth path to which the trend converges, characterised by constant nominal expenditure shares of GDP.

A permanent fall in foreign demand illustrates the difference. Exports fall, the current account deteriorates, and restoring external balance requires a *permanent* real depreciation and a *permanent* rise in the UIP premium. The higher premium raises real interest rates, which discourages investment, which lowers potential output. None of that is a business cycle; all of it happens in the trends.

3.3 The internal balance block

3.3.1 Aggregate demand

Nominal GDP is the sum of the nominal expenditure components, each the product of a real quantity and its own deflator:

$$\begin{aligned} Y_t = {} & P_t^{C,T} C_t^R + P_t^I I_t^R + P_t^{XNNR} X_t^{R,NNR} + P_t^{XNR} X_t^{R,NR} \\ & - P_t^{MNOIL,T} M_t^{R,NOIL} - P_t^{MOIL,T} M_t^{R,OIL} + P_t^G G_t^R \end{aligned} \tag{2}$$

Exports are split into natural-resource ($X^{R,NR}$) and non-natural-resource ($X^{R,NNR}$) components because the two are driven by different things — world commodity prices in the first case, foreign demand and competitiveness in the second. Imports are split into oil ($M^{R,OIL}$) and non-oil ($M^{R,NOIL}$) for the same reason.

The real output gap is the share-weighted sum of the component gaps, with weights given by the trend nominal shares in GDP:

$$\begin{aligned} \widehat{y}_t^R = {} & \frac{\overline{C}_t^Y}{100}\,\widehat{c}_t^R + \frac{\overline{I}_t^Y}{100}\,\widehat{\imath}_t^R + \frac{\overline{X}_t^{NNR,Y}}{100}\,\widehat{x}_t^{R,NNR} + \frac{\overline{X}_t^{NR,Y}}{100}\,\widehat{x}_t^{R,NR} \\ & - \frac{\overline{M}_t^{NOIL,Y}}{100}\,\widehat{m}_t^{R,NOIL} - \frac{\overline{M}_t^{OIL,Y}}{100}\,\widehat{m}_t^{R,OIL} + \frac{\overline{G}_t^Y}{100}\,\widehat{g}_t^R \end{aligned} \tag{3}$$

Potential growth is the same weighted sum applied to trend growth rates. The consumption weight is 0.72 in Uganda, 0.76 in Kenya and 0.63 in South Africa, so equation (3) is overwhelmingly a statement about consumption in all three: between three-fifths and three-quarters of any demand disturbance reaches the output gap through that single channel. The weight is a calibrated share rather than an estimate, which is why it differs across countries without any behaviour differing.

3.3.2 The generic gap and trend equations

Every demand component except government absorption follows the same two-equation template. The gap equation is

$$\widehat{x}_t = \rho_x\,\widehat{x}_{t-1} + (1-\rho_x)\,\mathbb{E}_t\widehat{x}_{t+1} + \textit{determinants}_t^{\widehat{x}} + \varepsilon_t^{\widehat{x}} \tag{4}$$

combining a lag term (real rigidity — habit formation, adjustment costs), a model-consistent expectation, component-specific determinants, and a structural shock. All gaps close in equilibrium, which is a property of the stabilising policy rules rather than an assumption imposed on the gaps themselves.

The trend equation is an AR(1) in the trend growth rate with an error-correction term that pins the level to the balanced growth path:

$$\begin{aligned}\Delta\overline{x}_t = \alpha_1\Delta\overline{x}_{t-1} + (1-\alpha_1)\left(\Delta\overline{y}_t^R + \Delta\overline{p}_t^{R,Y} - \Delta\overline{p}_t^{R,\mathcal{X}}\right)\\ - \alpha_2\left(\overline{\mathcal{X}}_t^Y - \left(\overline{\mathcal{X}}^{Y,SS} - \textit{determinants}_t^{\Delta\overline{x}}\right)\right) + \varepsilon_t^{\Delta\overline{x}}\end{aligned} \tag{5}$$

The second line is what makes the steady state a genuine attractor: if the component's trend share of GDP $\overline{\mathcal{X}}_t^Y$ drifts from its steady-state value $\overline{\mathcal{X}}^{Y,SS}$, trend growth adjusts to bring it back. The *determinants* term allows permanent shifters — a permanently higher tax rate, a permanent change in remittances — to move the destination itself.

3.3.3 Prices and the Phillips curves

Each component carries its own deflator, and each relative price is expressed against consumption and is non-stationary:

$$p_t^{\mathcal{X}} = 100\log P_t^{\mathcal{X}}, \qquad p_t^{R,\mathcal{X}} = p_t^{\mathcal{X}} - p_t^C \tag{6}$$

Relative-price trends matter a great deal in practice. For investment and government goods the trend follows an AR(1) toward a steady-state drift; for traded components it is tied to the component-specific equilibrium real exchange rate:

$$\Delta\overline{p}_t^{R,\mathcal{X}} = \begin{cases}\rho_{\mathcal{X}}\Delta\overline{p}_{t-1}^{R,\mathcal{X}} + (1-\rho_{\mathcal{X}})\Delta\overline{p}^{R,\mathcal{X},SS} + \varepsilon_t, & \mathcal{X}\in\{I^R, G^R\}\\ \Delta\overline{p}^{R,\mathcal{X},SS} + \alpha_{\mathcal{X}}\left(\Delta\overline{z}_t^{\mathcal{X}} - \Delta\overline{z}^{\mathcal{X},SS}\right) + \varepsilon_t, & \mathcal{X}\in\{X^{NNR}, X^{NR}, M^{NOIL}, M^{OIL}\}\end{cases} \tag{7}$$

Domestic inflation rates follow hybrid New Keynesian Phillips curves with an error-correction term in the relative-price gap:

$$\pi_t^{\mathcal{X}} = a_1 \pi_{t-1}^{\mathcal{X}} + (1-a_1)\mathbb{E}_t \pi_{t+1}^{\mathcal{X}} + a_2\, RMC_t^{\mathcal{X}} - a_3\, \widehat{p}_t^{R,\mathcal{X}} + \varepsilon_t^{\pi}, \qquad \mathcal{X} \in \{C^R, I^R, G^R\} \tag{8}$$

while traded-goods prices are driven by world prices converted into domestic currency:

$$\pi_t^{\mathcal{X}} = \left(\Delta \overline{p}_t^{R,\mathcal{X}} + \overline{\pi}_t^{C}\right) + a_4 \left(\pi_t^{\mathcal{X},*} + \Delta s_t^{US} - \left(\overline{\pi}_t^{C} + \Delta \overline{z}_t^{\mathcal{X}}\right)\right) - a_5\, \widehat{p}_t^{R,\mathcal{X}} + \varepsilon_t^{\pi} \tag{9}$$

Real marginal cost combines domestic and imported input costs:

$$RMC_t^{\mathcal{X}} = b_1 \widehat{x}_t + b_2 \left(\widehat{p}_t^{R,M^{NOIL}} - \widehat{p}_t^{R,\mathcal{X}}\right) + b_3 \left(\widehat{p}_t^{R,M^{OIL}} - \widehat{p}_t^{R,\mathcal{X}}\right) \tag{10}$$

The second and third terms are what a conventional model would capture with a single real-exchange-rate term. Splitting them means an oil shock and a general import-price shock have different inflationary consequences — which is exactly what a fuel-importing economy needs from a model.

3.4 Demand components in detail

3.4.1 Private consumption

$$\begin{aligned} \widehat{c}_t^{R} &= c_1 \widehat{c}_{t-1}^{R} + c_2 \mathbb{E}_t \widehat{c}_{t+1}^{R} - c_3 \widehat{r}_t^{R} + c_4 \widehat{y}_t^{R} \\ &\quad - c_5 \left(\tau_t^{C} - \tau_{t+1}^{C}\right) + c_6 \widehat{REM}_t^{Y} + c_7 \widehat{GE}_t^{Tr,Y} + \varepsilon_t^{\widehat{c}} \end{aligned} \tag{11}$$

Five channels, each with a policy interpretation. Tighter credit conditions (a positive real rate gap $\widehat{r}_t^{R}$) reduce consumption — the monetary transmission channel, governed by c_3. Cyclical income enters through the output gap. An anticipated *rise* in the consumption tax ($\tau_t^{C} - \tau_{t+1}^{C} < 0$) front-loads spending. Government transfers above trend raise disposable income. And cyclical remittances enter directly.

The remittance channel matters in East Africa and barely exists further south: remittances run 3.2 % of GDP in the Ugandan steady state and 2.7 % in Kenya's, against 0.2 % in South Africa's. They act on consumption without passing through the output gap first, so a remittance shock is not simply a smaller version of an export shock — and in South Africa it is not a meaningful shock at all.

The trend counterpart determines the consumption share of GDP:

$$\begin{aligned} \Delta \overline{c}_t^{R} &= c_1^{\Delta \overline{c}} \Delta \overline{c}_{t-1}^{R} + (1 - c_1^{\Delta \overline{c}}) \left(\Delta \overline{p}_t^{R,Y} + \Delta \overline{y}_t^{R}\right) \\ &\quad - c_2^{\Delta \overline{c}} \Big(\overline{C}_t^{Y} - \big(C^{Y,SS} - c_3^{\Delta \overline{c}}(\tau_t^{C} - \mathcal{S}(\tau^{C})) + c_4^{\Delta \overline{c}}(\overline{REM}_t^{Y} - REM^{Y,SS}) \\ &\quad + c_5^{\Delta \overline{c}}(\overline{GE}_t^{Tr,Y} - \mathcal{S}(\overline{GE}^{Tr,Y}))\big)\Big) + \varepsilon_t^{\Delta \overline{c}} \end{aligned} \tag{12}$$

A permanently higher consumption tax lowers the steady-state consumption ratio; permanently higher remittances or transfers raise it. This is the mechanism by which fiscal policy has level effects and not merely cyclical ones.

3.4.2 Private investment

$$\begin{aligned}\widehat{\imath}_t^R &= c_1^{\widehat{\imath}}\widehat{\imath}_{t-1}^R + c_2^{\widehat{\imath}}\mathbb{E}_t\widehat{\imath}_{t+1}^R - c_3^{\widehat{\imath}}\widehat{r}_t^R + c_4^{\widehat{\imath}}\mathbb{E}_t\widehat{c}_{t+1}^R \\ &\quad + c_6^{\widehat{\imath}}\mathbb{E}_t\widehat{x}_{t+1}^{R,NNR} - c_5^{\widehat{\imath}}\left(\widehat{z}_t^{Mw} - \widehat{p}_t^{R,I}\right) + \varepsilon_t^{\widehat{\imath}}\end{aligned} \tag{13}$$

Investment rises when firms expect stronger domestic demand ($\mathbb{E}_t\widehat{c}_{t+1}^R$) or stronger export demand ($\mathbb{E}_t\widehat{x}_{t+1}^{R,NNR}$), and falls when the marginal cost of capital rises — either through a wider real interest gap or through a more expensive imported capital good, captured by the import-weighted real effective exchange rate gap adjusted for the relative price of investment.

That last term has a policy consequence worth stating: **a depreciation is contractionary for investment in this model even while it is expansionary for exports**. Which effect dominates is a quantitative question the estimated parameters answer, not a qualitative one the structure settles.

3.4.3 Exports and imports

Non-natural-resource exports respond to foreign demand and competitiveness:

$$\widehat{x}_t^{R,NNR} = c_1\widehat{x}_{t-1}^{R,NNR} + c_2\mathbb{E}_t\widehat{x}_{t+1}^{R,NNR} + c_3\left(\widehat{z}_t^{Xw} - \widehat{p}_t^{R,X^{NNR}}\right) + c_4\,\widehat{y}_t^{R,*} + \varepsilon_t \tag{14}$$

where $\widehat{y}_t^{R,*}$ is the export-weighted foreign output gap and $\widehat{z}_t^{Xw}$ the export-weighted real effective exchange rate gap: an undervalued currency stimulates export production.

Non-oil imports depend on every domestic demand component separately, which is the disaggregation earning its keep:

$$\begin{aligned}\widehat{m}_t^{R,NOIL} &= c_1\widehat{m}_{t-1}^{R,NOIL} + c_2\widehat{x}_t^{R,NNR} + c_3\widehat{\imath}_t^R + c_4\widehat{c}_t^R + c_7\widehat{g}_t^R \\ &\quad - c_5\widehat{p}_t^{R,M^{NOIL}} - c_6\left(\widehat{p}_t^{R,M^{NOIL}} - \widehat{p}_t^{R,I}\right) - c_8\left(\tau_t^{M^{NOIL}} - \tau_{t+1}^{M^{NOIL}}\right) + \varepsilon_t\end{aligned} \tag{15}$$

Because consumption, investment, exports and government absorption enter with separate coefficients, the import content of a fiscal expansion differs from that of an export boom. This is precisely the channel that makes Uganda's estimated fiscal multiplier small (Section 9.2).

3.5 The investment–growth nexus

This is FINEX's most consequential departure from conventional gap models: **potential output is not exogenous**. A Cobb–Douglas production function relates potential GDP to private capital K_t^R, public capital $K_t^{R,G}$ and total factor productivity, so potential growth is

$$\Delta\overline{y}_t^R = c_1\Delta k_t^R + c_2\Delta k_t^{R,G} + \Delta\overline{a}_t^R \tag{16}$$

Private capital accumulates in the standard way, with depreciation rate δ^{K^R}:

$$I_t^R = K_t^R - (1 - \delta^{K^R})K_{t-1}^R \tag{17}$$

Public capital accumulates similarly, but with a realistic implementation lag — public investment becomes fully productive only after three years:

$$\begin{aligned} K_t^{R,G} = (1 - \delta^{K^{R,G}})K_{t-1}^{R,G} + (1 - c_2 - c_3 - c_4)\, GE_t^{R,I^G} \\ + c_2 GE_{t-1}^{R,I^G} + c_3 GE_{t-2}^{R,I^G} + c_4 GE_{t-3}^{R,I^G} \end{aligned} \tag{18}$$

Total factor productivity follows an exogenous process subject to both temporary and persistent shocks.

3.5.1 Why this matters for policy

Equations (16)–(18) close a loop that is absent from most operational models, and it runs in both directions.

- **Interest rates affect the trend.** A higher real rate reduces private investment (equation (13)), which slows private capital accumulation, which lowers potential growth. Monetary policy is not neutral with respect to the path the economy converges to during the transition.
- **Public investment is a growth instrument with a lag.** Debt-financed public investment raises potential output — but only with a three-year delay, and only if it survives the fiscal rule's response to the resulting debt.
- **Debt feeds back onto growth.** Higher debt raises the UIP premium (Section 3.6.3), which raises real rates, which lowers investment and hence potential output. Whether debt-financed public investment is self-defeating depends on whether its productivity exceeds this drag — a quantitative question, and one whose answer depends on the initial debt level.
- **Crowding in and crowding out coexist.** Higher public investment raises demand directly, spurring private investment; but it also raises inflationary pressure, calling forth tighter policy that reduces private investment. Both are present in the model at once.

3.6 The external balance block

3.6.1 The balance of payments identity

The identity binds each period, in both levels and trends:

$$0 = CA_t^Y + FA_t^Y, \qquad 0 = \overline{CA}_t^Y + \overline{FA}_t^Y \tag{19}$$

The current account is

$$CA_t^Y = X_t^Y - \frac{M_t^{NOIL,Y}}{1+\tau_t^{MNOIL}/100} - \frac{M_t^{OIL,Y}}{1+\tau_t^{MOIL}/100} + REM_t^Y$$
$$- \left(G_t^{NR,Y} - GR_t^{NR,Y}\right) - GE_t^{BFCY,Y} + I_t^{NFA^O,Y} + CA_t^{O,Y} \tag{20}$$

comprising the trade balance measured before import taxes, remittances, natural resource sector profits paid abroad net of royalties, interest on foreign-currency public debt, interest income on private net foreign assets, and a residual.

The financial account distinguishes endogenous private flows from reserve management and official borrowing:

$$FA_t^Y = FA_t^{O,Y} - FXI_t^Y - FXA_t^Y + \left(GF_t^{FCY,Y} + \varepsilon_t^{BFCY,Y}\right) \tag{21}$$

where FXI_t^Y is discretionary foreign exchange intervention and FXA_t^Y planned reserve accumulation. Private flows give rise to a stock of net foreign assets, revalued by the exchange rate and deflated by nominal growth:

$$NFA_t^{O,Y} = -FA_t^{O,Y} + NFA_{t-1}^{O,Y}\frac{1+\Delta s_t^{US}/100}{1+\Delta y_t/100} \tag{22}$$

3.6.2 Uncovered interest parity and endogenous capital flows

The UIP premium is the interest differential adjusted for expected real depreciation, split into trend and gap:

$$\gamma_t = r_t^R - r_t^{R,US} - \Delta\bar{z}_{t+1}, \qquad \gamma_t = \bar{\gamma}_t + \hat{\gamma}_t \tag{23}$$

Crucially, FINEX does *not* assume perfect capital mobility. When cyclical pressures push private flows away from trend, a premium gap must open to clear the market:

$$FA_t^{O,Y} = \overline{FA}_t^{O,Y} + c_1^{FA^O}\left(1 - \tau^{FA^O,adm}\right)\left(\hat{\gamma}_t - \varepsilon_t^{\hat{\gamma}}\right) \tag{24}$$

The product $c_1^{FA^O}(1-\tau^{FA^O,adm})$ *is* the degree of capital mobility. The first factor reflects the depth of domestic financial markets; the second reflects administrative capital-flow management measures. When the product is small, flows are unresponsive to the premium, and any disturbance to the rest of the balance of payments must be absorbed by a large movement in the premium and hence the exchange rate. When it is large, flows do the adjusting and the premium barely moves. **The same shock therefore has very different consequences in a financially open and a financially closed economy**, and the difference is one parameter.

3.6.3 Risk appetite and sudden stops

The trend premium depends on the stocks investors care about — public debt, foreign reserves and private net foreign assets — through three deliberately *nonlinear* functions:

$$\gamma_t^{B} = \exp\Big(c_1^{\gamma^B}\left(B_t^Y - B^{Y,SS}\right)\Big) - 1 \tag{25}$$

$$\gamma_t^{FXR} = c_4^{\gamma^{FXR}}\left(\gamma^{FXR,max} + \min\Big(c_1^{\gamma^{FXR}} + c_3^{\gamma^{FXR}}\exp\Big(-c_2^{\gamma^{FXR}} FXR_t^Y\Big) - \gamma^{FXR,max},\, \epsilon\Big)\right) \tag{26}$$

$$\gamma_t^{NFA^O} = \exp\Big(-c_1^{\gamma^{NFA^O}}\left(NFA_t^{O,Y} - \mathcal{S}(NFA^{O,Y})\right)\Big) - 1 \tag{27}$$

The exponential form is the point. Around the steady state these terms are locally linear and unremarkable. Far from it — high debt, depleted reserves, a deeply negative net foreign asset position — they rise steeply, so that a small further deterioration produces a large increase in the required return. This is how FINEX represents a **sudden stop**: not as a regime switch or a discrete threshold, but as a continuous nonlinearity that is quiescent in normal times and violent when buffers are thin.

It also means **the same shock is more dangerous when the starting position is worse**, which linear models cannot express. A country at 45% debt and one at 85% debt do not face the same risk premium response to an identical fiscal slippage.

3.7 Macroeconomic policy

3.7.1 Monetary policy

The unconstrained policy rate follows a forward-looking Taylor rule with smoothing:

$$i_t^{unc} = \rho\, i_{t-1} + (1-\rho)\Big(\overline{rr}_t + \pi_t^C + \phi_\pi\left(\mathbb{E}_t\pi_{t+1}^C - \pi_t^*\right) + \phi_y\, \widehat{y}_t^R\Big) + \varepsilon_t^i \tag{28}$$

subject to an effective lower bound:

$$i_t = \max\left(i_t^{unc},\, i^{ELB}\right) \tag{29}$$

Equation (29) is the model's principal nonlinearity and has two consequences used later in this report: it is why the shock decomposition requires an exact rather than a linear method (Section 13), and why the additivity of policy shocks must be verified rather than assumed (Section 12.1).

Beyond the interest rate, FINEX carries two instruments that conventional models omit. **Foreign exchange intervention** FXI_t^Y enters the financial account directly (equation (21)) and, by changing reserves, shifts the risk-appetite term γ_t^{FXR}. **Capital flow management** operates either administratively through $\tau^{FA^O,adm}$ in equation (24), or as a price-based inflow tax τ^{FA^O} that raises investors' required return. Neither is used in this report's Uganda application — the model is run in its inflation-targeting float configuration — but both are available, and for a managed-float economy they would be central.

3.7.2 Fiscal policy

The structural deficit follows a rule that reverts to target, responds to the debt gap, and may lean against the cycle:

$$d_t^s = \delta_1 d_{t-1}^s + (1-\delta_1) d_t^{s,*} - \delta_2 \, \widetilde{b}_t + \delta_3 \, \widehat{y}_t^R + \varepsilon_t^d \tag{30}$$

where the debt deviation is itself forward-looking,

$$\widetilde{b}_t = (1-\delta_4)\,(b_t - b_t^*) + \delta_4 \, \mathbb{E}_t \widetilde{b}_{t+1} \tag{31}$$

so the fiscal authority reacts to where debt is going, not only where it is.

Debt is split by currency, and the foreign-currency component is revalued by the exchange rate:

$$b_t = b_t^{lcy} + b_t^{fcy} \tag{32}$$

This is where the model's nonlinearity meets its accounting. A depreciation raises b_t^{fcy} mechanically; the higher debt ratio raises the risk premium through equation (25); the higher premium raises interest rates and depreciates the currency further. Whether that loop is stabilising or explosive depends on the foreign-currency share, the level of debt and the capital-mobility parameter — which is why Uganda's largely concessional external borrowing matters so much to the results in Section 6.

3.8 Policy regimes

Three regimes ship with the model and are selectable at solve time:

- `IT_FLOAT` — inflation targeting with a floating exchange rate. The interest rate follows equation (28) and the exchange rate clears the balance of payments. *This is the configuration used throughout this report.*
- `IT_MANAGED` — inflation targeting with a managed float, in which FX intervention smooths exchange-rate movements.
- `FIX_CLOSED_PART` — a peg with partial capital mobility, in which the exchange rate follows a target path and the interest rate becomes endogenous to defending it.

The regime is not a switch on a single equation but a change in which variables are endogenous, which is why each is a separate model instance rather than a parameter setting.

4. Data

4.1 The observable panel

74 annual series are built for Uganda, of which 70 carry the name of an actual FINEX variable and enter the estimation.

Table 1: Data sources by block

Block	Primary source
National accounts	UBOS via World Bank WDI; UBOS quarterly GDP by expenditure
Inflation	IMF CPI spliced onto WDI; Bank of Uganda composite CPI
Balance of payments	**Bank of Uganda, Analytical BPM6** (2001Q3–2024Q4)
Fiscal	**MoFPED central government operations** (monthly, 1997–2025)
Monetary	Bank of Uganda Depository Corporations Survey; IMF MFS
Debt	IMF WEO gross debt; World Bank IDS; World Bank QPSD
Foreign block	FRED — US and euro-area activity and prices, Brent, commodities

Notes: National sources are preferred where they are richer than the international aggregators, which is the case for the balance of payments and the fiscal accounts. Transforms follow the model's own definitions: $DL\_X = 100\Delta \log X$; L_X_GAP is the Hodrick–Prescott cycle of $100 \log X$ at $\lambda = 100$, the annual smoothing constant (Hodrick and Prescott, 1997); $X_RAT = 100X/$nominal GDP; rates are annual averages in percent.

Two mapping choices materially affect the results.

- **Exports** are split into natural-resource (gold and coffee — world-price driven) and everything else, and **imports** into oil and non-oil, because FINEX treats these with different price processes.
- **Private investment** is measured capital formation *less* the public capital budget, since the model carries public investment separately. Getting this wrong puts the steady-state net foreign asset position at -116% of GDP.

4.2 A known gap: the policy rate

The short-term interest rate ends in 2019M04. Uganda stopped reporting interest rates to the IMF's Monetary and Financial Statistics database, and the Bank of Uganda publishes the Central Bank Rate only through a dashboard with no machine-readable endpoint.

Because moments are computed pairwise-complete (Section 5.2), the rate series still contributes its variance, autocovariance and cross-covariances over the years it exists — but **the monetary block is estimated on 20 rather than 24 annual observations**, and the projection's monetary starting point is not anchored to observed 2020–2024 policy rates. This is a real weakness in the monetary results of Section 9.1, not a technicality.

5. Estimation by Simulated Method of Moments

5.1 Why SMM

FINEX has 436 variables and 83 structural shocks; Uganda's usable annual sample is 24 observations. That ratio rules out the obvious alternatives. **Maximum likelihood** requires a well-conditioned likelihood, which with 83 shocks and 24 observations is flat in most directions. **Bayesian estimation** would work but converts the identification problem into a prior-sensitivity problem: with this little data the posterior would largely reproduce the priors.

Simulated Method of Moments matches a chosen set of second moments and is explicit about which features of the data the estimate is asked to reproduce. That explicitness is the point — when the estimate fails, as it does here for the commodity block, SMM shows exactly which moments it failed on (Section 5.7) rather than burying the failure in a likelihood value.

The estimator matches second moments only. Means are handled separately by the steady-state calibration, so every series is demeaned before moments are formed: the steady state is identified by sample means and policy anchors, the dynamics by second moments.

5.2 The moment set

166 moment conditions on 2001–2024: the variance and first autocovariance of every observable, plus 24 selected cross-covariances.

The cross-covariances are chosen, not exhaustive. A full cross-moment matrix over 70 observables would be 2,415 moments on 24 observations — overwhelmingly redundant and numerically ill-conditioned. The 24 selected span the channels the model is meant to capture: the output gap against inflation (the Phillips curve), the real rate against consumption (the IS curve), the exchange rate against import prices (pass-through), the deficit against debt (the fiscal rule), and the current account against its components.

5.3 Concentrating out the shock variances

This step makes the problem tractable, and rests on an exact property of the model rather than an approximation.

Let $h_j(\ell)$ be the impulse response of the observable vector to shock j at horizon ℓ. For a linear model driven by mutually independent shocks with variances σ_j^2, the autocovariance of the observables at lag τ is

$$\Gamma(\tau) \;=\; \sum_j \sigma_j^2\, A_j(\tau), \qquad A_j(\tau) \;=\; \sum_\ell h_j(\ell)\, h_j(\ell+\tau)' \tag{33}$$

The model's autocovariances are therefore **exactly linear in the shock variances**, and the loadings $A_j(\tau)$ depend only on the behavioural coefficients. Split the parameter vector into 16 behavioural coefficients θ and 40 shock variances σ^2. For any trial θ, stack the loadings into $A(\theta)$ and the data moments into γ; the variances then solve

$$\min_{\sigma^2 \geq 0} \left\| W^{1/2} \big(A(\theta)\, \sigma^2 - \gamma \big) \right\|^2 \tag{34}$$

which is a **non-negative least squares** problem, solved *exactly* by the Lawson–Hanson active-set algorithm (Lawson and Hanson, 1974) rather than searched. The non-negativity constraint is not a convenience: a variance must be non-negative, and imposing it as a constraint rather than by reparameterisation is what allows an uninformative shock to sit cleanly at zero.

Only θ is searched numerically. Reducing a 56-dimensional search to a 16-dimensional one is what makes the estimation feasible: each outer evaluation costs one model solve plus one NNLS, and 272 evaluations sufficed.

5.4 The outer search and the stability restriction

The 16 behavioural coefficients are searched by Nelder–Mead (Nelder and Mead, 1965) over a logistic reparameterisation of their bounds, so the simplex works in an unconstrained space while every trial point remains admissible.

Three gap equations are written $x_t = c_1 x_{t-1} + c_2 \mathbb{E}_t x_{t+1} + \cdots$ — consumption, investment and non-resource exports. Since c_2 is calibrated and only c_1 estimated, the search must respect

$$c_1 + c_2 < 1 \tag{35}$$

At $c_1 + c_2 \geq 1$ the equation has a root on or outside the unit circle and the whole model is explosive. With $c_2 = 0.6$ for investment, c_1 is capped at 0.4, implemented as $\min(\bar{c}_1,\, 1 - c_2 - 0.05)$ before the search starts.

The SMM objective will not catch this. A near-explosive impulse response either still yields finite second moments over a 16-year truncation, or fails outright and is silently skipped. The first unrestricted Uganda run estimated $c_1 = 0.58$ for investment, scored a *better* objective, and produced a projection that diverged, reaching 252% GDP growth and -135% of GDP public debt by 2035. A good moment fit is not evidence of a stable model — the two must be checked separately.

The objective fell from 91.2 on the shipped generic calibration to 83.2 estimated, over 272 outer evaluations.

5.5 Standard errors and identification

Standard errors for the shock standard deviations come from a moving-block bootstrap of the data moments (Künsch, 1989): 400 replications, block length $T^{1/3}$, re-solving the NNLS at each replication. The block structure preserves the serial dependence the autocovariance moments are measuring — an i.i.d. bootstrap would destroy exactly the feature being estimated.

Identification is diagnosed by how much of each column of the weighted loading matrix A lies in the span of the others. Columns that are effectively collinear cannot be separately identified and are flagged *weak*: the data can say what their sum contributes but not how it divides between them. This diagnostic is what surfaces the central limitation of the exercise (Section 17).

5.6 Estimated parameters

The same 16 behavioural coefficients were searched for each country, from the same generic calibration and against the same moment set, so the three estimates are directly comparable and a difference between columns is a difference the data insisted on.

Table 2: Estimated behavioural parameters, all countries

Parameter	Meaning	Cal.	Uganda	South Africa	Kenya
`C1_L_CONS_GAP`	Consumption persistence	0.25	**0.45**	**0.46**	**0.51**
`C3_L_CONS_GAP`	IS real-rate elasticity	0.10	**0.58**	**0.02**	**0.59**
`C1_L_INV_GAP`	Investment persistence	0.15	0.15	0.16	0.16
`C1_L_EXP_NNR_GAP`	Export persistence	0.15	**0.74**	**0.09**	**0.01**
`C1_L_IMP_NOIL_GAP`	Import persistence	0.15	**0.00**	**0.06**	**0.77**
`C1_DL_P_CONS`	Inflation backward-lookingness	0.80	0.63	0.90	0.91
`C2_DL_P_CONS`	Import-price pass-through	0.30	0.32	**0.14**	**0.11**
`C4_DL_P_CONS`	Phillips slope on the output gap	0.30	**0.19**	**0.46**	**0.19**
`C1_RS`	Policy-rate smoothing	0.15	**0.73**	**0.47**	**0.68**
`C2_RS`	Taylor weight on inflation	1.00	**0.55**	0.79	0.95
`C3_RS`	Taylor weight on the output gap	0.20	**0.04**	**0.47**	**0.02**
`C1_OFA_RAT`	Private financial inflows	1.00	1.02	1.02	1.01
`C1_OCA_RAT`	Other current account persistence	0.20	**0.11**	**0.13**	**0.05**
`C1_REMIT_RAT_GAP`	Remittance persistence	0.50	0.50	0.47	0.51
`C1_DEF_S_RAT`	Structural-deficit persistence	0.30	**0.08**	**0.13**	0.25
`C1_GEXP_CONS_RAT`	Spending persistence	0.70	0.86	0.62	0.70

Notes: "Cal." is the shipped generic calibration, identical for all three countries — so a value in bold, marking a move of more than 30% from it, is a parameter the country's own data spoke about. Searched by Nelder–Mead over a logistic reparameterisation of the bounds, with the stability restriction of Section 5.4 imposed before the search.

Four contrasts stand out, and each matters for a later result.

- **Uganda and Kenya are far more interest-sensitive than the generic calibration assumes; South Africa is not.** The IS real-rate elasticity rises from 0.10 to 0.58 in Uganda and 0.59 in Kenya — roughly 5.8 and 5.9 times the calibrated value — but *falls* to 0.02 in South Africa. Two countries agreeing and the third moving sharply the other way is what leads Section 18 to read the South African value as an identification failure rather than as a finding about South African households.

- **Policy-rate smoothing separates the three central banks.** It is estimated at 0.73 for Uganda and 0.68 for Kenya against 0.47 for South Africa, all well above the calibrated 0.15. A high value means a rate change is expected to persist, so the *real* rate moves more than the nominal shock — which is why Uganda's policy rate barely moves across the projection while Kenya's unwinds a five-point gap.

- **Inflation persistence splits the sample the other way.** Uganda's backward-looking weight is 0.63, against 0.91 for Kenya and 0.90 for South Africa. Kenya sits with South Africa rather than with its neighbour: it combines South African price dynamics with a Ugandan interest channel, and that combination is what makes the third country worth estimating rather than redundant.

- **Three parameters are untouched in all three countries.** Private financial flows, remittance persistence and, nearly, government consumption stay at their calibrated values everywhere. That is informative in itself: on those channels the second moments carry no information in any of the three samples, and the reported figure is the calibration rather than a result.

Section 8 reports the full estimate — these coefficients together with the shock standard deviations — with the identification pattern set out shock by shock for each country.

5.7 Fit, and where it fails

The **cost-push shock is the largest domestic disturbance** and cuts growth 1.08 pp — comparable to the monetary shock — because it raises inflation, which triggers the Taylor rule, which raises the real rate, which hits the interest-sensitive consumption block. *The output cost of an inflation shock in this model is mostly the policy response to it.*

6. Forecasting

This section sets out how FINEX produces a forecast, what Uganda's baseline forecast to 2035 says block by block, how well the model tracks history, and what kinds of forecast the framework supports beyond the unconditional baseline.

6.1 What kind of forecast this is

Three quite different objects are all loosely called "forecasts", and the distinction governs how each should be read.

- An **unconditional model forecast** projects forward with no future shocks, from model-consistent initial conditions. It answers: *where does the estimated structure take the economy from its current position if nothing else happens?* This is the baseline in Section 6.4.

- A **conditional forecast** imposes a path on some variables and lets the model determine the rest, by exogenising those variables and endogenising corresponding shocks. It answers: *what is consistent with oil at $110 through 2028?* This is Section 6.6.

- A **judgemental forecast** adds off-model information — a known project coming on stream, an announced tax measure, a forecaster's view. FINEX supports this through the same conditioning machinery, but the Uganda baseline here contains *none* of it.

The baseline below is the first kind. It is not a prediction of what will happen to Uganda. Most obviously, it knows nothing about Ugandan oil production coming on stream, which is a first-order omission for the external and fiscal accounts in the late 2020s. It is the model's own statement about the dynamics implied by its estimated structure, and its value lies in being exactly that — a clean benchmark against which scenarios and judgement can be measured.

6.2 How the forecast is constructed

6.2.1 Step 1 — the steady state

The steady state is calibrated separately from the dynamics. Sample means over 2010–2024 supply the great ratios — consumption, investment, exports, imports, remittances, reserves and the foreign-currency share of debt — while policy anchors supply the rest: the Bank of Uganda's 5% inflation target and the authorities' 50%-of-GDP debt anchor.

Two calibration steps are less obvious and matter more than they look.

Effective tax rates are solved for, not assumed. The split of the tax take between income, consumption and trade taxes is a judgement; the total it must deliver is an observable. Rates are scaled iteratively until steady-state revenue reproduces Uganda's observed ratio. Hand-set rates left revenue at 12.3% of GDP against 15.0% observed, which squeezed government purchases three points below the data, which shrank $C + I + G$, which shrank the steady-state trade deficit and left the current account at -1.4% against a -5.4% mean. A 2.7-point error in one calibrated ratio propagated into a four-point error in another.

Public investment closes the government budget residually. Because public investment enters in logarithms, a non-positive value is inadmissible, so the solver searches downward over the government consumption ratio until the steady state solves with positive public investment, and reports the value it needed.

6.2.2 Step 2 — inverting the model over history

The model is inverted over 2006–2024. Each observable is exogenised at each date where data exist, and the structural shock in that observable's *own* equation is endogenised at the same date, so the counts balance period by period. Solving the system then recovers the sequence of structural shocks that reproduces Uganda's history exactly.

This serves two purposes. It produces the historical shocks that Section 13 decomposes. And — the operational point — it produces **model-consistent initial conditions** for 2025, including the debt stocks, the capital stocks and the net foreign asset position. Starting a forecast from raw data would be inconsistent with the model's own accounting; starting it from the inversion is not.

25 observables are anchored this way. Observables are added one at a time and any that makes the system singular is skipped and reported, so a single awkward series does not cost the rest. Partial coverage is allowed: a series is exogenised in the years it exists and left endogenous elsewhere, which

is what permits the policy rate — whose data end in 2019 — still to anchor the monetary block over 2006–2019.

6.2.3 Step 3 — projecting forward

From 2025 the model is simulated with all structural shocks set to zero. The simulation runs 60 years beyond the reporting horizon under natural terminal conditions, so the terminal condition does not contaminate the reported path. Each observable is reported as its steady-state level plus the simulated deviation.

6.3 In-sample performance

Exogenised series are reproduced exactly by construction, so they say nothing about fit. The informative test is how the model does on everything it was *not* told.

Over 2006–2019 the model tracks the output gap to a root-mean-square error of 3.0 percentage points, the current account to 2.6 points of GDP, and government revenue to 1.7 points of GDP — none of which were exogenised. These are genuine cross-equation predictions: the model is inferring them from the shocks recovered elsewhere plus its own structure.

Figure 1 shows the fitted path against the data. The visual impression matters as much as the statistics: the model reproduces the timing and sign of Uganda's major cyclical episodes — the 2011 inflation spike, the 2016–17 slowdown, the 2020 contraction and the subsequent recovery — while smoothing their amplitude, which is the expected behaviour of a model whose shocks are recovered from a subset of series.

Figure 1: Uganda: history, model fit and projection to 2035

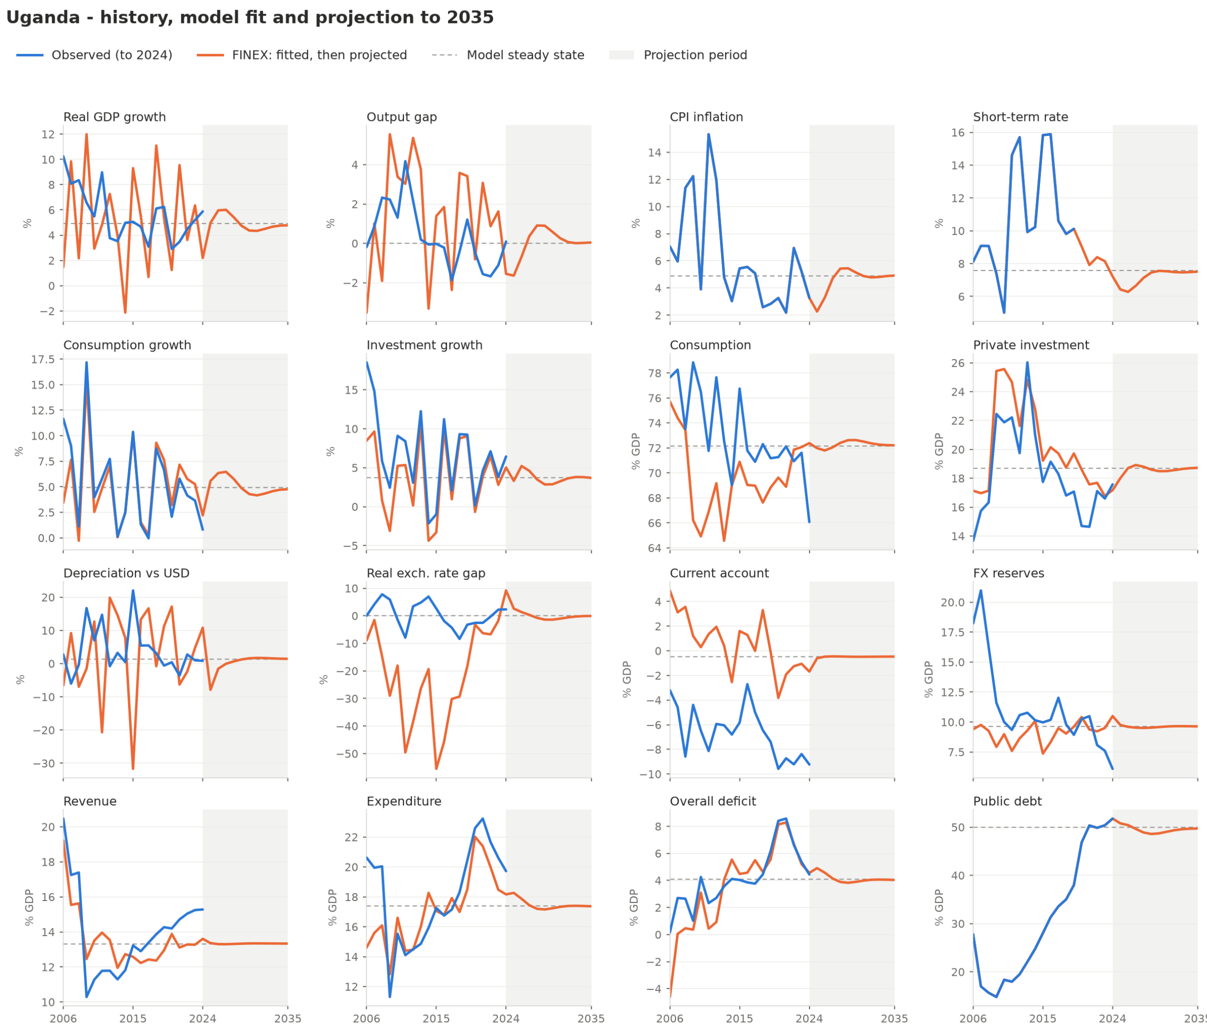

Notes: Blue is observed data through 2024; orange is the FINEX model — fitted over history, then projected. The dashed horizontal line in each panel is that variable's model steady state; the shaded band marks the projection period (2025–2035). Growth, inflation and the interest rate in percent; fiscal and external variables in percent of GDP. The short-term interest rate has no observations after 2019 (Section 4.2).

Figure 2: South Africa: history, model fit and projection to 2035

Notes: As Figure 1, for South Africa. Note the different vertical scales: South African growth runs near 1.3% against Uganda's 4.9%, and the debt ratio near 76% against 50%. The 2020 contraction is far deeper here, and investment falls roughly 30% in that year.

Figure 3: Kenya: history, model fit and projection to 2035

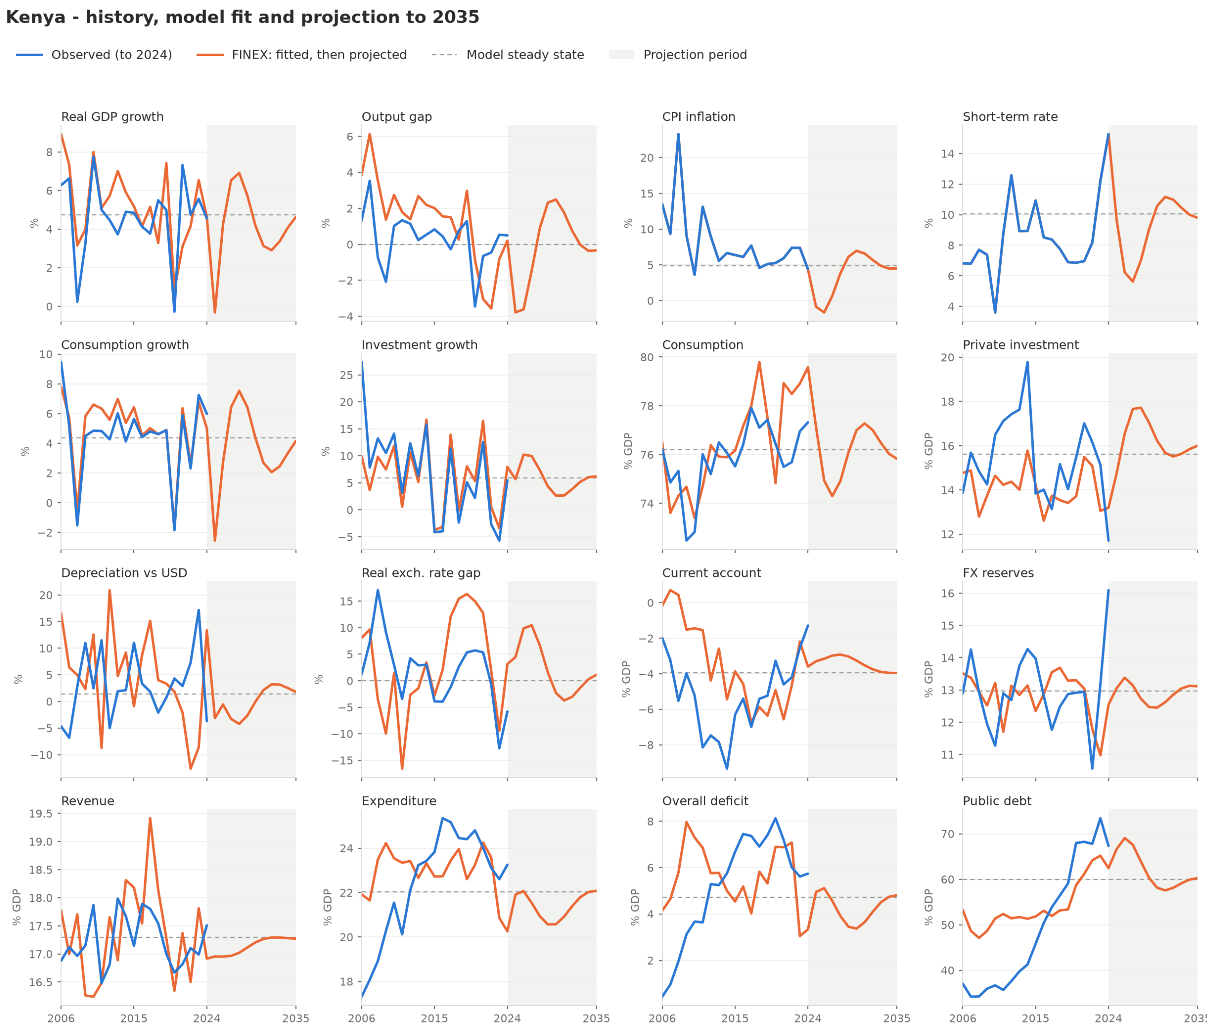

Notes: As Figure 1, for Kenya. The 2024 policy rate of 15.3% against a 10.1% steady state is the largest tightening in the three samples, and the projection is dominated by its unwinding. Kenya's estimate rests on a substantially weaker moment fit than the other two (Section 18.3).

6.4 The baseline in detail

The four blocks of the baseline — activity, prices and monetary conditions, the external accounts, and the fiscal accounts and debt — are reported in Section 7, with all three countries side by side in each table so that a comparison can be read across a row rather than held in the head across pages. The convergence diagnostics that qualify those projections follow immediately below.

6.5 Convergence diagnostics

A projection is only a forecast if it converges. FINEX contains lead terms and accumulating stocks, and a mis-specified or mis-estimated parameter can produce a path that drifts without ever obviously breaking.

Two checks are applied to every run in this project:

- **Terminal convergence:** every observable must lie within 25 units of its own steady state by 2035. Uganda's largest terminal gap is **1.28**, comfortably inside.

- **Path boundedness:** no observable may deviate from its steady state by more than 500 units anywhere in the projection.

These thresholds are deliberately loose — they are designed to catch divergence, not to certify accuracy. Their value is demonstrated by the cases they caught. An unrestricted estimation of the investment persistence parameter (Section 5.4) produced a projection reaching 252% GDP growth and -135% of GDP public debt by 2035, while scoring a *better* moment fit than the restricted estimate. Convergence and fit are independent properties and must be checked independently.

6.6 Conditional forecasting

The same inversion machinery that produces initial conditions can be pointed at the *future*, which turns FINEX into a conditional forecasting tool. This is how the model is used in practice, and it is the most operationally valuable capability in the framework.

To impose a path, the conditioned variable is exogenised over the chosen future dates and an instrument shock is endogenised at the same dates. The solver then finds the shock sequence that delivers the requested path, and reports what the rest of the economy must look like for that path to be coherent.

The counting must balance. One endogenised shock-date is required per exogenised variable-date. Fewer, and the system is over-determined and will not solve; more, and it is under-determined and the solution is not unique. This is the single easiest way to obtain silent nonsense from a conditional forecast, and it is checked before submission rather than diagnosed afterwards.

The choice of instrument is an economic statement, not a technicality. Conditioning inflation on a path via the price shock says "suppose cost pressures evolve so that inflation follows this path". Conditioning the same path via the policy-rate shock says "suppose the central bank acts so that inflation follows this path". The two produce different paths for everything else. The instrument defaults to the shock in the conditioned variable's own equation, which is right most of the time, but it is a default and not a law.

Typical uses in a Uganda context:

- **Imposing a world price path** — oil at a given level through 2028 — and reading off the implied inflation, fiscal and external consequences.

- **Imposing near-term official projections** for revenue or growth and letting the model complete the rest of the accounts consistently.

- **Imposing an announced policy path** — a stated disinflation trajectory — and recovering the interest-rate path that delivers it.

- **Nowcasting**: conditioning on partial-year data for the series that are already observed, and letting the model fill in the rest.

Infeasible conditioning is reported as non-convergence rather than as an approximate answer. If the requested path is not attainable, the honest output is a failure, not a near-miss silently presented as a result.

6.7 What the baseline omits

Four things are absent from the baseline by construction, and a reader comparing it against an institutional forecast should account for all four.

- **Oil production.** Uganda's oil is not in this forecast in any form. For the external accounts and the fiscal accounts from the late 2020s this is the largest single omission, and it runs in the direction of making the baseline too pessimistic on exports and revenue.
- **Future shocks.** The projection sets every structural shock to zero from 2025. Uganda's history contains large commodity and weather disturbances roughly every three to four years; a path with none is not a central expectation so much as a reference.
- **Announced policy.** No tax measure, spending plan or monetary decision after 2024 enters the baseline except through the estimated rules.
- **Uncertainty bands.** The baseline is a point path. Fan charts would require the estimated shock covariance, and roughly half those variances sit at the zero bound (Section 17), so bands built from them would understate risk — particularly foreign-driven risk — in a way that would mislead rather than inform. They are therefore not reported.

The remedy for the first three is conditional forecasting (Section 6.6); the remedy for the fourth is estimating the foreign block off foreign data, as Section 17 sets out.

7. Results

Section 6 set out how the forecast is constructed. This section reports the baseline projection for all three countries, block by block, with the countries side by side in each table so that the comparison can be read across a row rather than held in the head across pages.

All three are unconditional projections with no future shocks, from model-consistent initial conditions, converging on the steady states calibrated in Section 6.2. None contains off-model information.

7.1 Activity

Table 3: Baseline projection: activity

		Uganda				South Africa				Kenya			
		SS	2027	2030	2035	SS	2027	2030	2035	SS	2027	2030	2035
Real GDP growth	%	4.9	6.0	4.4	4.8	1.3	1.3	1.2	1.2	4.7	6.5	4.2	4.6
Output gap	%	0.0	0.4	0.6	0.0	0.0	-0.6	-0.4	-0.2	0.0	-1.5	2.5	-0.3
Consumption growth	%	4.9	6.5	4.3	4.8	1.5	1.4	1.3	1.4	4.4	6.4	4.4	4.2
Investment growth	%	3.7	4.6	2.9	3.7	1.0	1.9	1.3	1.0	5.9	10.0	2.6	6.2
Private consumption	% GDP	72.2	72.0	72.6	72.2	63.5	63.6	63.6	63.5	76.2	74.3	77.0	75.8
Private investment	% GDP	18.7	18.9	18.5	18.7	14.8	14.4	14.7	14.9	15.6	17.7	16.2	16.0

Notes: SS is each country's model steady state; the year columns are the model's own path with no future shocks. Growth and interest rates in percent per year, gaps in percent of the corresponding trend, ratios in percent of nominal GDP. Steady states are calibrated separately for each country (Section 6.2), so a difference in the SS column is a difference in calibration as well as in the data.

The growth ordering is the clearest single fact in the table. Uganda and Kenya converge on 4.8% and 4.6%; South Africa on 1.2%. That gap is not a cyclical difference the projection works off — it is the difference in calibrated potential, and it drives almost everything in the fiscal block below.

Uganda's growth hump — 6.1% in 2027 falling to 4.3% by 2030 — is the output gap closing from below and overshooting, not a growth acceleration. **South Africa's output gap never closes** within the horizon, reaching only -0.3 by 2035, while Uganda's and Kenya's do.

The 2027 peak is a cycle, not a growth rate. Reading it as a forecast of medium-term potential would overstate Uganda by 1.1 points and Kenya by 1.8 (6.0% against a 4.9% steady state, 6.5% against 4.7%). The model's view of potential is the SS column, and in both cases the peak is the output gap closing from below and overshooting. South Africa is the instructive exception: at 1.3% against a 1.3% steady state it has no hump to misread, because its gap never closes within the horizon.

Kenya's private investment at 15.6% of GDP in steady state sits between the others, but its 2024 position of 11.7% is the furthest below trend of the three, so a substantial part of Kenya's projected recovery is investment returning to where the model thinks it belongs.

7.2 Prices and monetary conditions

Table 4: Baseline projection: prices and monetary conditions

		Uganda				South Africa				Kenya			
		SS	2027	2030	2035	SS	2027	2030	2035	SS	2027	2030	2035
CPI inflation	%	4.9	4.7	5.1	4.9	4.4	3.9	4.5	4.5	4.9	0.6	7.0	4.5
GDP deflator inflation	%	4.9	4.8	5.0	4.9	4.6	4.2	4.7	4.7	4.5	1.4	5.9	4.3
Short-term rate	%	7.6	6.7	7.6	7.5	5.2	4.3	4.9	5.1	10.1	5.6	10.6	9.8
Real rate gap	pp	0.0	-0.7	-0.3	-0.1	0.0	-0.2	-0.3	-0.1	0.0	0.2	-1.6	0.3
Depreciation vs USD	%	1.4	-0.0	1.6	1.4	0.9	1.4	1.1	1.1	1.4	-3.3	-0.0	1.8
Real exchange-rate gap	%	0.0	0.3	-1.4	-0.1	0.0	-1.9	-1.2	-0.2	0.0	10.5	-2.3	1.1

Notes: SS is each country's model steady state; the year columns are the model's own path with no future shocks. Growth and interest rates in percent per year, gaps in percent of the corresponding trend, ratios in percent of nominal GDP. Steady states are calibrated separately for each country (Section 6.2), so a difference in the SS column is a difference in calibration as well as in the data.

Inflation converges near target in all three — 4.9, 4.5 and 4.7 — which is unsurprising given that each steady state is calibrated to the country's own anchor. The informative content is the transition.

Kenya's is the largest monetary adjustment in the three samples. Its policy rate starts 2024 at 15.3% against a 10.1% steady state and falls back through the projection. South Africa eases from 8.3% toward 5.2%. Uganda's rate barely moves, which is what a smoothing coefficient of 0.73 produces.

The real rate gap tells the story the nominal rate hides: Uganda's is *positive* at 1.7 in 2025 despite below-target inflation, because the nominal rate is slow to follow inflation down. A heavily smoothed rule delivers a real tightening it never intended.

7.3 External accounts

Table 5: Baseline projection: external accounts

		Uganda				South Africa				Kenya			
		SS	2027	2030	2035	SS	2027	2030	2035	SS	2027	2030	2035
Current account	% GDP	-0.5	-0.5	-0.5	-0.5	-0.6	-0.0	-0.2	-0.4	-3.9	-3.0	-3.3	-4.0
Non-resource exports	% GDP	13.2	13.4	13.2	13.3	16.1	17.1	16.8	16.4	9.8	10.5	9.3	9.7
Resource exports	% GDP	3.4	3.4	3.4	3.4	12.8	13.2	13.0	12.9	5.0	5.1	4.6	5.1
Non-oil imports	% GDP	16.7	17.0	16.8	16.8	14.8	15.4	15.2	15.0	15.3	15.0	14.4	15.4
Oil imports	% GDP	2.9	2.9	2.9	2.9	5.5	5.6	5.5	5.5	4.0	4.0	3.8	4.0
Remittances	% GDP	3.2	3.3	3.2	3.2	0.2	0.3	0.3	0.2	2.7	2.8	2.7	2.7
Gross FX reserves	% GDP	9.6	9.5	9.6	9.6	13.6	13.8	13.7	13.6	13.0	13.1	12.5	13.1

Notes: SS is each country's model steady state; the year columns are the model's own path with no future shocks. Growth and interest rates in percent per year, gaps in percent of the corresponding trend, ratios in percent of nominal GDP. Steady states are calibrated separately for each country (Section 6.2), so a difference in the SS column is a difference in calibration as well as in the data.

The external blocks converge quickly and then barely move, which is the balance-of-payments identity combined with calibrated steady-state ratios rather than a finding about any of the three economies.

Two structural differences are real. **South Africa's current account is near balance** at -0.6% of GDP in steady state where Uganda and Kenya run -2.8 and -2.7. And **remittances matter for Uganda in a way they do not for South Africa** — 3.2% of GDP against a much smaller share — which is why the remittance channel of equation (11) carries first-order weight in the Ugandan results and not elsewhere.

Uganda's and Kenya's projections carry a level step between the 2024 data and the 2025 projection, discussed in Section 17.2. It is a data reconciliation problem, not an economic adjustment.

7.4 Fiscal accounts and debt

Table 6: Baseline projection: fiscal accounts and debt

		Uganda				South Africa				Kenya			
		SS	2027	2030	2035	SS	2027	2030	2035	SS	2027	2030	2035
Revenue	% GDP	13.3	13.3	13.3	13.3	25.9	25.8	25.9	25.9	17.3	17.0	17.2	17.3
Expenditure	% GDP	17.4	17.4	17.2	17.4	29.8	29.6	29.6	29.7	22.0	21.5	20.6	22.1
Government purchases	% GDP	12.1	12.2	12.0	12.1	13.0	12.5	12.6	12.7	12.7	11.5	11.1	12.8
Interest bill	% GDP	2.6	2.7	2.6	2.6	6.5	6.7	6.7	6.6	5.6	6.2	5.7	5.5
Overall deficit	% GDP	4.1	4.1	3.9	4.0	3.9	3.8	3.7	3.8	4.7	4.6	3.4	4.8
Structural deficit	% GDP	4.1	4.2	4.0	4.1	3.9	3.4	3.5	3.7	4.7	3.8	4.0	4.7
Public debt	% GDP	50.0	49.6	48.7	49.7	75.0	77.8	77.1	76.1	60.0	67.6	58.1	60.3
External public debt	% GDP	30.6	27.5	28.9	30.2	27.9	27.2	27.9	28.2	27.8	30.5	26.4	28.1

Notes: SS is each country's model steady state; the year columns are the model's own path with no future shocks. Growth and interest rates in percent per year, gaps in percent of the corresponding trend, ratios in percent of nominal GDP. Steady states are calibrated separately for each country (Section 6.2), so a difference in the SS column is a difference in calibration as well as in the data.

This is where the three economies differ most, and where the interest bill does the work.

South Africa raises and spends roughly twice what Uganda does — 25.9% and 29.8% of GDP against 13.3% and 17.4% — yet runs a similar overall deficit. The difference is absorbed by debt service: **6.5% of GDP against 2.9%**. Kenya sits between on revenue and spending but close to South Africa on interest, at 5.6% on a debt ratio 15 points lower, because Kenya borrows substantially at market rates where Uganda borrows concessionally.

The debt paths converge on their anchors from different directions and at different speeds. Uganda's falls from 51.8% to 48.8% by 2030 and settles at 50%. Kenya's closes a seven-point gap from 67.3% toward its 60% anchor — the largest adjustment of the three, on the second-largest interest bill. South Africa's *rises* to 77.7% by 2027 before edging back toward 75%, because its growth–interest differential is the least favourable.

The interest burden, not the growth rate, is what separates Uganda from Kenya. The two grow at nearly the same pace and face very different fiscal arithmetic.

8. Estimates

This section reports the full estimate for all three countries: the 16 behavioural coefficients searched by Nelder–Mead, and the shock standard deviations the non-negative least squares identified. Section 18 interprets the contrasts; here they are laid out.

8.1 Behavioural parameters

Table 7: Estimated behavioural parameters, all countries

Parameter	Meaning	Cal.	Uganda	South Africa	Kenya
C1_L_CONS_GAP	Consumption persistence	0.25	**0.45**	**0.46**	**0.51**
C3_L_CONS_GAP	IS real-rate elasticity	0.10	**0.58**	**0.02**	**0.59**
C1_L_INV_GAP	Investment persistence	0.15	0.15	0.16	0.16
C1_L_EXP_NNR_GAP	Export persistence	0.15	**0.74**	**0.09**	**0.01**
C1_L_IMP_NOIL_GAP	Import persistence	0.15	**0.00**	**0.06**	**0.77**
C1_DL_P_CONS	Inflation backward-lookingness	0.80	0.63	0.90	0.91
C2_DL_P_CONS	Import-price pass-through	0.30	0.32	**0.14**	**0.11**
C4_DL_P_CONS	Phillips slope on the output gap	0.30	**0.19**	**0.46**	**0.19**
C1_RS	Policy-rate smoothing	0.15	**0.73**	**0.47**	**0.68**
C2_RS	Taylor weight on inflation	1.00	**0.55**	0.79	0.95
C3_RS	Taylor weight on the output gap	0.20	**0.04**	**0.47**	**0.02**
C1_OFA_RAT	Private financial inflows	1.00	1.02	1.02	1.01
C1_OCA_RAT	Other current account persistence	0.20	**0.11**	**0.13**	**0.05**
C1_REMIT_RAT_GAP	Remittance persistence	0.50	0.50	0.47	0.51
C1_DEF_S_RAT	Structural-deficit persistence	0.30	**0.08**	**0.13**	0.25
C1_GEXP_CONS_RAT	Spending persistence	0.70	0.86	0.62	0.70

Notes: "Cal." is the shipped generic calibration, identical for all three countries — so a value in bold, marking a move of more than 30% from it, is a parameter the country's own data spoke about. Searched by Nelder–Mead over a logistic reparameterisation of the bounds, with the stability restriction of Section 5.4 imposed before the search.

The calibration column is the same for all three countries, which makes the table readable as a statement about what each country's data had to say. **Bold marks a move of more than 30% from the generic calibration — a parameter the data spoke about.**

Uganda's data move the estimate furthest: the IS real-rate elasticity rises nearly six-fold, policy-rate smoothing more than four-fold, import persistence collapses to near zero. South Africa's estimate is the most orthodox, with the sharp exception of the IS elasticity, which *falls* to 0.02 — the only key parameter in this report that moves toward zero, and the one Section 18 argues is an identification failure rather than a finding. Kenya sits between the two on most parameters and close to Uganda on the one that matters most for that argument.

Three parameters are effectively untouched in all three countries — C1_OFA_RAT, C1_REMIT_RAT_GAP and, nearly, C1_GEXP_CONS_RAT. That is informative in its own right: on those channels the second moments carry no information in any of the three samples, and the reported values are the calibration rather than a result.

8.2 Shock standard deviations

Table 8: Shock standard deviations, all countries

Shock	Uganda	South Africa	Kenya
RES_ABS_EZ_DL_P_CONS	1.57	1.58	1.57
RES_ABS_EZ_L_GDP_GAP	1.42	1.37	1.44
RES_ABS_EZ_L_RER_GAP	5.13	4.93	5.11
RES_ABS_L_RP_OIL_GAP	–	10.9	–
RES_ABS_US_DL_P_CONS	1.39	1.36	0.838
RES_ABS_US_L_GDP_GAP	1.36	1.44	1.69
RES_ABS_US_RS	1.41	1.12	1.26
RES_DEF_S_RAT	0.805	–	–
RES_DL_P_CONS	4.36	0.992	–
RES_DL_P_EXP_NR	13.8	–	8.79
RES_DL_P_GOV	4.22	–	2.97
RES_DL_P_IMP_NOIL	2.94	–	2.94
RES_DL_P_IMP_OIL	–	4.53	–
RES_DL_P_INV	–	2.3	5.38
RES_FIN_FCY_RAT	1.4	–	0.571
RES_GEXP_CONS_RAT	0.214	0.0337	0.157
RES_GREV_OTH_RAT	0.513	–	–
RES_L_CONS_GAP	–	1.08	–
RES_L_EXP_NNR_GAP	–	–	1.22
RES_L_EXP_NR_GAP	15.2	–	8.83
RES_L_IMP_NOIL_GAP	3.56	–	–
RES_L_IMP_OIL_GAP	7.18	8.68	–
RES_L_INV_GAP	1.67	–	3.13
RES_RR_PREM_GOV_FCY	–	0.951	2.13
RES_RR_PREM_GOV_LCY	2.72	–	–
RES_RS	–	1.37	–

Notes: A dash marks a variance the non-negative least squares drove to the zero bound; that shock reverts to its calibrated value, so any response scaled by it has an estimated shape and an assumed magnitude. Shocks at the bound in all three countries are omitted from this table and listed in Appendix I. Identified out of 40: Uganda 19, South Africa 15, Kenya 16.

The dashes are the substance of this table. A dash marks a variance driven to the zero bound, meaning the estimator could not separate that shock's contribution from the others and the shock reverts to its calibrated value.

Identified out of 40: Uganda 21, Kenya 16, South Africa 15. The pattern of which shocks fail differs by country in ways that matter for reading the impulse responses — Uganda's monetary shock scale is calibrated while South Africa's is estimated, and the world oil price is identified in South Africa but not in Uganda or Kenya *despite being the same underlying series.* That last observation is the identification problem of Section 17 in its clearest form: an estimator that identifies a common world price in one country's sample and not in another's is reporting something about the samples, not about oil.

The largest identified shocks are the commodity-trade disturbances in Uganda and Kenya, and the domestic demand and price shocks in South Africa — which is what one would expect of two commodity-exporting economies and one large diversified one.

9. Policy simulation: monetary and fiscal

All impulse responses are to a **one-standard-deviation** shock — the SMM-estimated standard deviation where the data identify it, the calibrated one otherwise, labelled on each row of the figures. Because shocks are scaled by their own dispersion, responses are comparable *within* a country: they show what a typical shock of each kind does there.

They are not directly comparable across countries, because the two countries' standard deviations differ. Where a cross-country statement is made below it is made per unit of standard deviation, and says so.

Figure 4: Uganda: responses to policy and demand shocks

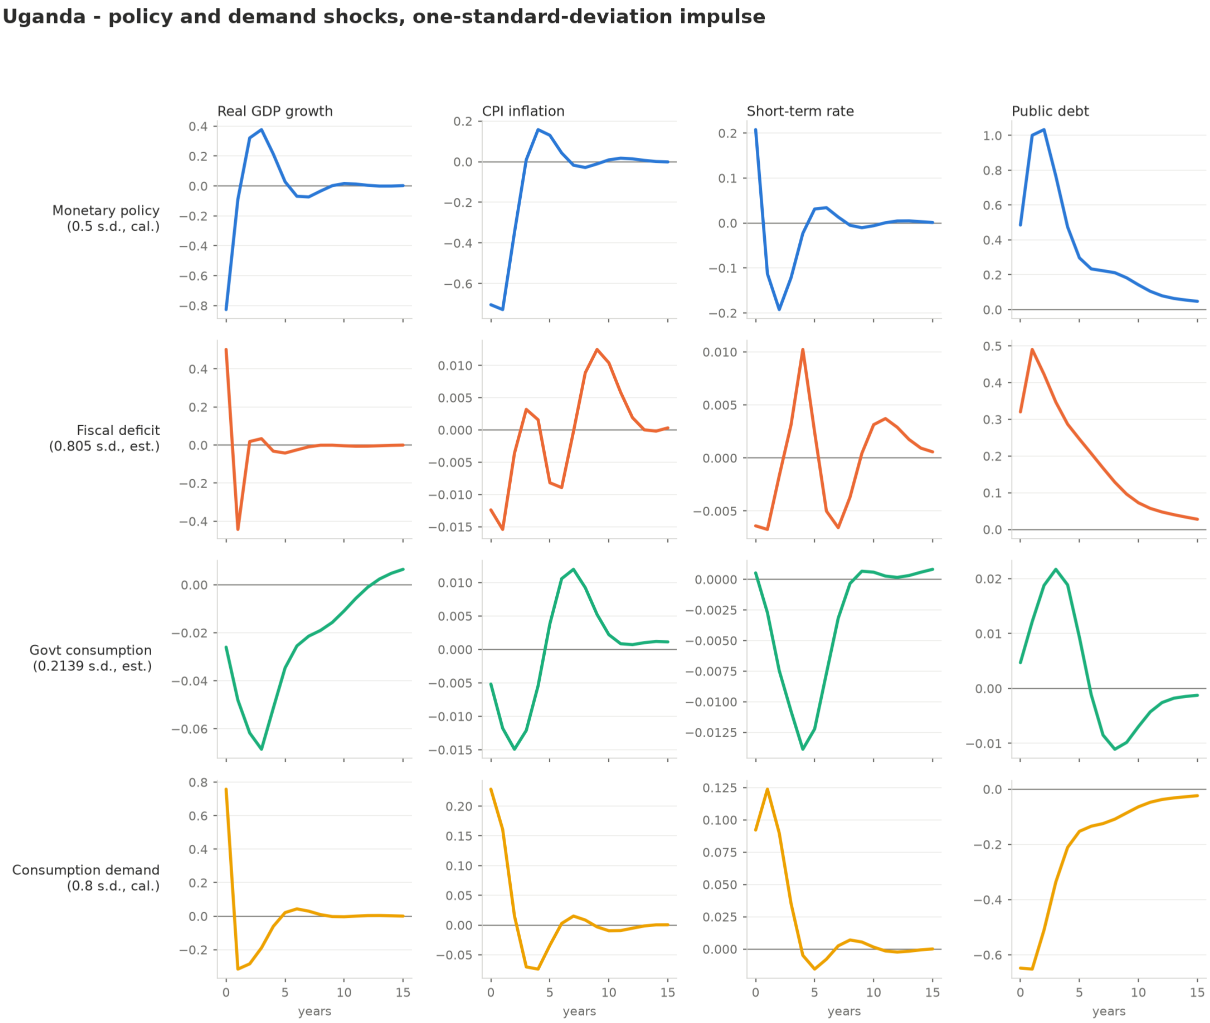

Notes: Rows are shocks, columns are observables; the horizontal axis is years after impact. Row labels give each shock's standard deviation and whether it was estimated (*est.*) or retained at its calibrated value (*cal.*). Growth, inflation and the rate in percentage points; debt in percent of GDP.

Figure 5: South Africa: responses to policy and demand shocks

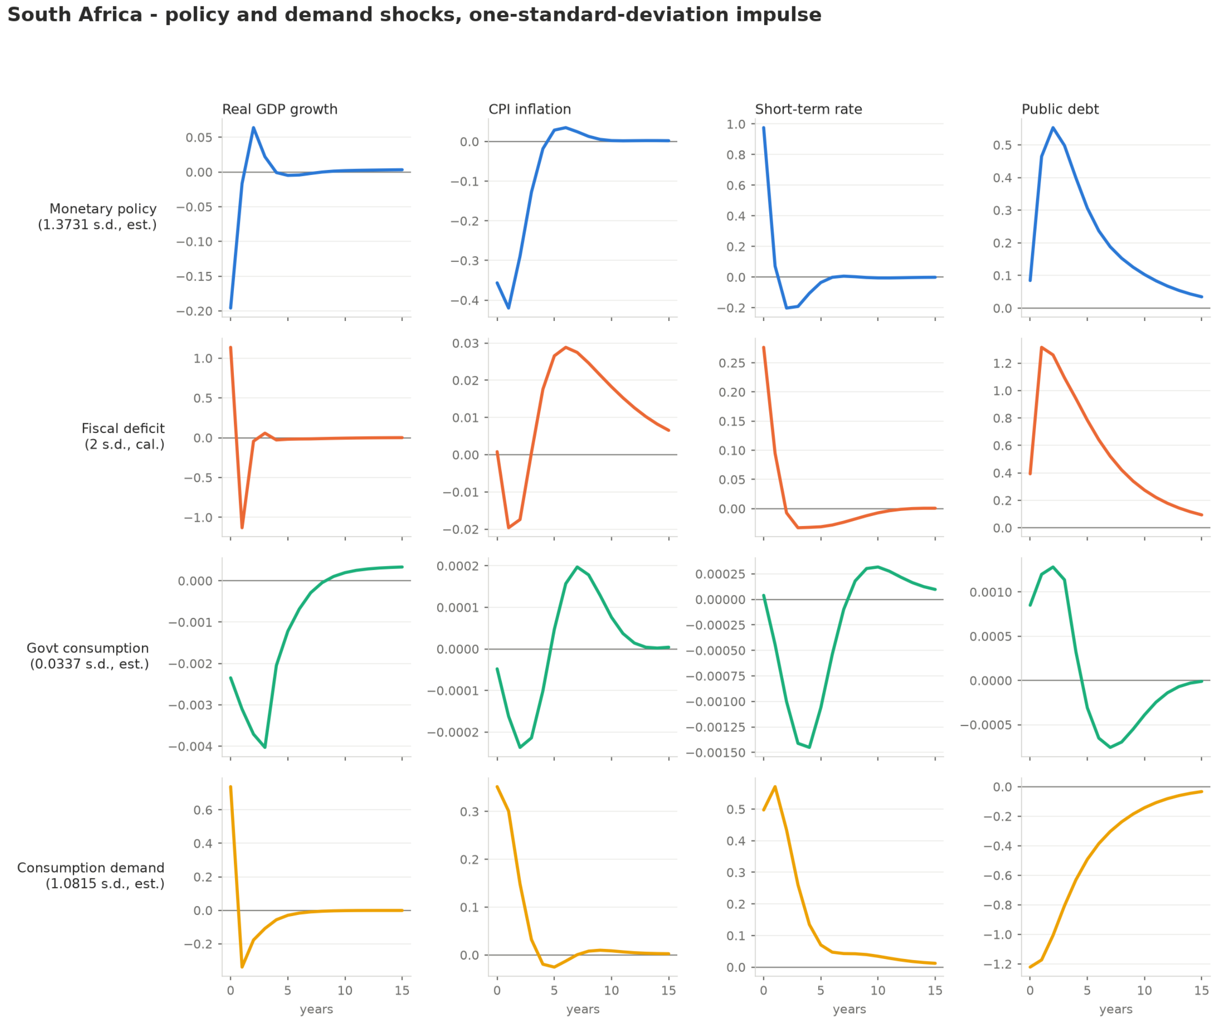

Notes: As Figure 4, for South Africa. Note the different standard deviations in the row labels: South Africa's monetary shock is 1.36 against Uganda's 0.50, so the two panels are not on a common scale.

Figure 6: Kenya: responses to policy and demand shocks

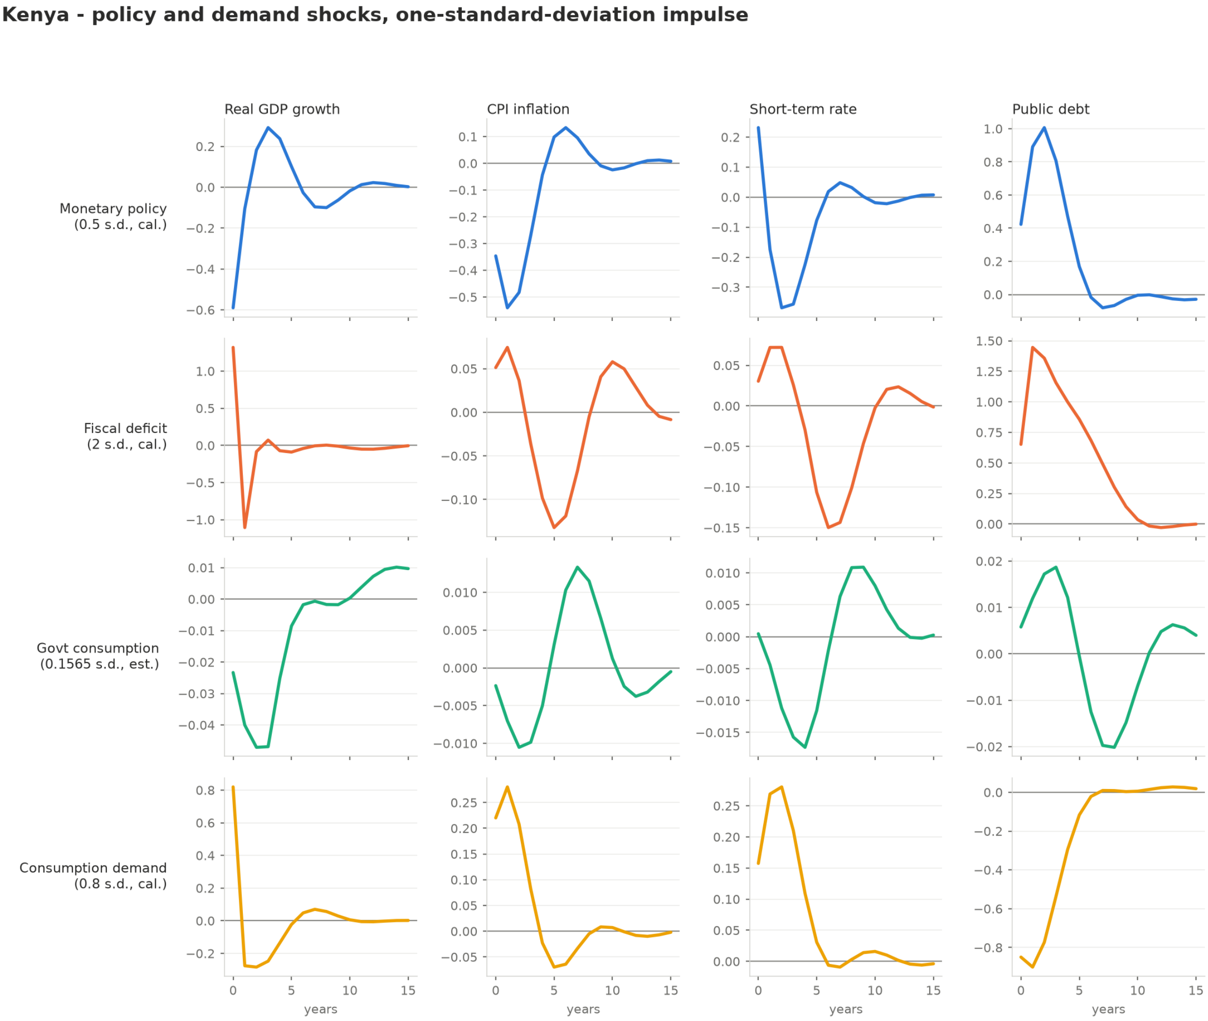

Notes: As Figure 4, for Kenya. Kenya's monetary shock scale is calibrated, as Uganda's is, so the two are on a common footing with each other but not with South Africa's estimated scale.

Table 9: Peak responses to policy and demand shocks

	Uganda				Kenya				South Africa			
Shock	σ	src	GDP gr.	Debt	σ	src	GDP gr.	Debt	σ	src	GDP gr.	Debt
Monetary policy	0.50	cal.	-0.83	$+1.03$	0.50	cal.	-0.59	$+1.01$	1.37	est.	-0.20	$+0.55$
Fiscal (struct. deficit)	0.81	est.	$+0.50$	$+0.49$	2.00	cal.	$+1.32$	$+1.45$	2.00	cal.	$+1.14$	$+1.31$
Government consumption	0.21	est.	-0.07	$+0.02$	0.16	est.	-0.05	-0.02	0.03	est.	-0.00	$+0.00$
Consumption demand	0.80	cal.	$+0.76$	-0.65	0.80	cal.	$+0.82$	-0.90	1.08	est.	$+0.74$	-1.22
Cost-push	4.36	est.	-1.03	-1.50	1.00	cal.	$+0.24$	-0.53	0.99	est.	-0.04	-0.40

Notes: Peak absolute response over a 15-year horizon, in percentage points for growth and inflation and percent of GDP for debt. "src" records whether σ was estimated (*est.*) or reverted to its calibrated value (*cal.*) after reaching the zero bound. **The sources differ by country in ways that matter**: a response whose scale is calibrated has an estimated shape and an assumed magnitude. Because σ differs by country the columns are comparable within a country but not across one; per-unit comparisons in the text divide by σ.

9.1 Monetary policy

Uganda. A one-standard-deviation tightening — 0.50 percentage points, **calibrated not estimated** — cuts real GDP growth by 0.83 percentage points and inflation by 0.73. Public debt *rises* 1.03 points of GDP: the tightening is fiscally costly through lower nominal GDP in the denominator and weaker revenue.

Kenya. The same 0.50-point tightening, also on a calibrated scale, cuts growth 0.59 and inflation 0.54, with debt rising 1.01 points. Kenya sits between the other two in level but close to Uganda in kind, which follows from an IS elasticity the estimation put at almost exactly Uganda's value.

South Africa. A one-standard-deviation tightening — 1.37 percentage points, larger than either — cuts growth by only 0.20 points and inflation by 0.42.

Per unit of standard deviation the ordering is stark: -1.65 for Uganda and -1.18 for Kenya against -0.14 for South Africa — a factor of 12 between the extremes. This traces directly to the IS real-rate elasticity, estimated at 0.58 for Uganda and 0.59 for Kenya against 0.02 for South Africa (Section 18), and inherits that parameter's identification problem in full. **That two of the three agree closely is the reason the South African value is read as an identification failure rather than as a finding about South African households.**

The three results are not equally anchored. Uganda's and Kenya's monetary shock scales both reverted to a calibrated 0.50, so the *magnitude* of those responses is an assumption while their shape is estimated; Uganda's rate series also ends in 2019 (Section 4.2). South Africa's monetary shock scale *is* estimated and its rate series is complete — so its near-zero response is the best-measured of the three, and is therefore the most troubling, since it is also the least economically credible.

9.2 Fiscal policy

The multipliers are strikingly similar, which is the more interesting result given how different the two economies are.

Uganda: a deficit expansion of 0.87% of GDP raises growth 0.54 points — an impact multiplier of about 0.62. South Africa: an expansion of 2.00% of GDP raises growth 1.14 points — about 0.57. Both are well below one, and in both the effect on private consumption and investment is close to zero.

The totals agree, and the mechanisms usually offered to explain them do not.

- **Import leakage is not the explanation.** Non-oil import persistence is estimated at 0.004 in Uganda, 0.767 in Kenya and 0.059 in South Africa. Ugandan imports adjust almost within the year and Kenyan ones do not, yet the two multipliers are 0.62 and 0.66. The channel changes where the demand goes without changing how much output it buys.

- **Monetary offset is not the explanation either.** The largest policy-rate response to the fiscal impulse over ten years is 0.01 percentage points per standard deviation in Uganda, 0.07 in Kenya and 0.14 in South Africa. Against IS elasticities of 0.58, 0.59 and 0.02, none of those rate moves is large enough to matter for consumption.

- **The debt response is the same once it is scaled.** The expansion raises debt 0.61 points of GDP per standard deviation in Uganda, 0.72 in Kenya and 0.66 in South Africa. The unscaled figures differ by a factor of three, but only because two of the three shock scales are calibrated rather than estimated.

- **What the three share is the reversal.** Growth rises on impact and gives the rise back the following year, so the cumulative effect on the level of output after five years is +0.05 %, +0.02 and -0.02. Section 14 works this through.

Neither country should read this as "fiscal policy does not work". The result is conditional on an aggregate deficit shock rather than a specific instrument, on import elasticities estimated from 24 annual observations, and — for South Africa — on a monetary transmission parameter that Section 18 argues is not credible. Section 14 takes the fiscal question further.

9.3 Demand and cost-push shocks

Cost-push is where the three economies differ most. In Uganda a one-standard-deviation cost-push shock ($\sigma = 4.36$, the largest estimated domestic shock in its data) raises inflation 2.93 points and *cuts* growth 1.03. In South Africa the same shock ($\sigma = 0.99$, also estimated) raises inflation 0.67 and leaves growth essentially unchanged at -0.04. Kenya's shock scale is calibrated rather than estimated ($\sigma = 1.00$), so only the shape of its response is a result: inflation +0.75 and growth +0.24.

The Ugandan output cost of an inflation shock is almost entirely *the policy response to it*: inflation rises, the Taylor rule raises the real rate, and the interest-sensitive consumption block contracts. In South Africa the same sequence occurs but the final link is missing — the rate rises and consumption barely responds — so inflation is more persistent (backward-lookingness 0.90, against Uganda's 0.63) and output is more nearly insulated.

This is a coherent pair of stories, and it is worth stating what it would mean if taken at face value: disinflation is cheap in output terms in South Africa and expensive in Uganda. That conclusion rests entirely on the disputed elasticity, and should not be used without resolving it.

10. Policy simulation: external shocks

Figure 7: Uganda: responses to external shocks

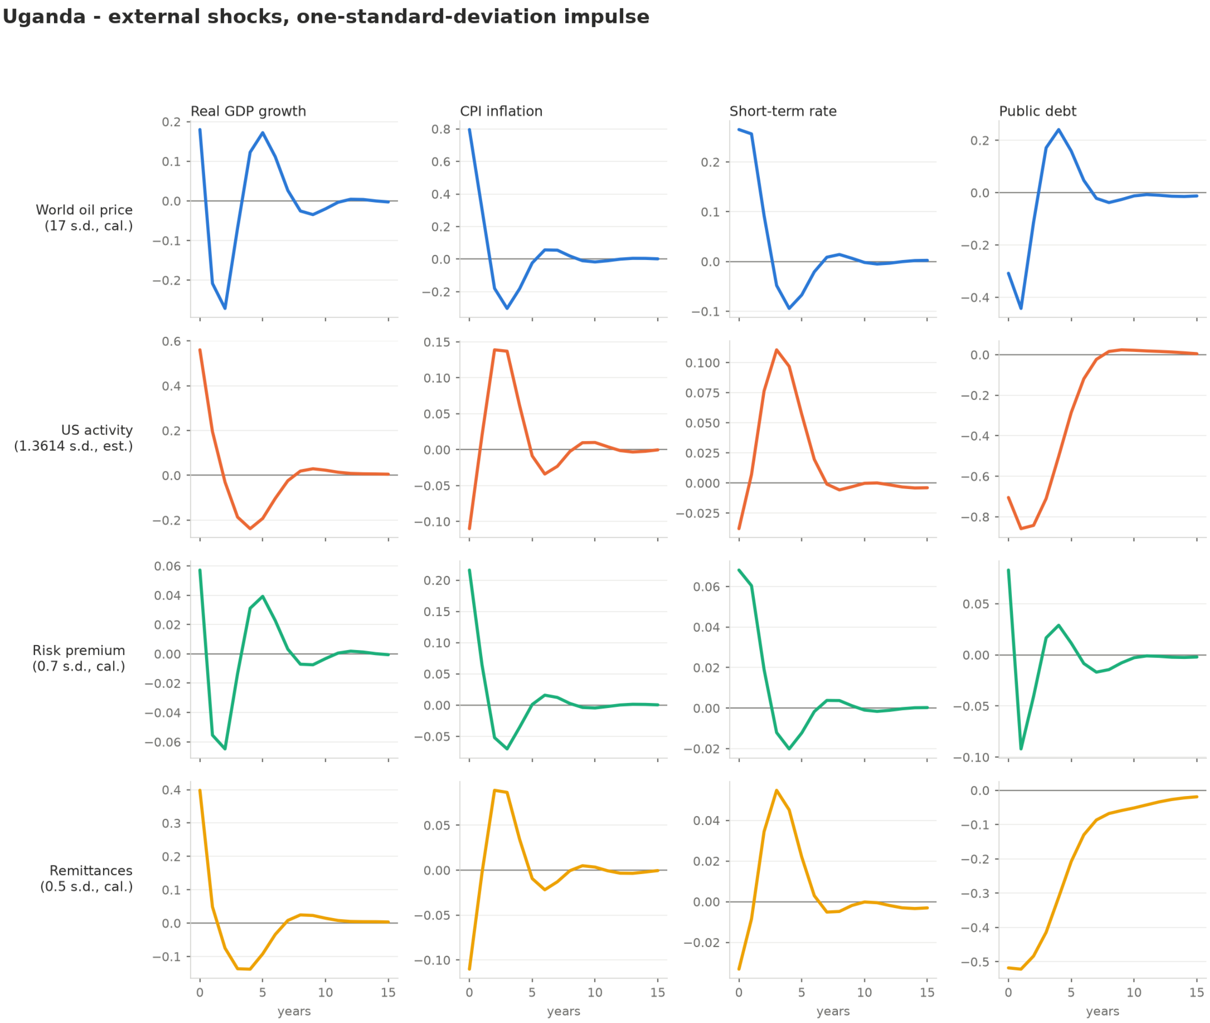

Notes: Rows are shocks, columns are observables; horizontal axis in years after impact. Three of the four shock scales are calibrated rather than estimated — they are among the variances the single-country sample cannot identify (Section 5.7) — so their shapes are model results and their scales are assumptions.

Figure 8: South Africa: responses to external shocks

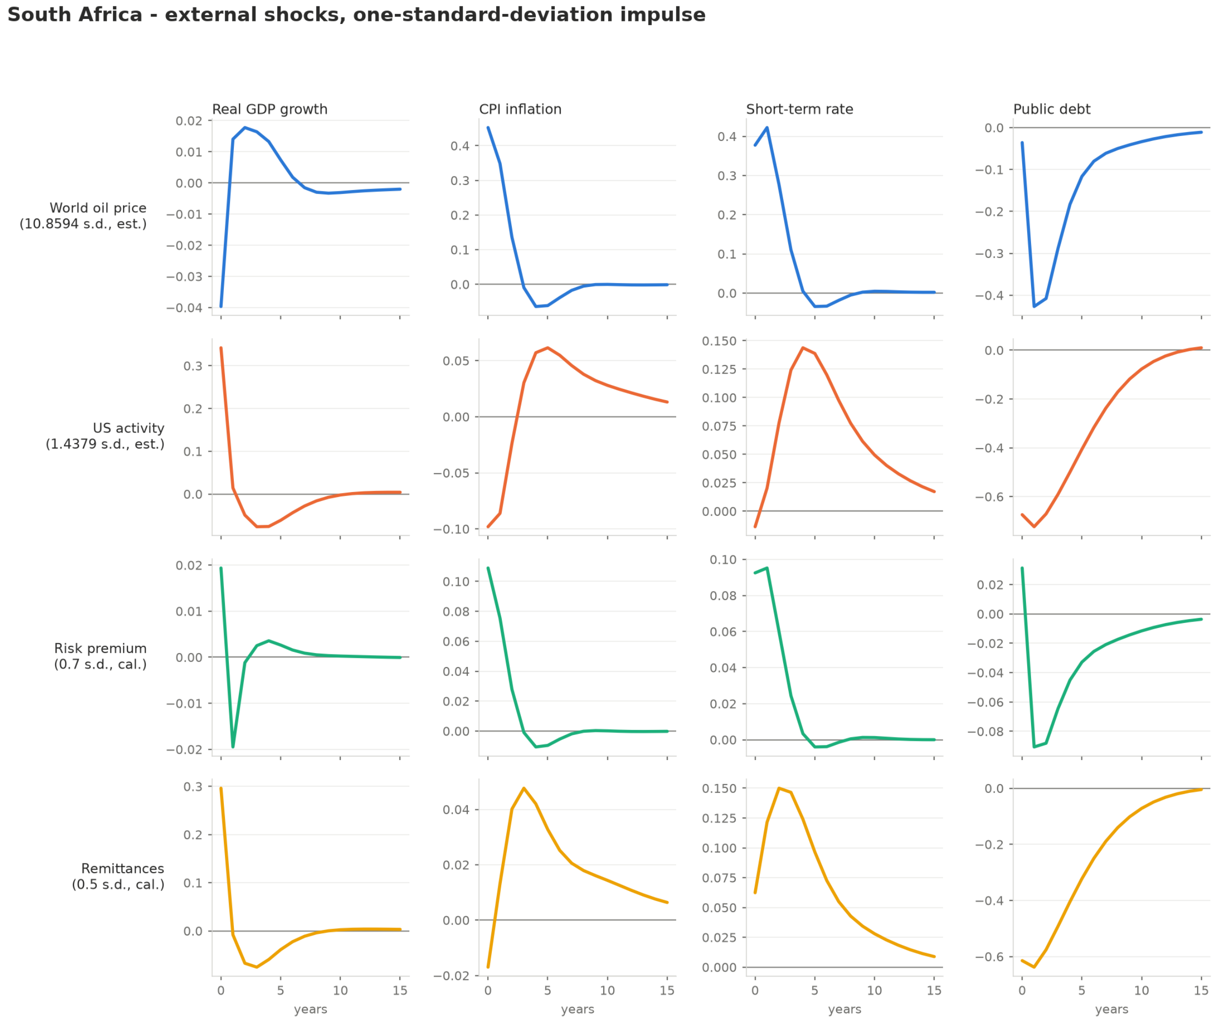

Notes: As Figure 7, for South Africa. The oil price shock is smaller here ($\sigma = 10.9$ against Uganda's 17.0) because the world-price process is estimated against each country's own moments even though the underlying series is identical — itself a symptom of the identification problem of Section 17.

Figure 9: Kenya: responses to external shocks

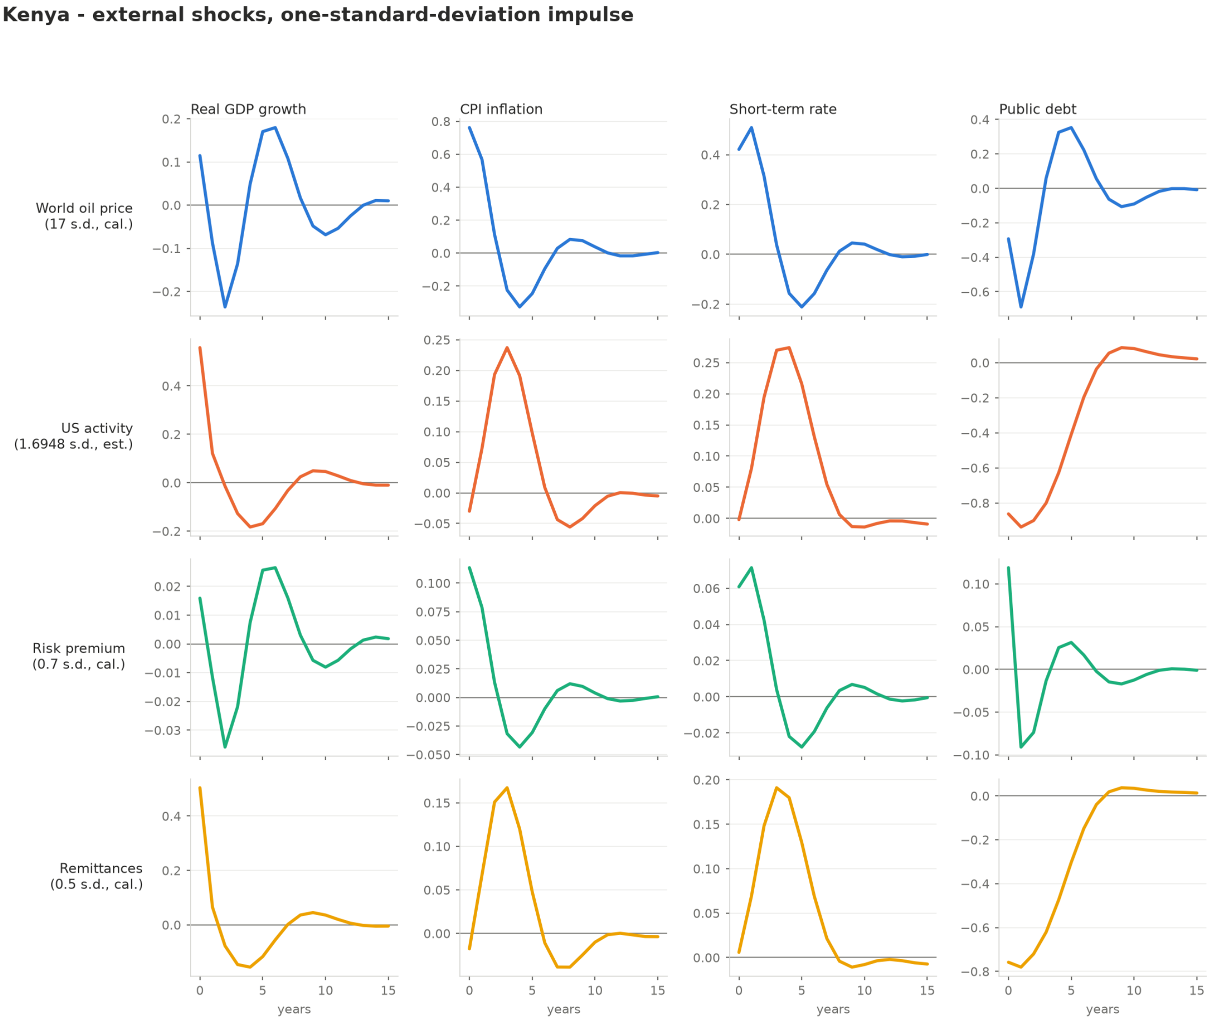

Notes: As Figure 7, for Kenya. The oil block is where Kenya's estimate is weakest (Section 18.3); the shape of the oil response here is the model's, and its scale is calibrated.

Table 10: Peak responses to external shocks

	Uganda				Kenya				South Africa			
Shock	σ	src	GDP gr.	Infl.	σ	src	GDP gr.	Infl.	σ	src	GDP gr.	Infl.
World oil price	17.00	cal.	-0.27	$+0.80$	17.00	cal.	-0.24	$+0.76$	10.86	est.	-0.04	$+0.45$
US activity	1.36	est.	$+0.56$	$+0.14$	1.69	est.	$+0.56$	$+0.24$	1.44	est.	$+0.34$	-0.10
Country risk premium	0.70	cal.	-0.06	$+0.22$	0.70	cal.	-0.04	$+0.11$	0.70	cal.	-0.02	$+0.11$
Remittances	0.50	cal.	$+0.40$	-0.11	0.50	cal.	$+0.50$	$+0.17$	0.50	cal.	$+0.30$	$+0.05$

Notes: Peak absolute response over a 15-year horizon, in percentage points for growth and inflation and percent of GDP for debt. "src" records whether σ was estimated (*est.*) or reverted to its calibrated value (*cal.*) after reaching the zero bound. **The sources differ by country in ways that matter**: a response whose scale is calibrated has an estimated shape and an assumed magnitude. Because σ differs by country the columns are comparable within a country but not across one; per-unit comparisons in the text divide by σ.

Oil. A one-standard-deviation world oil price rise raises Ugandan inflation 0.80 points and cuts growth 0.27; Kenya's inflation rises 0.76 and growth falls 0.24; South African inflation rises 0.45 and growth barely moves. **Uganda and Kenya are the oil-exposed pair here**, consistent with their stronger import-price pass-through — but all three effects are modest relative to the size of the shock, because oil imports are a small share of GDP in each.

This is also the block where the model fits Kenya worst. The oil moments carry the majority of Kenya's total moment misfit (Section 18.3), so the responses above should be read as the model's structure rather than as a measured elasticity for Kenya.

The oil shock is calibrated for Uganda and estimated for South Africa, which is worth pausing on: the underlying world price series is *identical* for both countries, so an estimator that identifies its variance in one country and not the other is telling us about the two samples, not about oil. This is the identification problem of Section 17 appearing in its clearest form, and it is why the fix is to estimate these variances once off the foreign series rather than country by country.

US activity. A US demand expansion raises growth 0.45 in Uganda and 0.34 in South Africa, transmitted through non-resource exports and, second-round, through investment's dependence on expected export demand. This is the one external shock whose scale is estimated in both countries, and the two responses are close — which is what one would expect of a common shock reaching two economies with similar trade openness.

Remittances. A 0.5%-of-GDP shock raises Ugandan growth 0.39 and South African growth 0.29. The Ugandan response is larger relative to the shock because remittances enter the consumption equation *directly* (equation (11), coefficient c_6) and run above 3% of GDP there. For South Africa the channel exists in the model but the flows are smaller.

Country risk premium. Small for growth in both countries — -0.07 and -0.02 — because the premium works mainly through the exchange rate and the debt-revaluation channel rather than through demand. **Its fiscal consequence exceeds its real consequence in both**, which is the point: a sudden stop shows up in the debt ratio well before it shows up in the growth rate, and South Africa's higher foreign-currency debt makes that channel the more material of the two.

11. What the model is for: shocks, transmission and interpretation

A model of this size earns its keep only if the whole of it is wired up and can be interrogated. This section establishes that, then works through the model's response to each family of shocks and the historical decomposition of each economy. Each chart is read where it appears, and the published evidence is brought in at the point the result it bears on is being discussed.

11.1 The model is complete; its shocks are not all identified

The implementation carries all 83 structural shocks of Berg et al. (2023) and all 353 equations of its appendix, in a system that is exactly square: 353 equations in 353 endogenous variables. Every variable, shock and parameter name printed in the paper's appendix tables is present in the code. Appendix C sets out equations 1–82 in full.

Completeness of the equation list is not identification of the shock list, and the distinction governs everything below.

Figure 10: The structural shocks, by block and by whether the estimation identified their scale

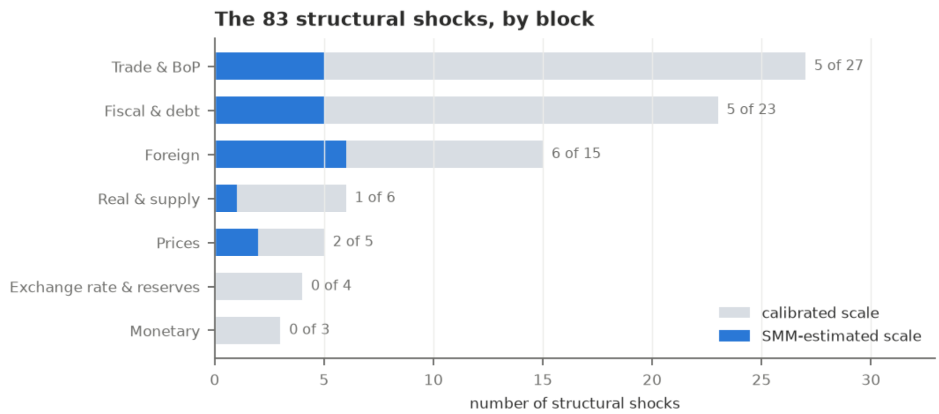

Each bar is one block of the model and its length is the number of structural shocks in that block. The blue portion counts the shocks whose standard deviation the Simulated Method of Moments identified from Uganda's data; the grey remainder kept its calibrated value. The label gives the two counts.

Figure 10 is the first thing to read, because it qualifies every impulse response that follows. Of 83 shocks, only 19 have a standard deviation the estimation could recover from twenty-four annual observations. The distribution matters more than the total: the trade and balance-of-payments block gets 5 of 27 and the fiscal block 5 of 23, while the monetary block gets **none of 3** and the exchange-rate and reserve block **none of 4**.

The instruments FINEX is most distinctive for are precisely the ones this country's data cannot pin down. It would be easy to read that as an estimator failure, and in part it is — with a sample this short, weak identification is the expected condition, and the small-sample distribution of a simulated-moments estimator is poorly approximated by its asymptotic normal, so inference should be conservative (Ruge-Murcia, 2012). But the pattern across blocks points at something else as well. Monetary transmission in low-income countries is found to be weak and unreliable for structural reasons — small banking sectors, illiquid financial markets, uncompetitive banking and lending concentrated at short maturities, all of which blunt the pass-through from the policy rate to the rates households and firms face (Mishra et al., 2010; Mishra and Montiel, 2012). If transmission is genuinely weak, monetary shocks leave little signature in the second moments, and an estimator that declines to identify them is reporting that absence rather than manufacturing it. The same reading applies to the exchange-rate block: the Fund's Integrated Policy Framework concludes that intervention has limited benefit under normal conditions with liquid markets and becomes effective mainly under large shocks combined with specific frictions — unhedged currency mismatches, shallow market depth, or de-anchoring expectations (International Monetary Fund, 2023) — and a sample containing few such episodes offers little to identify an intervention shock from.

The practical consequence runs through the rest of this section. The *shape* of these responses is a property of the model and is unaffected; the *size* is a calibration choice. The charts support statements of the form "these instruments act through different channels and on different horizons", but not “intervention of this size is worth a rate rise of that size”.

11.2 Transmission, family by family

11.2.1 Fiscal spending and revenue

The four fiscal disturbances separate the instruments a finance ministry actually controls: a financed spending impulse (the structural deficit), two changes in the composition of spending at a given deficit, and a revenue windfall.

Table 11: Fiscal spending and revenue shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
Structural deficit	+0.50	+0.08	+1.32	+0.13	+1.14	-0.01
Government consumption	-0.03	-0.26	-0.02	-0.18	-0.00	-0.02
Transfers	+0.26	-0.70	+0.27	-0.35	-0.04	-0.22
Non-tax revenue	+0.15	+0.26	+0.19	+0.37	+0.17	+0.49

One standard deviation per shock, percentage points of annual growth. “impact” is the response in the year of the shock; “5-yr” is the sum over the first five years. A dash marks a shock the country’s run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

Figure 11: Fiscal spending and revenue shocks — Uganda

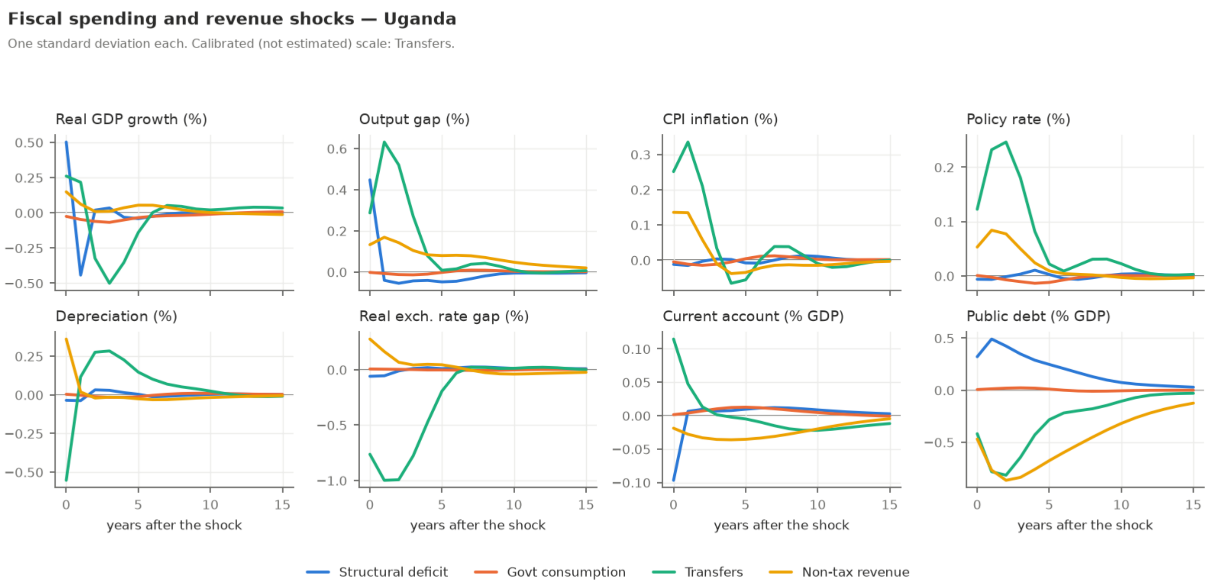

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model’s, the size is an assumption.

Read the top-left panel of each figure first. The deficit shock (blue) jumps on impact and then falls

Figure 12: Fiscal spending and revenue shocks — Kenya

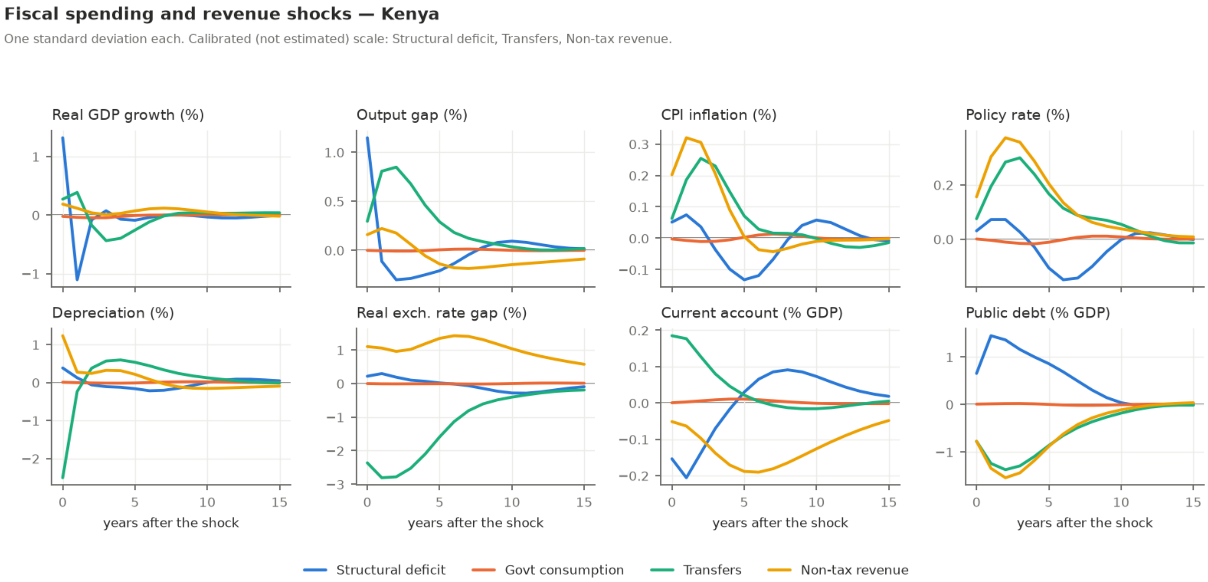

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 13: Fiscal spending and revenue shocks — South Africa

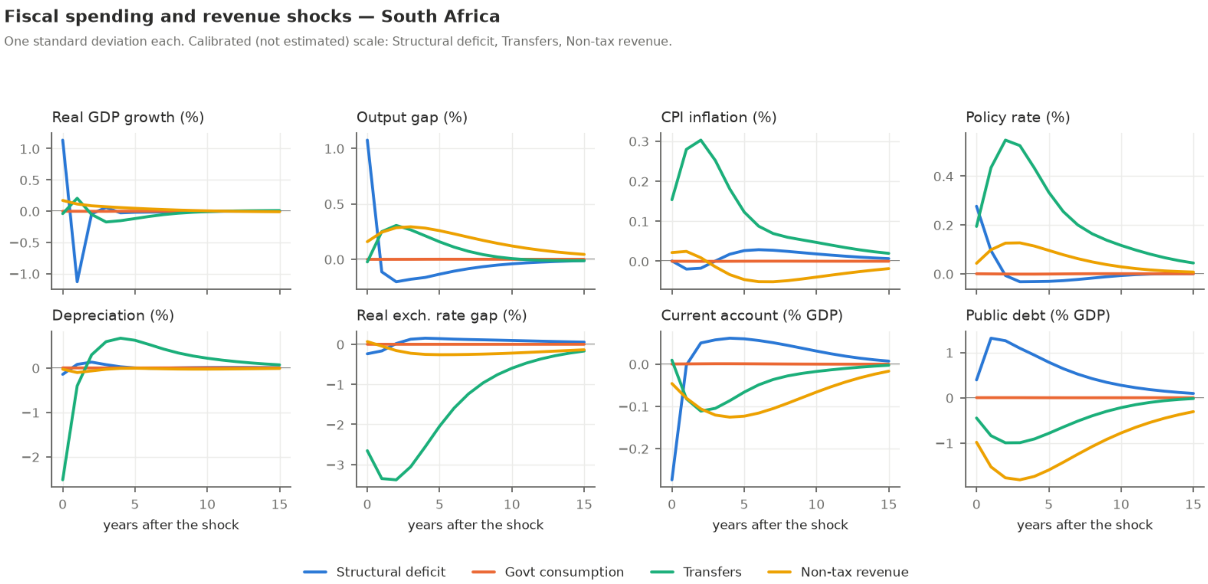

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

back through zero — the signature of a multiplier below one acting on a one-year impulse. Table 11 puts numbers on it: real GDP growth rises +0.50, +1.32 and +1.14 percentage points on impact in Uganda, Kenya and South Africa, but the five-year cumulative responses are only +0.08, +0.13 and -0.01. The bottom-right panel shows where the impulse persists instead: public debt, which in Uganda peaks at +0.49 per cent of GDP and decays far more slowly than the output effect.

Multipliers of this size are the consensus for these economies rather than an artefact of our calibration. The IMF's Independent Evaluation Office puts the overall fiscal multiplier for the low-income (PRGT) group at 0.6, noting that this is *higher* than both the earlier literature and the multipliers used in programme design (Independent Evaluation Office of the IMF, 2021); a recent Fund study of sub-Saharan Africa finds a consolidation of one per cent of GDP reduces output by about 0.54 per cent after two years (International Monetary Fund, 2026). Our figures are the same order, reached structurally rather than by reduced-form projection.

Two cautions. The cross-country literature finds enormous heterogeneity — government-consumption multipliers across twenty low-income sub-Saharan economies span 0.01 to 3.72 at five years (Woldu, 2023) — so a clustering as tight as ours across three countries more likely reflects a shared calibration than a shared economy. And the near-flat lines for government consumption and transfers in these panels are not a finding that composition does not matter; they are one-standard-deviation shocks, and those standard deviations are small.

11.2.2 Tax-rate instruments

FINEX carries four separate tax rates — on consumption, on income, on non-oil imports and on resource output — each entering through a different equation: the consumption tax through the consumption ratio of equation (25), the income tax through the investment trend, the tariff through the import ratio of equation (66), and the royalty through resource-producer profit in equation (59). This is what lets a tax-reform question be put to the model rather than only a deficit question.

Table 12: Tax-rate instruments shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
Consumption tax	+0.29	+0.23	+0.33	+0.19	+0.13	+0.08
Income tax	+0.55	+0.03	+0.53	-0.14	+0.20	-0.10
Non-oil tariff	+0.04	+0.07	+0.05	+0.08	+0.02	+0.02
Resource royalty	+0.01	+0.01	+0.01	+0.01	+0.01	+0.02

One standard deviation per shock, percentage points of annual growth. "impact" is the response in the year of the shock; "5-yr" is the sum over the first five years. A dash marks a shock the country's run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

These panels contain the most easily misread result in the report, and it is visible directly in the top-left

Figure 14: Tax-rate instruments — Uganda

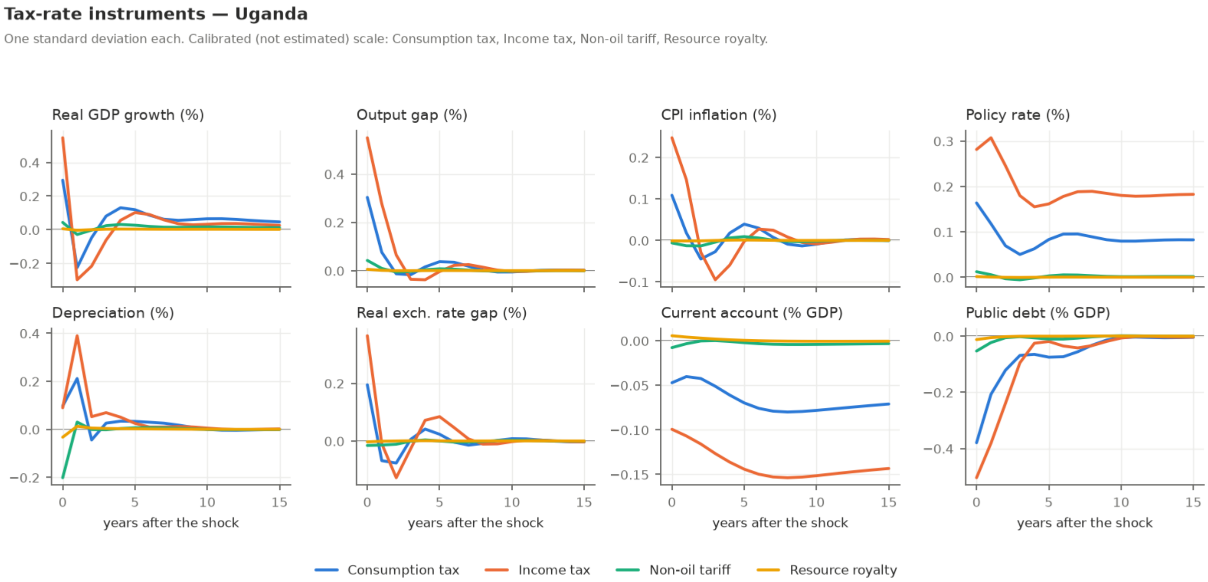

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 15: Tax-rate instruments — Kenya

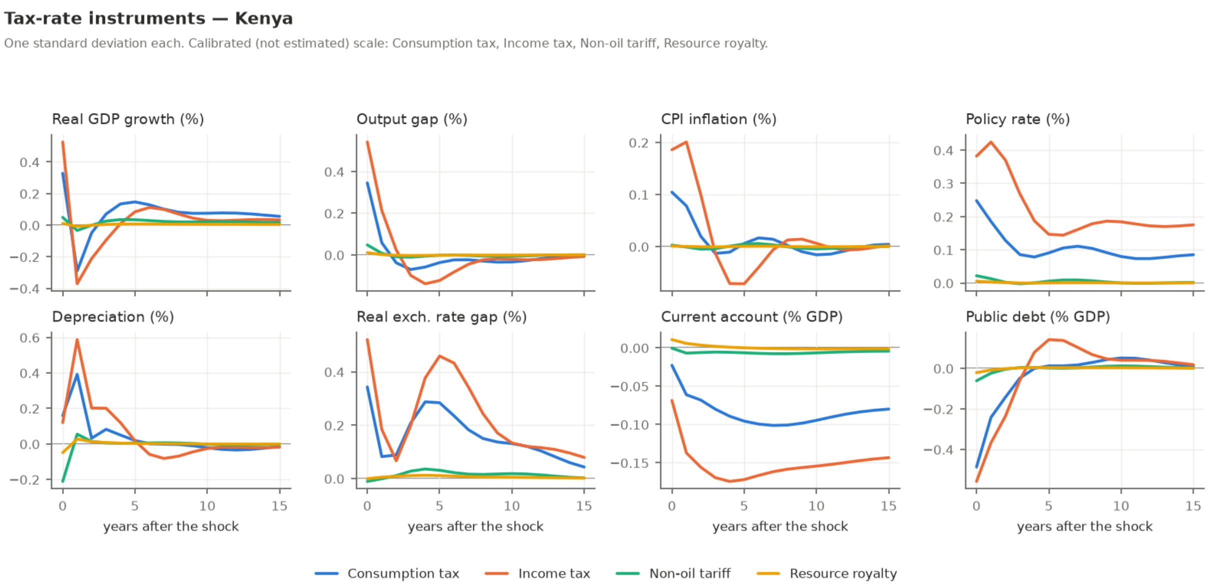

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 16: Tax-rate instruments — South Africa

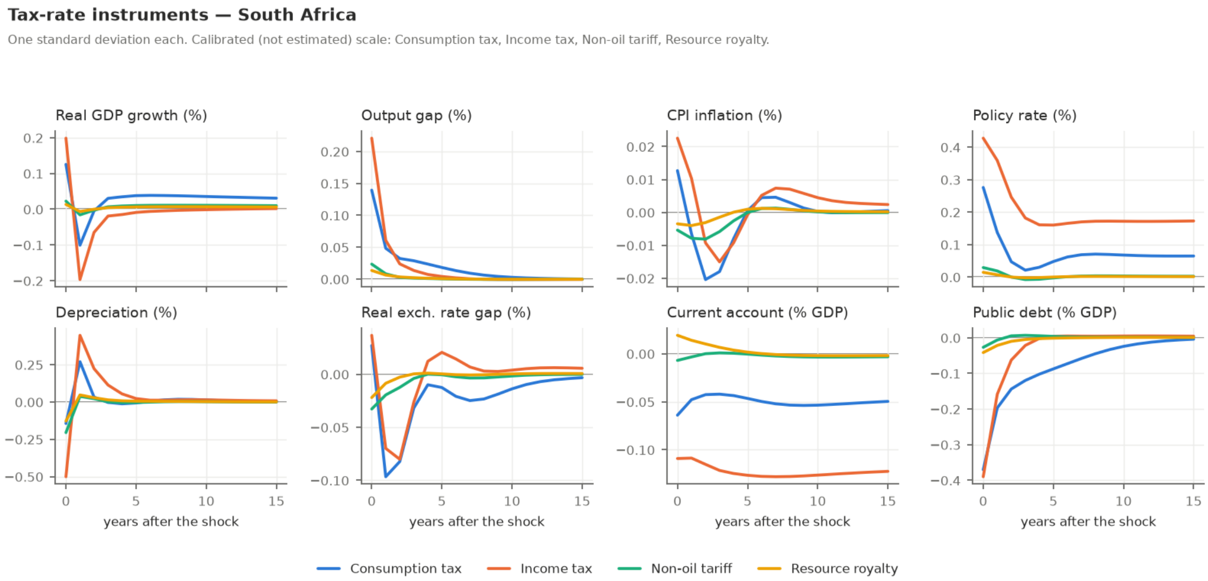

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

one: a *rise* in the income tax raises GDP growth on impact, by +0.55 percentage points in Uganda. That is not a sign error. The budget closes in FINEX, so revenue that arrives is spent: government expenditure rises +0.31 per cent of GDP on impact, and the spending lands immediately while the tax disincentive works through investment with a lag. The same panel shows the reversal in the following year, and private investment growth falls -0.57 on impact with a larger fall after that. Cumulated over five years the GDP response is only +0.03.

So the chart has to be read as a whole rather than at its peak: the peak reports the financing and says nothing about the tax. What the model reproduces here, endogenously, is exactly the composition asymmetry the empirical literature documents — tax-based adjustment carrying a smaller short-run output cost than spending-based adjustment (International Monetary Fund, 2026).

The tariff and royalty lines sit almost on the axis in every panel. That is a statement about the size of a one-standard-deviation change in those rates, not about their policy irrelevance: the royalty acts on resource-producer profit, which is a small share of GDP in all three economies.

11.2.3 Monetary policy and prices

This family separates the policy instrument from the risk premium sitting between the policy rate and the rate the private sector faces, and from the cost-push shock that forces the policy trade-off. The distinction between a temporary premium shock and a permanent one matters for the exchange rate, because the two imply different expected depreciation paths from the same current cost of funds.

Table 13: Monetary policy and prices shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
Policy rate	-0.83	-0.00	-0.59	+0.02	-0.20	-0.13
Risk premium (gap)	+0.06	-0.05	+0.02	-0.05	+0.02	+0.00
Risk premium (trend)	-0.00	-0.04	+0.00	-0.04	-0.01	-0.03
Cost-push	+0.85	-0.34	+0.24	-0.16	+0.04	+0.04

One standard deviation per shock, percentage points of annual growth. "impact" is the response in the year of the shock; "5-yr" is the sum over the first five years. A dash marks a shock the country's run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

Figure 17: Monetary and price shocks — Uganda

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 18: Monetary and price shocks — Kenya

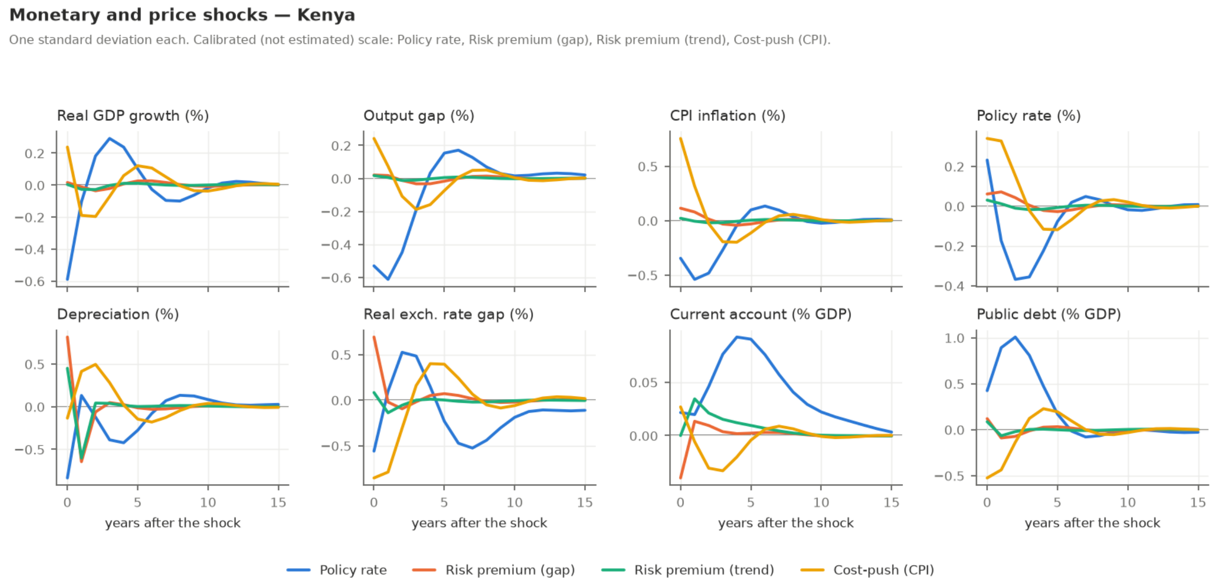

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

The policy-rate line (blue) behaves as a textbook tightening on impact: GDP growth -0.83, CPI inflation -0.73 in the following year, and the currency appreciating — depreciation falls 1.36 percentage points. The oscillation afterwards, with growth crossing above zero around years three to four, is the usual overshoot of a gap model with a forward-looking component, not a second policy action.

The cost-push line (orange) is the one to compare it against, because it is the only shock in this family whose scale the estimation identified. It raises inflation +2.93 on impact and cuts growth -1.03 in the following year: the trade-off the policy rule then has to resolve, visible in the policy-rate panel where the rate rises in response.

Everything else in these panels should be read as shape, not size. As Figure 10 showed, none of the three monetary shock scales was identified, which is consistent with the structural weakness of monetary transmission documented for low-income economies (Mishra et al., 2010; Mishra and Montiel, 2012) rather than with a coding or data failure. The current-account panel makes the point concretely: the response to a rate shock is two orders of magnitude smaller than the response of public debt in the panel beside it, and on a calibrated scale that difference should not be taken as a quantitative claim.

11.2.4 Exchange rate and reserves

This family is FINEX's distinctive contribution and the reason it exists alongside conventional gap–trend projection models: foreign exchange intervention, the reserve position and its target, and a capital inflow tax are *separate* instruments. A central bank facing pressure on the currency can be represented as choosing between raising the policy rate, selling reserves and taxing inflows, and the model answers

Figure 19: Monetary and price shocks — South Africa

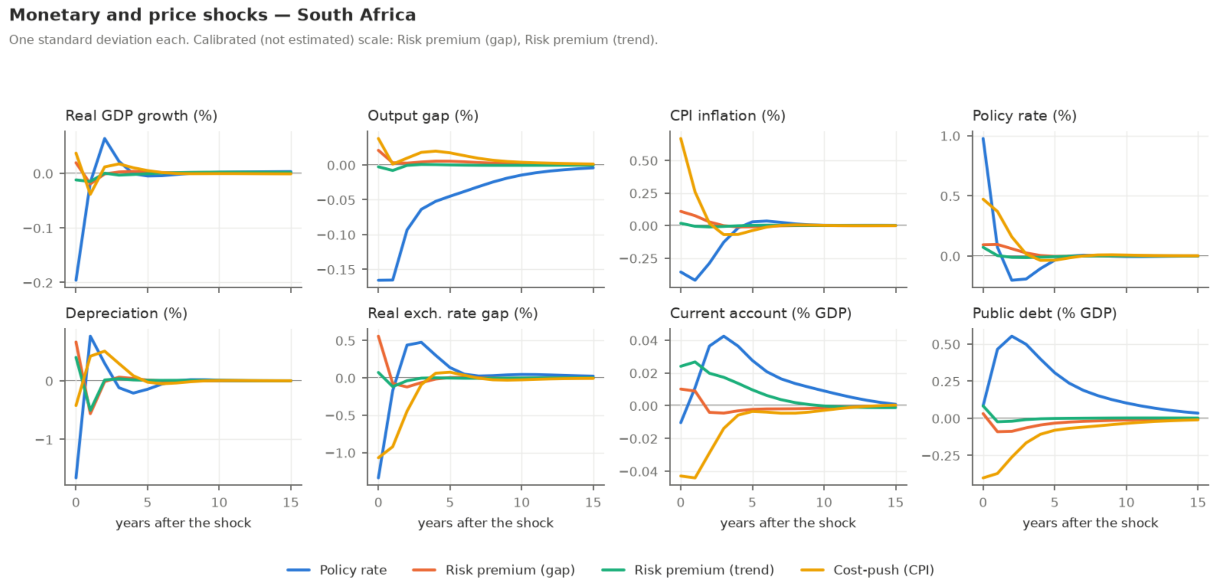

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

differently for each.

Table 14: Exchange rate and reserves shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
FX intervention	-0.00	+0.00	+0.00	+0.00	+0.00	-0.00
Reserve position	-0.00	+0.00	-0.00	+0.00	-0.00	+0.00
Reserve target	+0.00	-0.00	+0.00	-0.01	+0.00	-0.00
Capital inflow tax	+0.04	-0.08	+0.02	-0.10	-0.00	-0.06

One standard deviation per shock, percentage points of annual growth. "impact" is the response in the year of the shock; "5-yr" is the sum over the first five years. A dash marks a shock the country's run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

These are the flattest panels in the report, and the reason is worth stating plainly rather than leaving the reader to infer it. Every scale in this family is calibrated, and two of the four are calibrated at or near zero, so the lines carry almost no amplitude. The intervention shock moves the depreciation rate by 0.03 percentage points on impact; the capital inflow tax, the largest mover here, by +0.34, with the current account improving +0.07 per cent of GDP at its peak.

What the figure does show, and what it is legitimate to take from it, is that the four instruments have visibly different *profiles*: the inflow tax acts on the exchange rate and the current account and persists, while intervention acts on the exchange rate and dies within two years. That is the qualitative

Figure 20: Exchange-rate and reserve shocks — Uganda

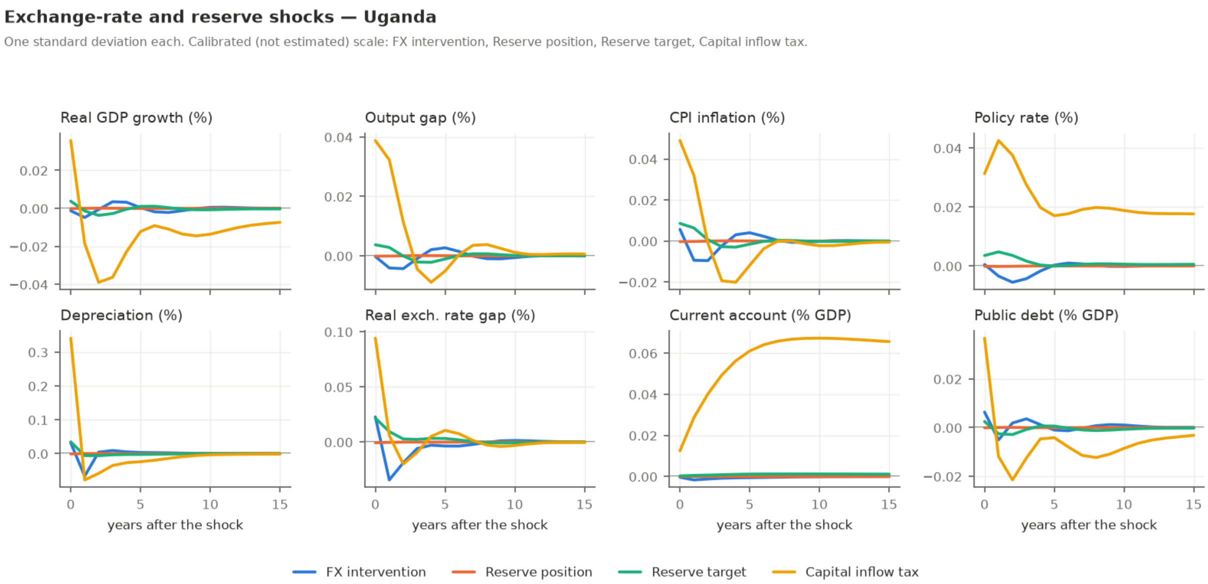

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 21: Exchange-rate and reserve shocks — Kenya

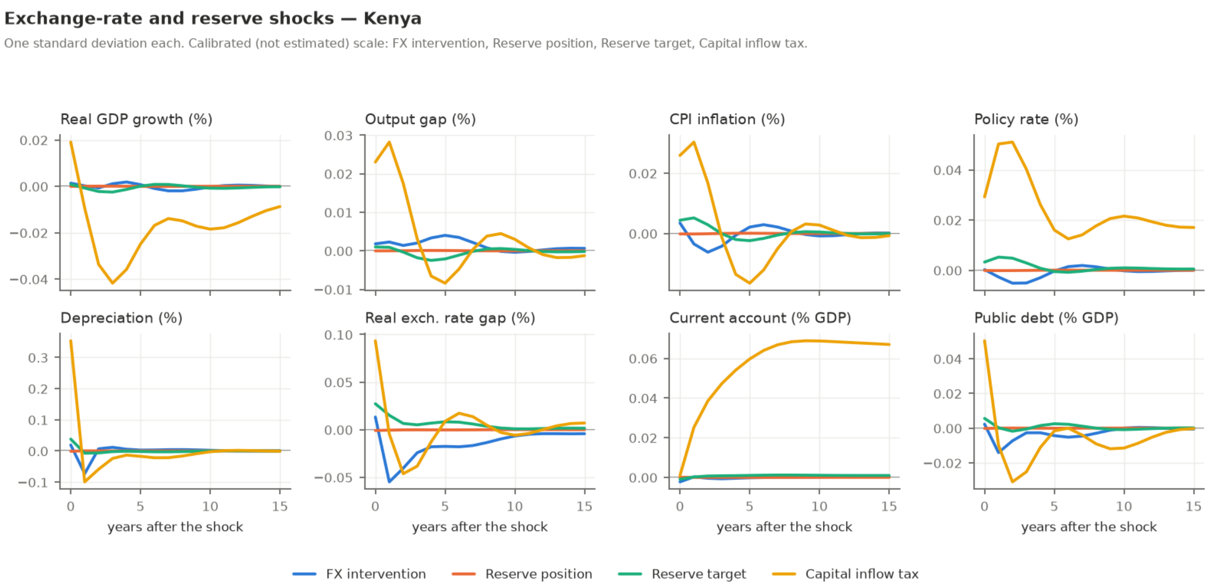

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 22: Exchange-rate and reserve shocks — South Africa

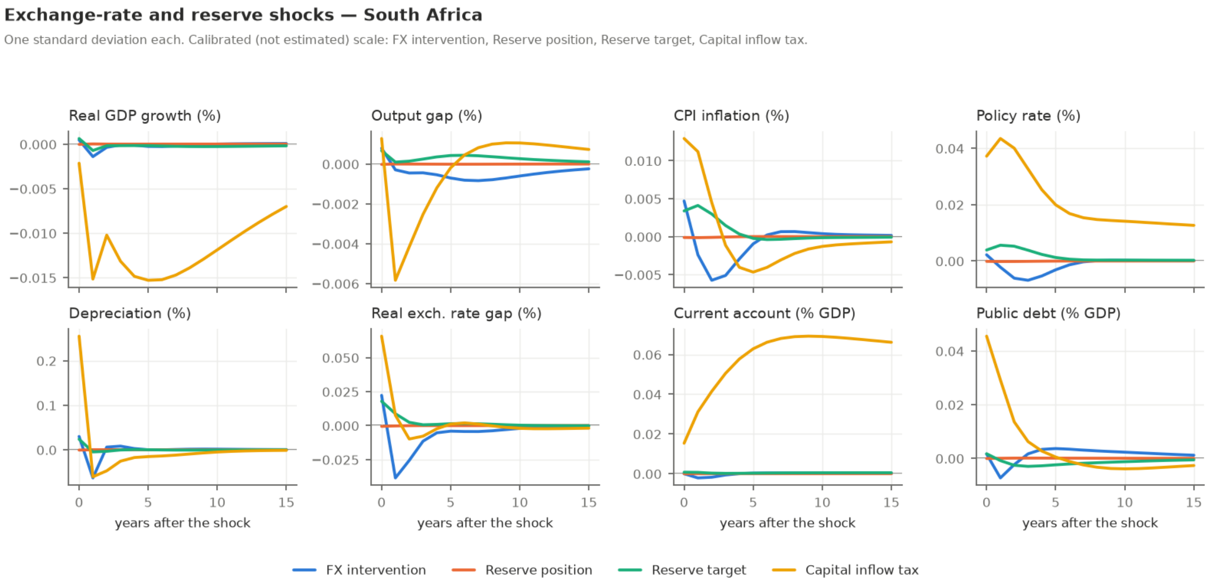

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

instrument-choice question the model can answer.

The quantitative one it cannot, and the reason is not only the short sample. The Integrated Policy Framework concludes that intervention's benefits under normal conditions with liquid markets are limited, and that it becomes effective mainly under large shocks combined with unhedged currency mismatches, shallow market depth, or de-anchoring expectations (International Monetary Fund, 2023). A sample with few such episodes contains little information about the size of an intervention shock — which is exactly what the estimation reported when it left all four scales at their calibrated values.

11.2.5 Balance of payments

This is where an African application is most exposed. Remittances, non-resource export demand, non-oil import demand and the foreign output gap are shocks these economies receive rather than choose, and the first is large enough in Uganda and Kenya to be a macroeconomic event rather than a line in the current account.

The remittance line (blue) is the one to follow across the panels, because it shows both sides of a phenomenon the literature treats as two separate findings. In the growth panel it raises GDP growth +0.40 percentage points and in the current-account panel improves the balance +0.48 per cent of GDP — the shock absorber. In the depreciation panel it appreciates the currency, cutting the depreciation rate by 1.17 percentage points — the cost.

Table 15: Balance of payments shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
Remittances	+0.40	+0.10	+0.50	+0.19	+0.30	+0.09
Non-resource exports	+0.59	+0.13	+0.04	-0.04	+0.27	+0.00
Non-oil imports	-0.46	-0.01	-0.43	-0.49	-0.35	+0.00
US output gap	+0.56	+0.30	+0.56	+0.35	+0.34	+0.15

One standard deviation per shock, percentage points of annual growth. “impact” is the response in the year of the shock; “5-yr” is the sum over the first five years. A dash marks a shock the country’s run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

Figure 23: Balance-of-payments shocks — Uganda

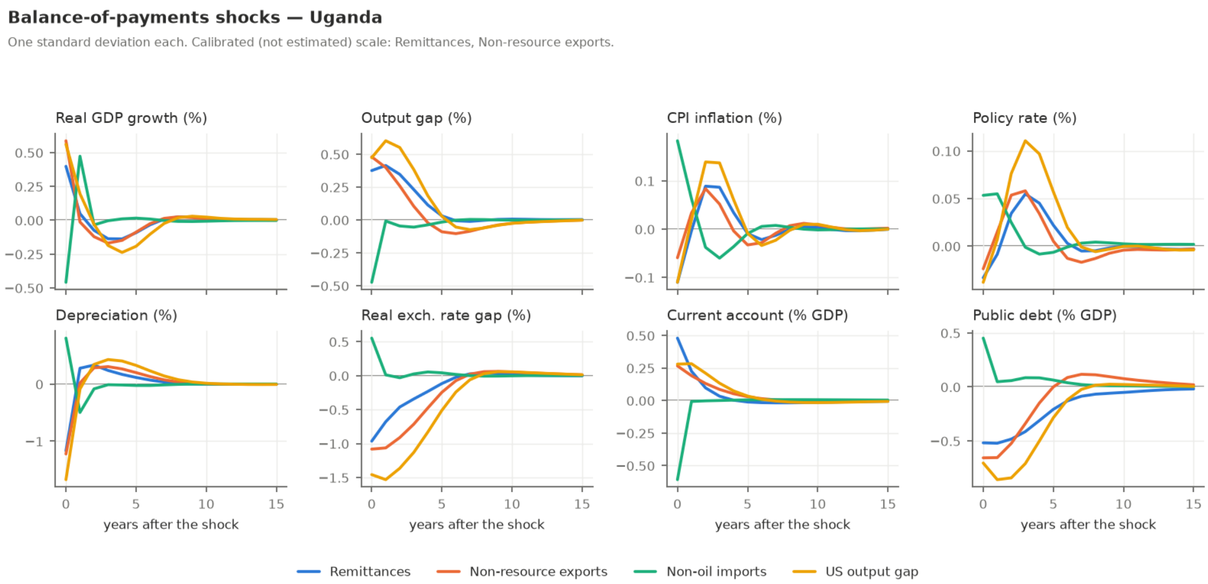

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model’s, the size is an assumption.

Figure 24: Balance-of-payments shocks — Kenya

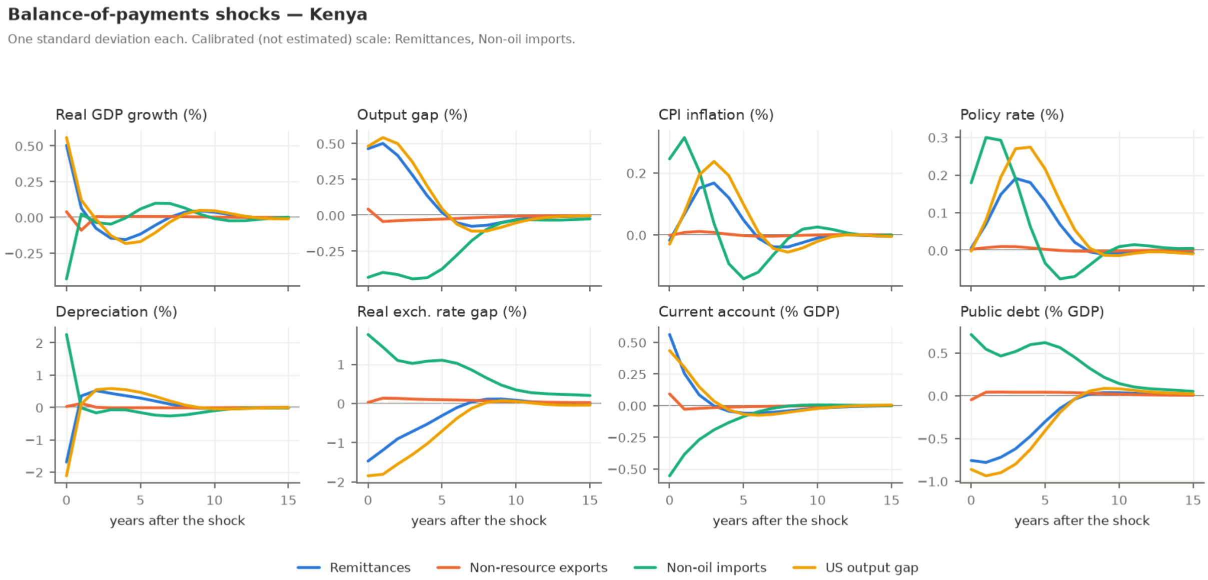

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 25: Balance-of-payments shocks — South Africa

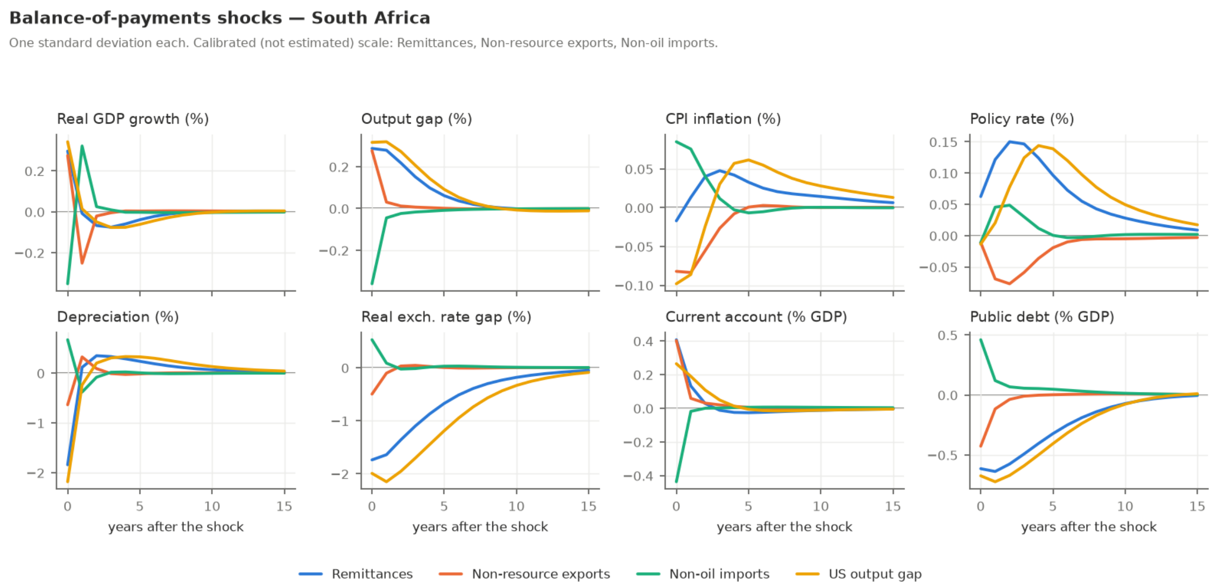

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Both match the record. Remittances to Uganda are more stable than its other foreign-exchange sources, behave countercyclically, and are found to reduce output volatility (International Monetary Fund, 2016); the appreciation is the same mechanism seen from the other side, and is the Dutch-disease concern that recurs in this literature, with rising recipient spending lifting domestic prices and the real exchange rate at the expense of the tradable sector. That the model produces the stabilising quantity effect and the relative-price cost from one shock, with neither imposed, is a point in favour of the external block.

The US output gap (yellow) is the largest single mover in the depreciation panel at -1.67 percentage points, against +0.56 for GDP growth — for these economies the external cycle arrives through the exchange rate at least as much as through export volumes.

11.2.6 Debt and sovereign financing

This family separates the authorities' targets — for the debt ratio and its currency composition — from the financing conditions they face. The sovereign premium shock is the channel through which a loss of market confidence reaches the real economy; the external financing shock represents a change in the availability of official borrowing rather than its price.

Table 16: Debt and sovereign financing shocks: response of real GDP growth

	Uganda		Kenya		South Africa	
Shock	impact	5-yr	impact	5-yr	impact	5-yr
Debt target	+0.15	+0.24	+0.16	+0.22	+0.14	+0.17
FX-debt share target	-0.03	+0.02	-0.00	+0.04	-0.01	-0.01
External financing	+0.01	+0.05	-0.01	-0.01	-0.01	+0.02
Sovereign premium	+0.04	+0.05	-0.00	+0.02	+0.04	+0.03

One standard deviation per shock, percentage points of annual growth. "impact" is the response in the year of the shock; "5-yr" is the sum over the first five years. A dash marks a shock the country's run does not contain. Where the two columns differ it is the cumulative column that answers a policy question.

The public-debt panel is the informative one here, and it looks different from every other panel in this section: the debt-target line does not return to zero within the horizon, peaking at +0.96 per cent of GDP at year 15 and still rising at the end of the window. That is the intended behaviour — a change in the target is a change in where the debt ratio is heading, not a disturbance around an unchanged path — and it is a useful visual reminder that not every shock in this model is transitory.

The sovereign premium line shows the confidence channel: GDP growth -0.19 at its trough, with the effect building over several years rather than landing at once, because it reaches activity through borrowing costs and investment rather than directly. The external financing shock moves the exchange rate most, +1.18 percentage points on the depreciation rate.

One limitation belongs here rather than in a footnote. FINEX has no sovereign-default mechanism.

Figure 26: Debt and sovereign financing shocks — Uganda

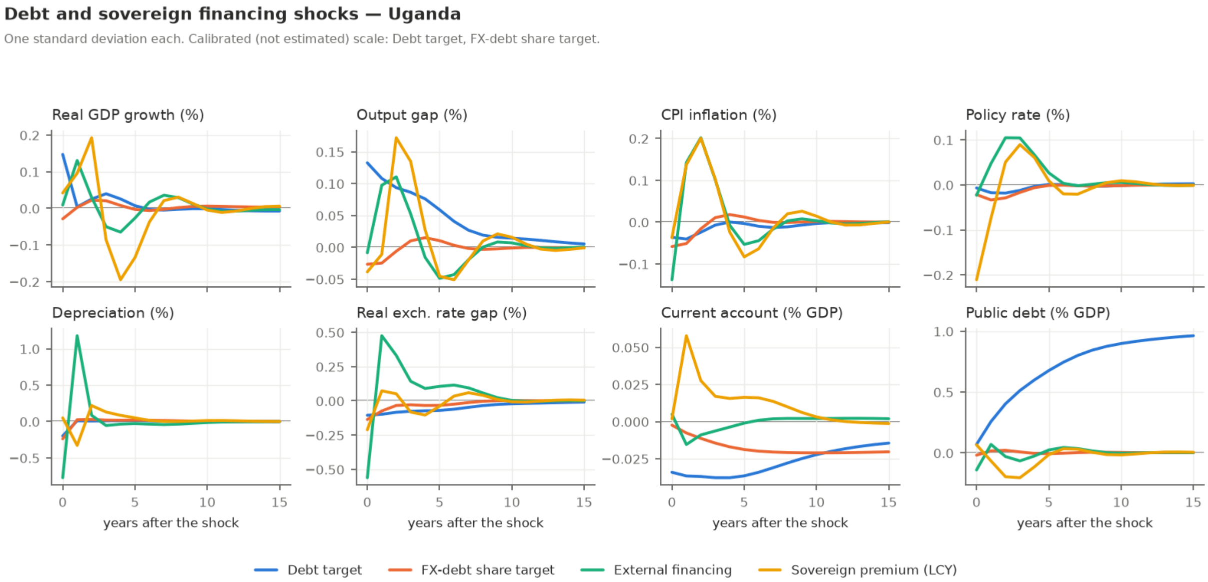

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 27: Debt and sovereign financing shocks — Kenya

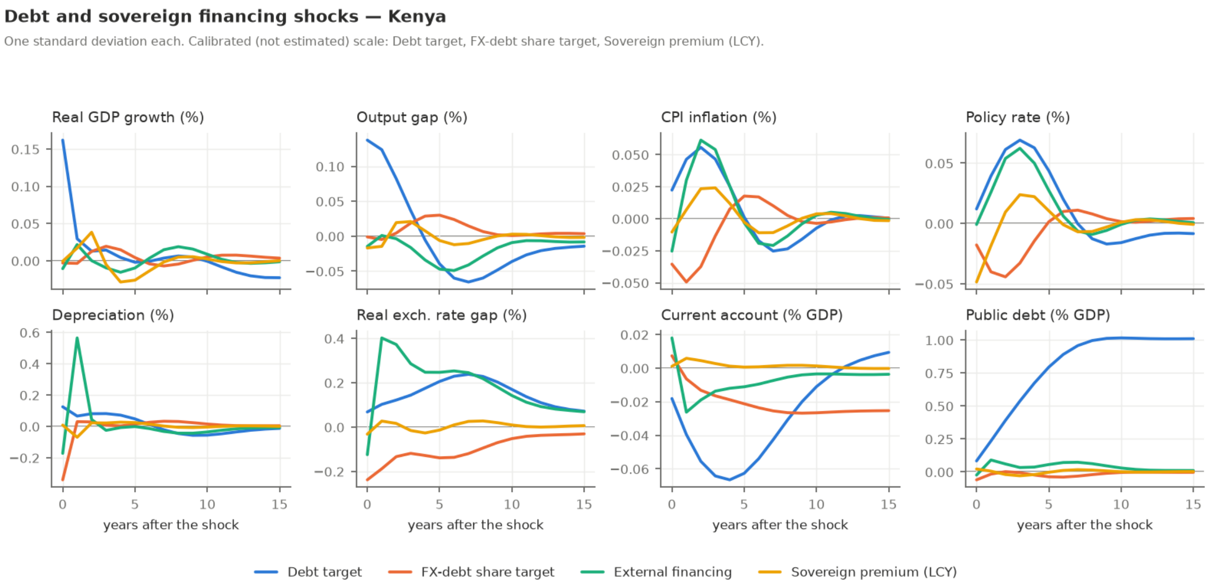

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

Figure 28: Debt and sovereign financing shocks — South Africa

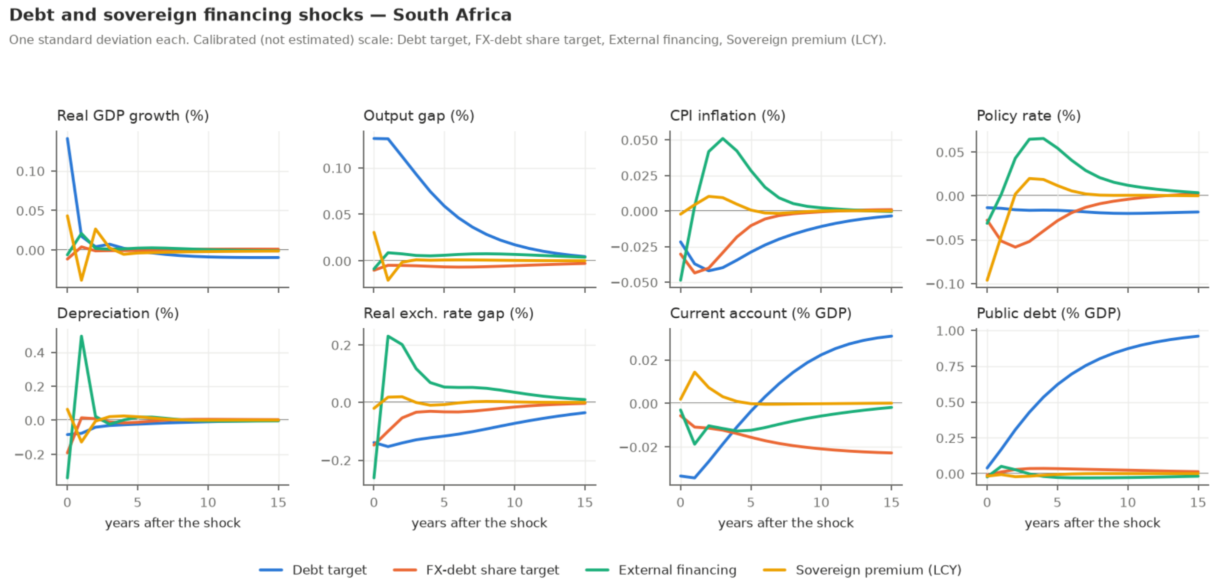

Each panel is one macroeconomic variable; each coloured line is one shock, hit with a single standard deviation in year 0. The horizontal axis is years since the shock, and zero on the vertical axis is the no-shock baseline, so a line above zero means the variable is higher than it would otherwise have been. Scales listed as calibrated in the panel note were not identified by the estimation: the shape of those responses is the model's, the size is an assumption.

In the wider set of African estimates reported elsewhere in this project, the debt moments dominate the misfit for a country that restructured inside the sample, because the debt block cannot reproduce a discontinuity it has no representation for. These panels describe a sovereign under stress, not a sovereign in default.

11.3 Historical decomposition: which channel moved the data

An impulse response says what the model does to a shock. A historical decomposition says which shocks the model thinks actually happened, and for constructing a narrative about the past it is the more useful of the two. The charts below attribute each observable's deviation from steady state to nine channels by cumulative group-wise re-simulation, so the contributions sum exactly to the total.

Two things about these charts have to be said before any of them is read.

First, the contributions are large and strongly offsetting. Individual channels reach twenty percentage points on a growth series whose own variation is a few points, because splitting a small total out of a nonlinear model leaves big pieces that nearly cancel. A bar that is tall in both directions is not a sign of a violent year; it is a sign that the decomposition has apportioned a quiet year between channels that largely offset. The right way to read the chart is the net balance within a year.

Second, the method is order-dependent: the groups are added in a fixed sequence and a different order would apportion the overlap differently. The ordering used here is stated in the appendix and held constant across countries so that the three are comparable with each other, which is the comparison

Figure 29: Contributions by channel — Uganda

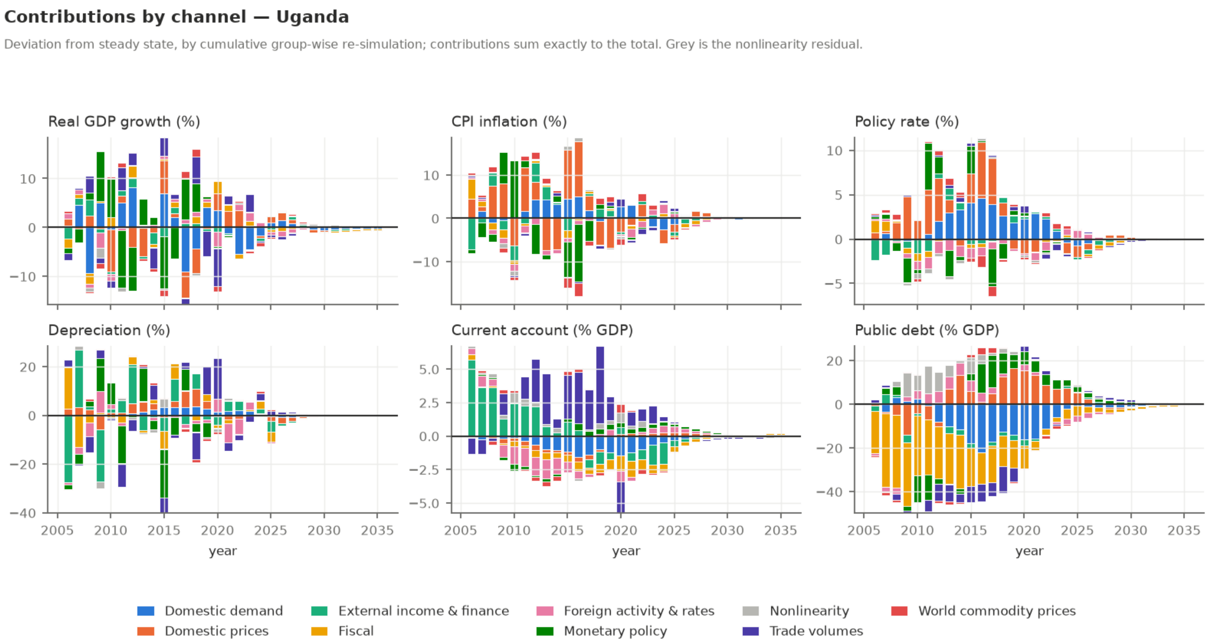

Each panel is one variable; each bar is one year, split into coloured segments, one per channel. Segments above the axis push the variable up that year and segments below push it down, so what the reader should take from a bar is the balance between the two, not the height of any one segment. Grey is the nonlinearity residual rather than a channel. The ordering is fixed and the method order-dependent.

Figure 30: Contributions by channel — Kenya

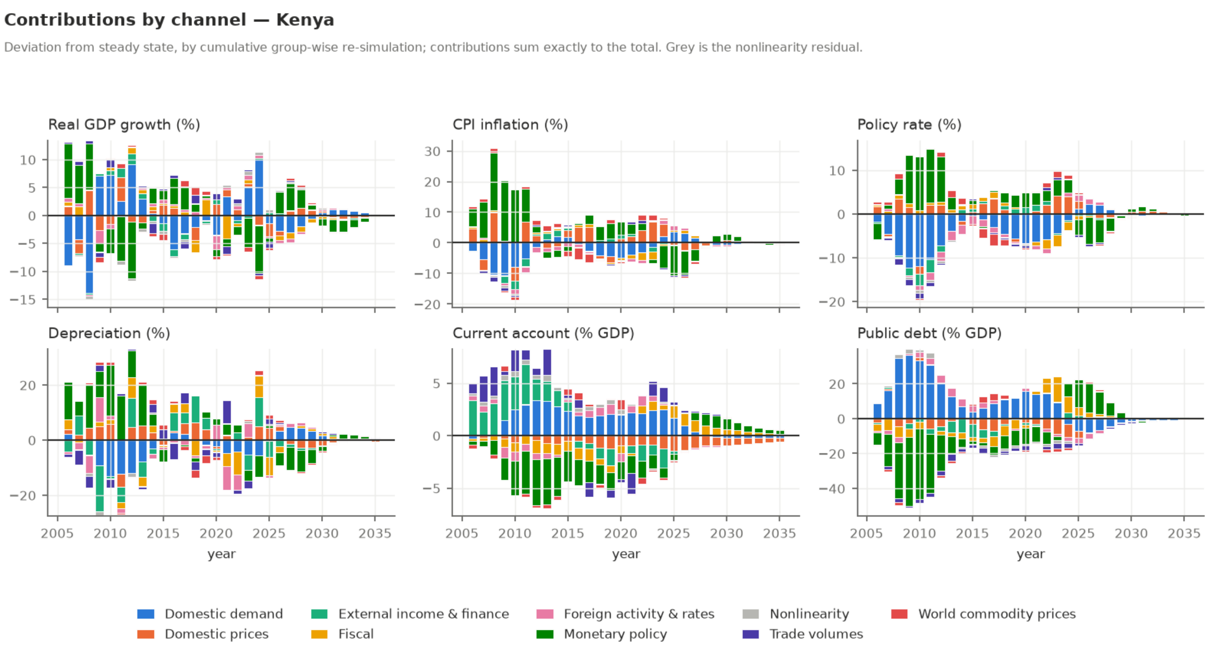

Each panel is one variable; each bar is one year, split into coloured segments, one per channel. Segments above the axis push the variable up that year and segments below push it down, so what the reader should take from a bar is the balance between the two, not the height of any one segment. Grey is the nonlinearity residual rather than a channel. The ordering is fixed and the method order-dependent.

Figure 31: Contributions by channel — South Africa

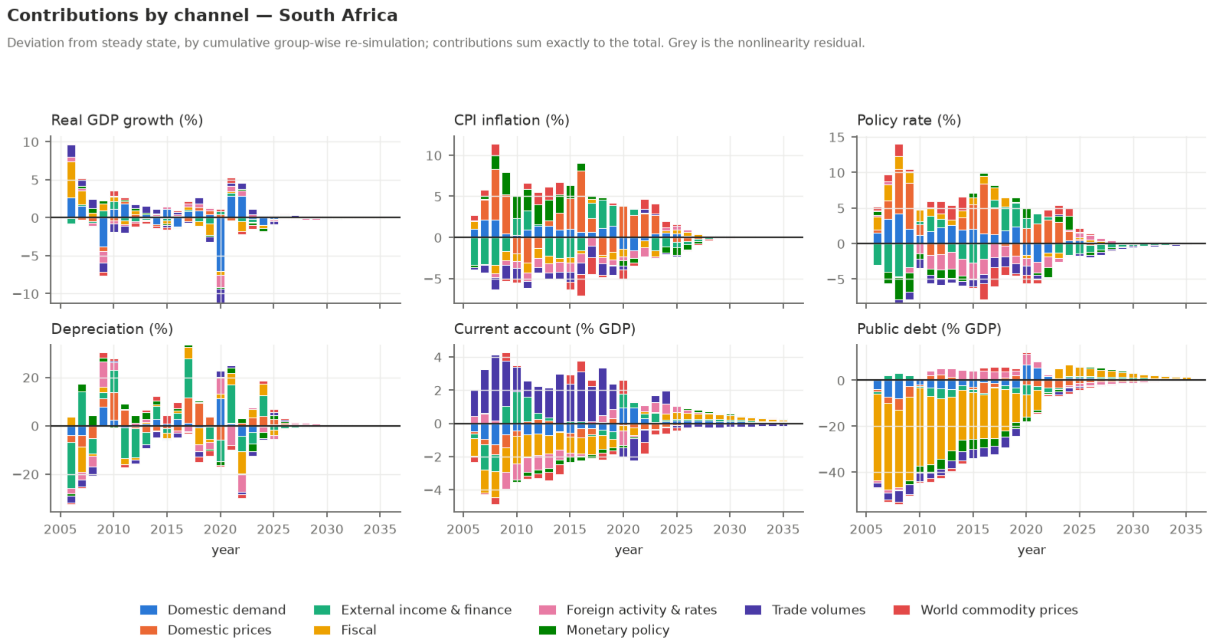

Each panel is one variable; each bar is one year, split into coloured segments, one per channel. Segments above the axis push the variable up that year and segments below push it down, so what the reader should take from a bar is the balance between the two, not the height of any one segment. Grey is the nonlinearity residual rather than a channel. The ordering is fixed and the method order-dependent.

the chart is for.

Because the nine-channel view is hard to read for exactly the first reason, the next figure folds the channels into four and draws the total being explained.

With the total drawn as a line, these panels become legible. The clearest reading is in public debt, where the policy group (orange) sits persistently below the axis through the 2010s in Uganda while the total climbs — fiscal policy holding debt down relative to what the other channels were doing, not adding to it. The current-account panel shows the mirror image: external income and finance (green) dominating the early years, then giving way.

What these charts are *for* is narrative discipline. Any account of the past decade in these economies can be checked against them, and the check is a strong one, because the contributions are constrained to add to the observed deviation. An explanation that requires a channel the decomposition gives no weight to is an explanation the data do not support.

This last figure asks a question the per-country charts cannot: when the same channel moves in all three economies at once, it is unlikely to be a domestic policy choice. The external income and finance panel is the one where the three lines move together most visibly, which is what one would expect of a channel driven by world conditions; the fiscal panel is where they diverge most, which is what one would expect of a channel each government controls separately.

Figure 32: Domestic, external and policy contributions — Uganda

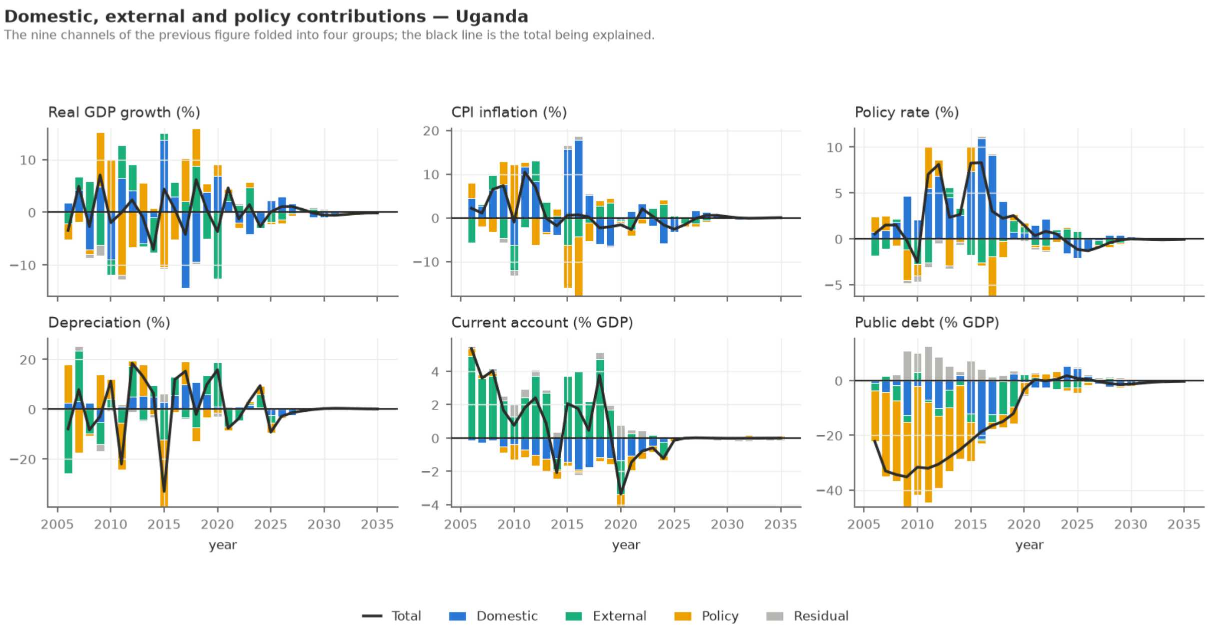

The nine channels of the previous figure folded into four groups, a strict partition so the two figures are consistent by construction. The black line is the total deviation from steady state being explained; the bars are what the model attributes it to.

Figure 33: Domestic, external and policy contributions — Kenya

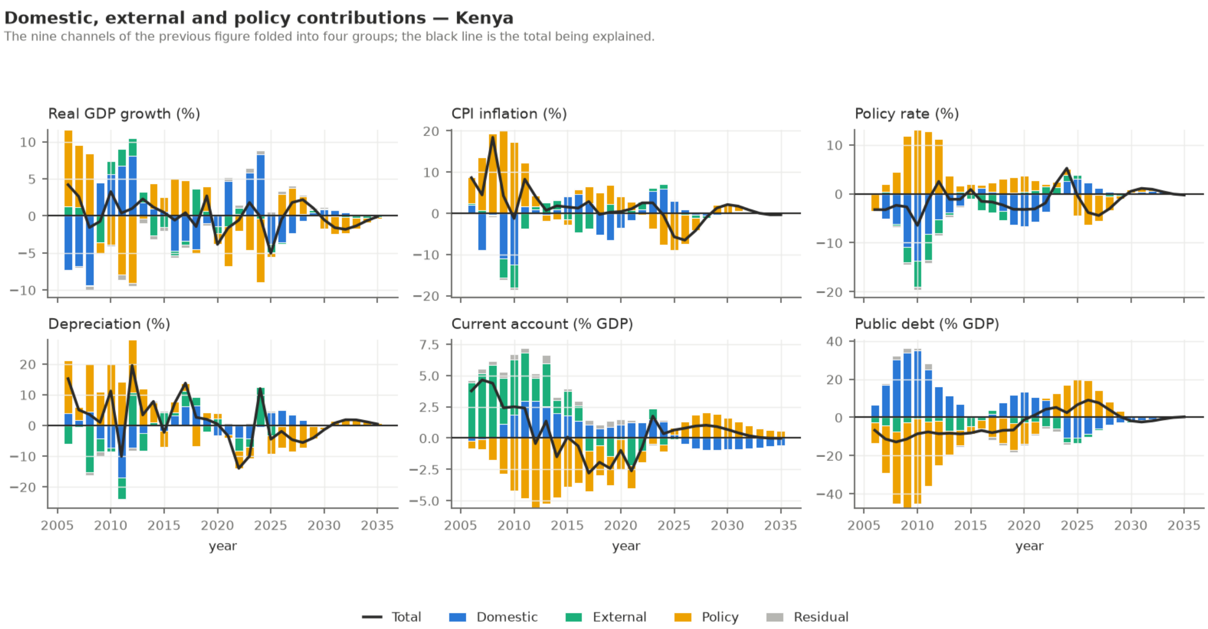

The nine channels of the previous figure folded into four groups, a strict partition so the two figures are consistent by construction. The black line is the total deviation from steady state being explained; the bars are what the model attributes it to.

Figure 34: Domestic, external and policy contributions — South Africa

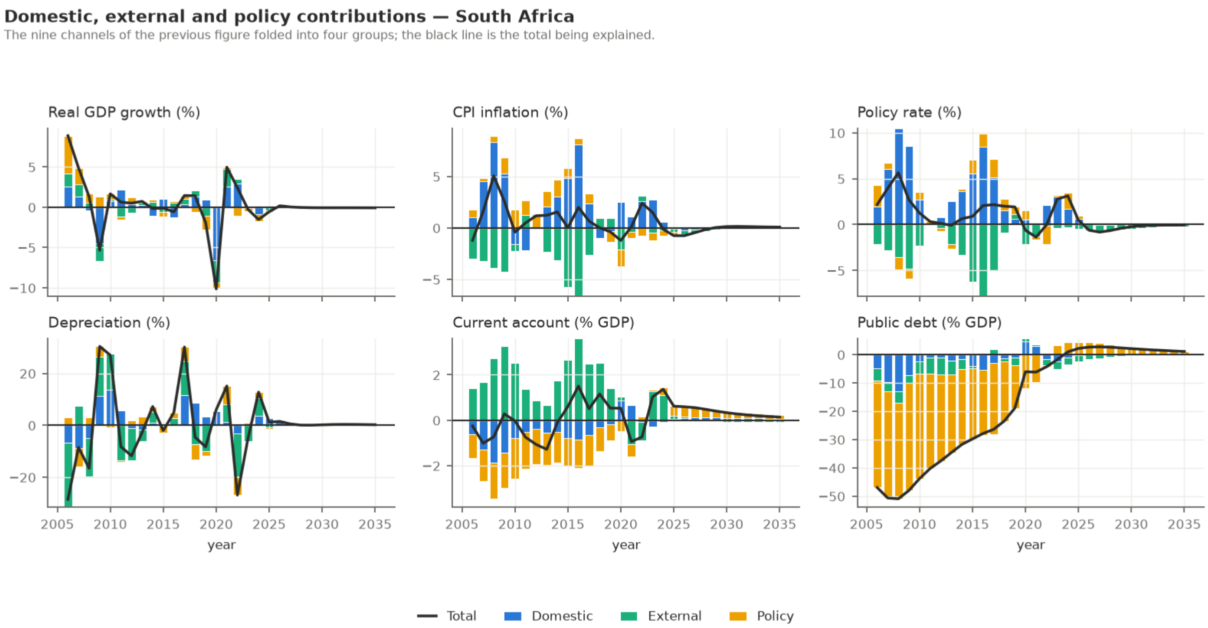

The nine channels of the previous figure folded into four groups, a strict partition so the two figures are consistent by construction. The black line is the total deviation from steady state being explained; the bars are what the model attributes it to.

Figure 35: One channel at a time, across the three economies

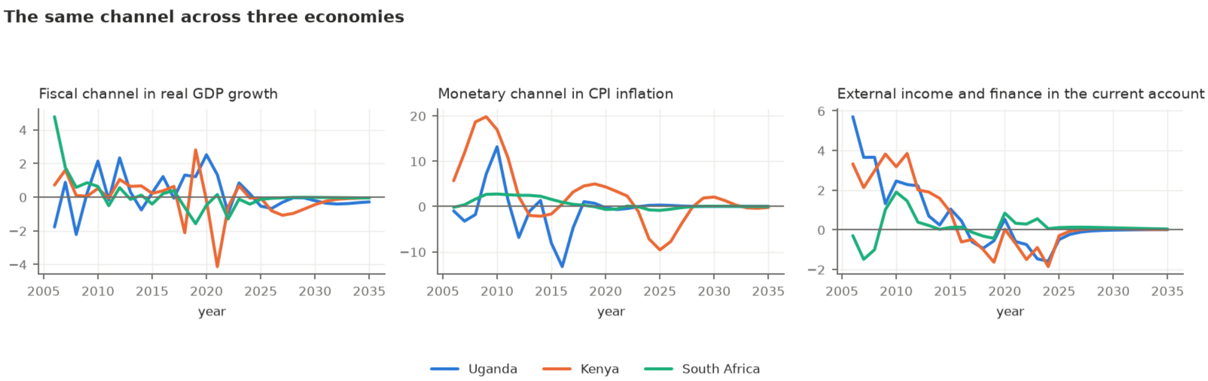

The same channel from the same decomposition, plotted for all three countries on one axis, so that a common external force can be told apart from a country-specific policy choice. One line per country, not one per channel.

11.4 Where the model runs out: the oil block

The moment-fit chart in Section 5.7 shows Kenya well below the other two, and the decompositions above give no hint why, because the weakness is not diffuse. The oil moments — 13 of 168 conditions — carry 57 per cent of Kenya's total absolute misfit, against 26 per cent for Uganda and 37 per cent for South Africa. In the worst cases the model produces essentially no variance where the data show a great deal.

The literature explains why a smooth gap–trend representation should be expected to fail here. What passes from world crude to the domestic economy depends on whether fuel is administered or market-priced, on the fiscal space available to absorb the difference, and on the state of reserves: governments hold domestic prices flat through subsidies and tax adjustments until fiscal capacity runs out, at which point prices move sharply, and the adjustment is asymmetric because those same structures make domestic prices rigid downward after favourable shocks (Nwosa, 2019). Even in South Africa, where pricing is more market-determined, pass-through is found to depend on the prevailing inflation environment rather than being a stable coefficient (Kabundi and Mlachila, 2025).

A block whose oil imports are a mean-reverting gap driven by a relative price cannot generate a series that is administratively flat and then discontinuous. This is a limitation of the structure, not of the estimation or the data, and it is the most valuable direction for extending FINEX for African use. It is also why Kenya's headline moment correlation of 0.191 should be read as a statement about the oil block rather than about the model's grip on Kenya's domestic economy: re-estimating with oil held out raises the fit over the remaining moments to 0.532, but the full-sample parameters already scored 0.528 on those same moments, so almost none of the apparent improvement is a better estimate.

11.5 What the model is good for, and what it is not

The charts above support four kinds of question and not a fifth.

Instrument choice. Because intervention, the reserve target, the inflow tax and the policy rate act through separate equations, the model returns materially different profiles for each, as the exchange-rate panels show. It answers which instrument acts where and over what horizon.

Composition rather than aggregate. The fiscal and tax panels distinguish a deficit from the spending or tax change producing it, so "consolidate by two points of GDP" becomes a question with several different answers — which is the asymmetry the empirical literature also finds (International Monetary Fund, 2026).

The investment–growth nexus. Public capital enters potential output in equation (7), so a capital programme has a supply-side consequence and not only a demand one. Its coefficient is smaller than private capital's, and that is defensible rather than arbitrary: Fund work finds average public-investment inefficiency of around thirty per cent in developing countries, with the most efficient investors obtaining roughly twice the growth return per unit invested (International Monetary Fund, 2015). The parameter is best read as an efficiency-adjusted return. It also marks what the model is silent on: FINEX represents the quantity of public investment, not its quality, so a reform raising execution standards can only be represented by changing the coefficient — that is, by assumption.

Narrative discipline. The decompositions force an account of the past to add up, which is a stronger constraint on a story than it sounds.

What the charts do not support is a quantitative trade-off between instruments whose scales were never identified. Of 83 shocks only 19 have an estimated standard deviation, and none of them are monetary or exchange-rate shocks. The appropriate inference from a weakly identified parameter is that the data are uninformative about it, not that its value is zero, and the small-sample behaviour of this estimator counsels conservative inference throughout (Ruge-Murcia, 2012).

12. A combined policy scenario

Single-shock responses are informative but not how policy questions arrive. A realistic question combines instruments at different dates and asks what the joint path looks like.

Scenario. A monetary tightening in 2027 (one standard deviation, dissipating the following year) followed by a fiscal consolidation in 2030 (one standard deviation, gone by 2031). Both unanticipated. Sign conventions follow the equations: a positive monetary residual raises the policy rate; a negative fiscal residual lowers the structural deficit.

Figure 36: Combined policy scenario, with attribution

Notes: All series are deviations from the no-shock baseline. Solid orange is the combined scenario; dashed lines are each shock run in isolation; the grey line is the interaction term — what the combined path does that the two shocks separately do not. Shaded vertical bands mark the shock years, 2027 and 2030.

Table 17: Combined scenario: deviation from baseline

	2027	2028	2030	2031	2035
Real GDP growth, pp	**−1.19**	−0.13	−0.04	+0.80	+0.01
Output gap, pp	−1.08	−1.16	−0.76	+0.07	−0.03
CPI inflation, pp	**−1.10**	−1.04	+0.10	+0.25	−0.03
Short-term rate, pp	+0.20	−0.09	−0.09	+0.03	−0.01
Depreciation vs USD, pp	**−2.03**	−0.09	+0.02	−0.02	+0.01
Current account, % GDP	+0.06	+0.05	+0.14	+0.03	+0.01
Fiscal deficit, % GDP	+0.30	+0.31	**−0.60**	−0.10	+0.01
Public debt, % GDP	+0.75	+1.42	+0.71	−0.00	+0.06

Notes: Deviations from the no-shock baseline of Section 6. Negative depreciation denotes appreciation. Both shocks are unanticipated and last one year.

The monetary tightening (2027). Growth falls 1.19 pp, inflation 1.10 pp, and the shilling appreciates 2.0%. Note that the deficit *widens* 0.30 pp of GDP and debt rises 1.4 points by 2028 — the tightening is fiscally costly, through lower nominal GDP in the denominator and weaker revenue. *A disinflation in this model is paid for partly out of the debt ratio.*

The fiscal consolidation (2030). The deficit falls 0.60 pp of GDP and by 2031 debt is back to baseline. The output cost is real but brief.

Together. By 2031 growth is 0.80 pp *above* baseline — the recovery from both shocks arriving together. By 2035 every variable is within 0.06 of baseline: the model returns to its steady state and the shocks leave no permanent mark.

12.1 The interaction term

The scenario was run three ways — each shock alone, and both together — so that

$$\text{combined} = \text{monetary alone} + \text{fiscal alone} + \text{interaction} \tag{36}$$

The largest interaction term across all observables and all years is 0.030, the flat grey line in Figure 36.

This matters more than it looks. It says that over this range Uganda's responses to monetary and fiscal shocks are **very nearly additive**: the effective lower bound does not bind and the nonlinearities in debt accumulation are second-order at this shock size. Policy analysis at this scale can therefore treat instruments separately and add them up.

That conclusion is *conditional on shock size*. Push shocks past roughly two standard deviations, or start from a position where the policy rate is near its lower bound, and the interaction stops being negligible. This is why it is computed and reported explicitly rather than assumed away — it is the diagnostic that says whether additivity is safe, and must be re-checked for each scenario rather than inherited from this one.

13. Shock decomposition

13.1 Method

The historical inversion recovers the structural shocks that reproduce Uganda's history exactly. The decomposition attributes each observable's deviation from steady state to those shocks, grouped into eight channels.

The method is group-wise and exact. Groups are switched on cumulatively and the model re-simulated each time; each group's contribution is the change in the path when it is added. FINEX is nonlinear — debt accumulation is a product of stocks, rates and the exchange rate, and the policy rule has a lower bound — so this is the only way to make contributions sum to the total. **They do, to 2×10^{-12}.** The cost is order dependence, so the order is fixed and reported: domestic demand, domestic prices, monetary policy, trade volumes, external income and finance, fiscal, world commodity prices, foreign activity and rates.

Figure 37: Historical shock decomposition

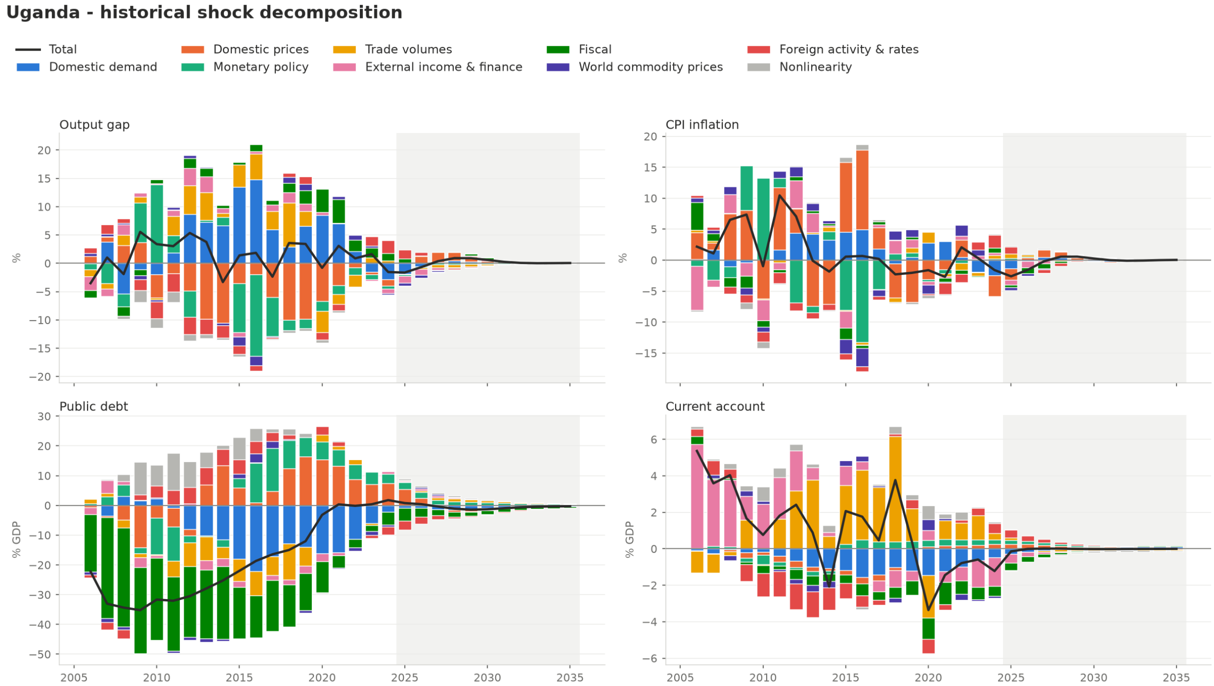

Notes: Bars are group contributions to each observable's deviation from its steady state; the black line is the total. The shaded region is the projection period. *Nonlinearity* is a residual — what the grouped contributions failed to explain — not a channel, hence its neutral grey. Contributions are large relative to the total and heavily offsetting: that is a real property of the exact cumulative method, and means individual bars are more uncertain than the total and that the fixed ordering matters.

Figure 38: South Africa: historical shock decomposition

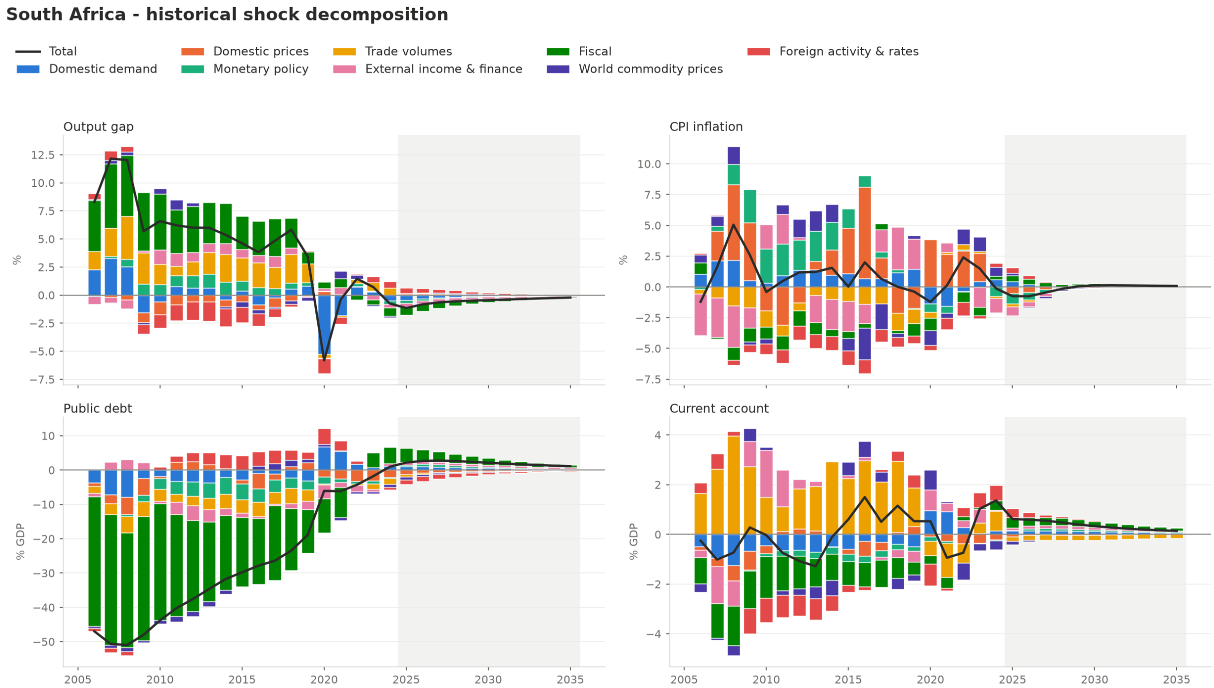

Notes: As Figure 37, for South Africa. The 2020 output-gap collapse is far deeper than Uganda's, and the debt panel shows the sustained fiscal contribution that carried the ratio from under 30% to over 75% across the sample.

Figure 39: Kenya: historical shock decomposition

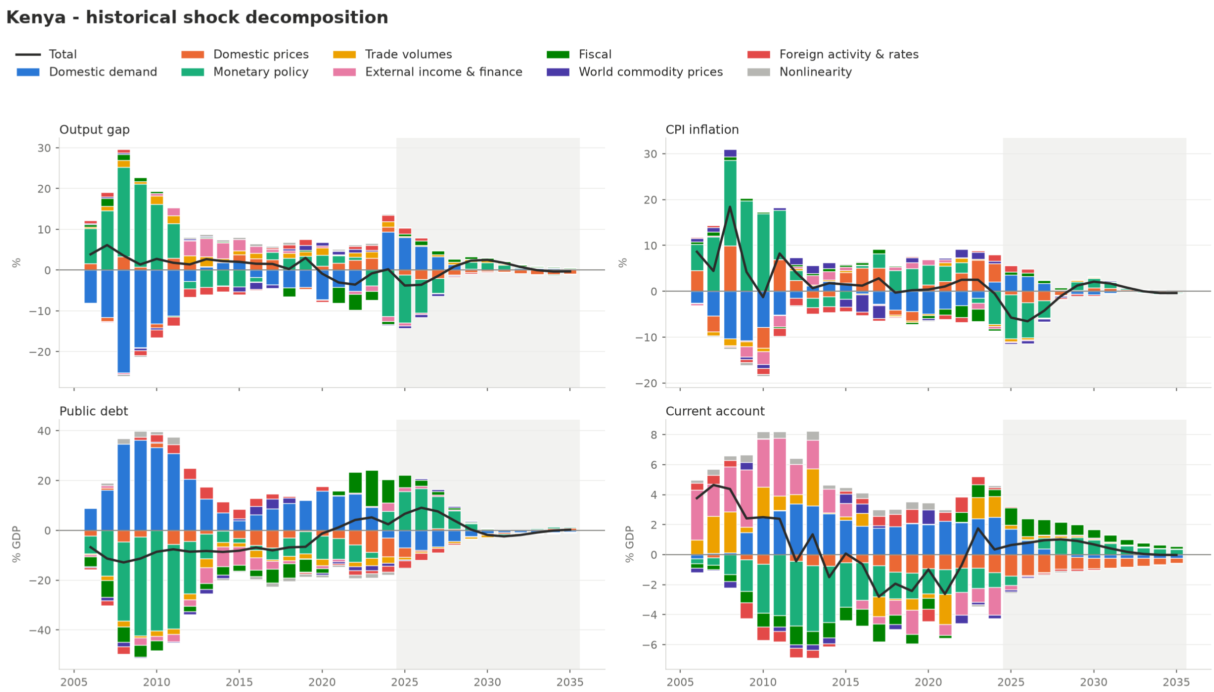

Notes: As Figure 37, for Kenya. Read the world commodity-price contribution here with the caveat of Section 18.3: the estimator put no variance on the world oil price, so that channel is attributed almost entirely to other groups.

13.2 What it says about Uganda's history

Table 18: Decomposition of selected episodes

Episode	Driver
Output gap 2020 (-0.7)	Fiscal support $(+2.7)$ offsetting collapsing trade volumes (-2.5) and foreign demand (-1.3)
Output gap 2024 (-1.7)	Weak domestic demand (-3.6), part offset by domestic prices $(+1.7)$
Inflation 2022 $(+2.1)$	World commodity prices $(+1.5)$ and domestic cost-push $(+1.4)$
Public debt 2024 $(+1.8$ above anchor$)$	Domestic demand $(+5.5)$ net of fiscal consolidation (-4.1)
Current account 2024 (-3.3)	External income and finance (-1.8) — weak remittances and other income

Notes: Contributions in percentage points (output gap, inflation) or percent of GDP (debt, current account), measured as deviations from steady state.

The 2020 pandemic year is the clearest result. Fiscal support contributed $+2.7$ pp to the output gap against a -2.5 pp collapse in trade volumes and -1.3 pp from foreign demand. The net gap of -0.7 understates both the shock and the response by a factor of four. This is exactly what a decomposition is for.

Debt is a fiscal-versus-demand story. The 1.8-point overshoot in 2024 decomposes into $+5.5$ from domestic demand — which raised spending and the deficit — against -4.1 from active consolidation. The consolidation was doing most of what was asked of it; demand pressure was simply larger.

13.3 What carries the projection

The decomposition runs through 2035, so it also says what drives the baseline. The 2025 growth pickup is domestic prices $(+1.8)$ and domestic demand $(+1.0)$, and every group fades within two years. After roughly 2028 the bars are visually empty: the projection is not being carried by anything, which is the correct behaviour for a no-shock forecast converging on its steady state.

14. Fiscal policy analysis in depth

Fiscal policy is where FINEX's disaggregation earns its cost, so this section treats it more fully than the single impulse response of Section 9.2. Four questions are addressed: how large the multiplier is in each of the three economies, why it is small, whether the instrument matters, and what debt-financed public investment does in this model.

Estimating three countries rather than one changes the answer to the second question. Read on Uganda alone, the small multiplier had two natural explanations — demand leaking into imports, and the central bank offsetting the impulse. Neither survives being asked of three economies at once, and saying so is more useful than repeating an explanation that happens to fit one case.

14.1 How large is the multiplier?

Table 19: The four estimated parameters that govern the fiscal result

Parameter	Uganda	Kenya	S. Africa	What it governs
Import-gap persistence	**0.004**	**0.767**	**0.059**	How long a demand impulse stays in domestic output before leaking abroad
IS real-rate elasticity	**0.582**	**0.590**	**0.021**	How strongly a rate change reaches consumption, hence the monetary offset
Structural-deficit persistence	**0.082**	**0.245**	**0.134**	How fast the fiscal rule pulls the deficit back to target
Inflation persistence	**0.634**	**0.909**	**0.903**	How much of a demand impulse shows up as inflation rather than output

Notes: All four are searched by the Simulated Method of Moments rather than calibrated, on 24 annual observations per country. The generic calibration each starts from is the same in all three: 0.15 for import persistence, 0.10 for the IS elasticity, 0.30 for deficit persistence and 0.80 for inflation persistence, so every difference below is something the data produced, not something assumed.

Table 20: The fiscal impulse, per standard deviation of the shock

	Uganda	Kenya	S. Africa
Shock size (% of GDP)	0.81	2.00[†]	2.00[†]
Impact effect on GDP growth (pp)	+0.62	+0.66	+0.57
Effect one year later (pp)	-0.55	-0.55	-0.57
Cumulative: level of GDP after 5 years (%)	**+0.05**	**+0.02**	**-0.02**
Peak policy-rate response (pp)	+0.01	-0.07	+0.14
Peak private consumption (pp)	+0.08	+0.12	+0.08
Peak private investment (pp)	-0.10	+0.13	-0.05

Notes: The first row is the shock scale itself; every response below it is divided by the country's own shock standard deviation, which is what makes the three columns comparable; the unscaled responses are not, because two of the three scales are assumptions. A dagger marks a standard deviation that reverted to its calibrated value because the estimation could not identify it. The cumulative row sums the growth responses over the impulse year and the five that follow, which is the effect on the *level* of GDP.

Figure 40 draws the same result.

Figure 40: The fiscal impulse in three economies: growth, the level of GDP, and the policy rate

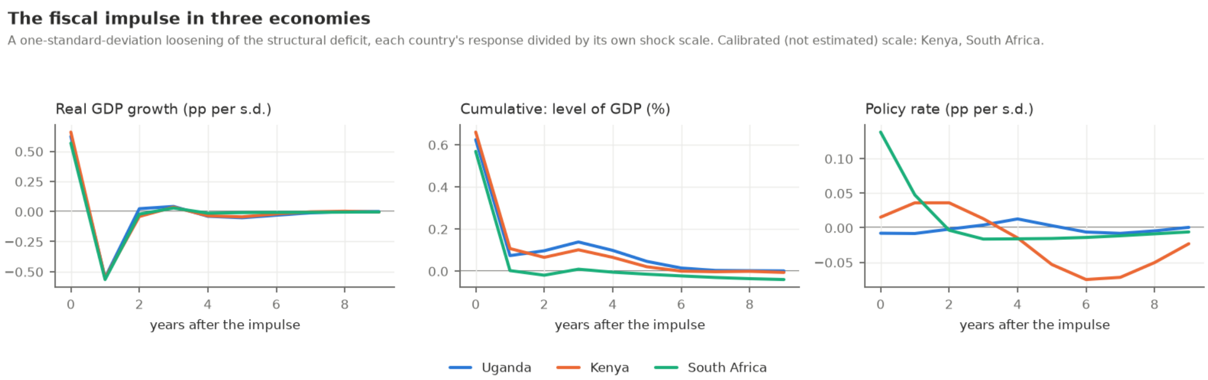

A one-standard-deviation loosening of the structural deficit, with each country's response divided by its own shock scale so that the three are comparable. The left panel is the growth rate, the middle panel its running sum — the effect on the level of GDP — and the right panel the policy rate. Kenya's and South Africa's shock scales are calibrated rather than estimated, so the *shape* of their responses is estimated and the *size* is an assumption; dividing by the scale removes that assumption from the comparison but not from the level.

The impact multiplier is near 0.6 in all three economies: 0.62 percentage points of growth per standard deviation in Uganda, 0.66 in Kenya and 0.57 in South Africa. The agreement is the finding. Three economies of different size, different openness and different monetary regimes, estimated separately on their own data, return an impact multiplier within a tenth of each other, and one that sits squarely in the range the empirical literature reports for low-income and emerging economies.

The cumulative effect is not near 0.6. It is near zero. Summing the growth responses over the impulse year and the five that follow gives +0.05 % on the level of GDP for Uganda, +0.02 for Kenya and -0.02 for South Africa. In every case the growth impulse of the first year is given back almost exactly in the second: +0.62 then -0.55 in Uganda, +0.66 then -0.55 in Kenya, +0.57 then -0.57 in

South Africa.

This is the single most important thing the fiscal block says, and it is easy to misreport. A temporary deficit impulse raises the *level* of output for about a year and then the level returns; a growth rate that rises and falls by the same amount describes exactly that. **Quoting the impact multiplier without the reversal overstates what a one-off fiscal expansion achieves by roughly an order of magnitude.** It is not that fiscal policy does nothing — a sustained change in the deficit is a different experiment, and one the scenario engine can run — but that a one-year impulse buys one year of activity and no more.

14.2 Why the multiplier is small: two explanations that do not survive

14.2.1 Leakage into imports

Non-oil imports respond to every domestic demand component separately (equation (15)), so a demand impulse that leaks abroad within the year never accumulates into domestic output. In equation (3) the import term enters negatively with weight $\overline{M}^{NOIL,Y}/100 \approx 0.18$, which is large enough for the channel to matter.

The estimates say the channel is real but that it is not what makes the multiplier small. **Import persistence is estimated at 0.004 in Uganda and 0.767 in Kenya — a factor of 207 — and the two impact multipliers are 0.62 and 0.66.** If fast import adjustment were the reason Uganda's multiplier is small, Kenya's, where imports adjust slowly, would be visibly larger. It is not. South Africa, at 0.059, sits close to Uganda and returns 0.57.

The import channel therefore shapes *where* the demand goes — Kenya's non-oil import ratio rises +0.11 points of GDP against Uganda's +0.03 — without changing how much output it buys. This is worth stating plainly because the single-country version of this section named import persistence as the most consequential estimate in the fiscal results, and the cross-country evidence does not support that.

14.2.2 Monetary offset

The second natural explanation is that the central bank leans against the impulse: the fiscal expansion raises the output gap, the Taylor rule (equation (28)) raises the policy rate, and the rate rise reaches consumption through the IS elasticity.

The rates do not move enough for this to be the explanation either. The largest policy-rate response over ten years is +0.01 percentage points per standard deviation in Uganda, -0.07 in Kenya and +0.14 in South Africa. Against IS elasticities of 0.58, 0.59 and 0.02, the implied effect on consumption is smaller than a rounding error in the growth response. In Uganda the rate initially *falls*.

The reason is visible in the inflation persistence estimates of Table 19: at 0.63 for Uganda and around 0.91 for the other two, inflation in this model is driven far more by its own past than by the current output gap, so a one-year demand impulse barely disturbs it — and a rule that responds mainly to inflation therefore barely responds at all.

14.2.3 What does explain it

The explanation that survives is the one the reversal already gives: the impulse is temporary, and nothing in the model converts a temporary deficit into a permanent change in the level of output. The structural deficit itself mean-reverts — persistence is estimated at 0.082 for Uganda, 0.245 for Kenya and 0.134 for South Africa, so in all three the fiscal rule pulls the deficit back towards target within a few years without any discretionary decision. Private demand does not take over: consumption responds +0.08, +0.12 and +0.08 points on impact, and investment -0.10, +0.13 and -0.05. There is no crowding in to sustain the expansion once the impulse ends.

There is also a slower channel, which matters for the level of debt rather than for this experiment. Higher debt raises the interest premium through the exponential risk-appetite term of equation (25), which raises real rates, which reduces investment and hence potential output (Section 3.5). At the debt levels of all three countries this term is in its flat region and the effect is small within ten years; at 85 % of GDP it would not be.

14.3 Does the instrument matter?

FINEX distinguishes fiscal instruments rather than treating "fiscal policy" as a single lever, and they act through different equations:

Table 21: Fiscal instruments and their transmission channels

Instrument	Enters through	Consequence
Government purchases	Demand identity directly, equation (3)	Immediate output effect; import leakage; no direct private-sector response
Public investment	Demand identity *and* public capital stock, equation (18)	Demand effect now, potential-output effect after three years
Transfers	Consumption equation, coefficient c_7	Works through household demand; subject to the same import leakage
Consumption tax	Consumption equation, intertemporal term c_5	Anticipated changes shift demand across time
Import duties	Import equation, term c_8	Affects the composition of demand, not only its level
Direct taxes	Investment trend, equation (5)	Level effect on the investment-to-GDP ratio

Notes: Each instrument is a distinct shock in the model with its own estimated or calibrated variance (Appendix I). Multipliers therefore differ by instrument as a result of the model's structure rather than by assumption. The table is the same for all three countries: the transmission channels are the model's, not a country's.

Two consequences follow that a single-multiplier model cannot express.

Composition matters as much as size. A consolidation of a given size has different consequences depending on whether it falls on purchases, transfers or public investment. The first two are cyclical; the third has a permanent component. This also qualifies the reversal result above: it was computed on the *aggregate* deficit, and public investment is precisely the instrument for which the reversal would not be expected to hold.

Timing matters for taxes. Because the consumption equation contains $(\tau_t^C - \tau_{t+1}^C)$, an *announced future* tax rise raises consumption today. A pre-announced value-added tax increase is expansionary before it is contractionary — a genuinely useful prediction, and one that follows from the structure rather than being imposed.

14.4 Debt-financed public investment

This is the case where FINEX's investment–growth nexus does work no conventional gap model can do, and where the answer is genuinely ambiguous rather than assumed.

Public investment financed by borrowing sets two forces against each other:

- **It raises potential output.** Through equations (16)–(18), public capital enters the production function. The effect is delayed — public investment becomes fully productive only after three years — but permanent.
- **It raises the cost of capital.** Higher debt raises the interest premium (equation (25)), which raises real interest rates, which depresses *private* investment (equation (13)) and hence private capital accumulation.

Which dominates is a quantitative question, and its answer depends on the starting debt level. Because the risk-appetite term is exponential, the drag is mild at low debt and severe at high debt. The same public investment programme can be growth-enhancing for a country at 40 % of GDP and self-defeating for one at 90 %, with no change in the productivity of the spending itself.

The three countries sit at different points on that curve, which is the clearest practical use of estimating them together. Uganda's debt is near its 50 % anchor and its external borrowing is largely concessional — the external premium is calibrated at zero — so the risk term is in its flat region and additional external debt is unusually cheap. Kenya and South Africa carry higher debt and borrow on market terms, so the same programme meets a steeper premium. Nothing in the estimation determines where the curve turns; the exponential form is imposed and its parameters are calibrated, so this is the model's structure speaking rather than the data.

This report does not run that experiment. It requires a scenario on public investment specifically rather than on the aggregate deficit, and the public-investment shock's standard deviation is among those the estimation does not identify in any of the three countries. Both are straightforward to add, and the result would be more useful than the aggregate multiplier reported above — but it would be a scenario built on a calibrated scale, and should be labelled as such.

14.5 What the fiscal results do and do not establish

Established. The impact multiplier on growth is near 0.6 in all three economies, estimated separately, and the cumulative effect on the level of output is close to zero in all three because the first year's impulse is given back in the second. The agreement across three different economies is stronger evidence than any one of them, and it is the reversal rather than the impact figure that should be quoted. The debt path in each baseline is stabilising without a large consolidation.

Established negatively. Import leakage and monetary offset do not explain the small cumulative multiplier. Import persistence differs by a factor of 207 between Uganda and Kenya with no visible consequence for the multiplier, and no country's policy rate moves by more than 0.14 percentage points per standard deviation. A reader who has seen either explanation offered for a small multiplier should ask for the cross-country check before accepting it.

Not established. That fiscal policy is ineffective in any of the three. The result is conditional on an aggregate deficit shock rather than a specific instrument; on a *temporary* impulse rather than a sustained change in the fiscal stance; on parameters estimated from 24 annual observations; and, for Kenya and South Africa, on a shock scale the estimation could not identify. Public investment in particular is excluded from the reversal result by construction, since its effect runs through the capital stock rather than through demand. Each of these is a reason the true effect of a fiscal programme could be larger than the number above, and none of them is a reason it could be smaller.

15. Macroeconomic history through the model

The inversion of Section 6.2 recovers the structural shocks that reproduce each country's history exactly, and the decomposition of Section 13 attributes each observable's path to them. Together they let the model tell a coherent story about the last two decades — which is, the working paper argues, the main thing a model of this kind is for (Berg et al., 2023). The narrative below is the model's, not the authors': each attribution is a number in Appendix M.

15.1 2006–2011: the commodity and credit boom

The period opens with public debt at roughly 22% of GDP — the post-HIPC, post-MDRI position, some 28 points below the model's 50% anchor — and a steadily positive contribution from domestic demand to the output gap.

The model reads the 2011 inflation episode, when headline inflation reached double digits, as predominantly a **domestic price** event with a **world commodity price** contribution on top. That is a meaningful distinction for policy: the domestic component is what the Taylor rule can act on, while the commodity component works through relative prices whose trend the model treats as largely exogenous. A monetary response calibrated to the headline number would have been acting substantially on a shock outside its reach.

15.2 2012–2016: consolidation and external drag

Debt rises steadily through this period, and the decomposition attributes the rise principally to **domestic demand** — higher spending and a wider deficit — rather than to interest costs or valuation effects. The **fiscal** channel contributes negatively throughout, which is the fiscal rule's response to a debt ratio moving away from its anchor: consolidation was active and continuous, and was simply smaller than the demand pressure it was offsetting.

Trade volumes contribute negatively over much of the period, consistent with weak external demand for non-resource exports, while remittances hold up.

15.3 2020: the pandemic

This is the clearest single result the decomposition produces, and the most useful illustration of why decomposition is worth doing at all.

The output gap in 2020 was -0.7 — a modest number that, taken alone, would suggest a mild disturbance. The decomposition shows it is a small residual of large offsetting forces:

- **Fiscal support:** $+2.7$ **percentage points.** The single largest positive contribution to the output gap in the sample.
- **Trade volumes:** -2.5 **percentage points.** The collapse in exports and the disruption to import-dependent production.
- **Foreign demand:** -1.3 **percentage points.** The recession in Uganda's trading partners.

The net gap understates both the shock and the response by roughly a factor of four. A reader looking only at the realised output gap would conclude 2020 was unremarkable for Uganda. The model's reading is that it was a very large negative external shock met by a very large fiscal response, and that the two approximately cancelled in the aggregate while leaving quite different footprints on the debt ratio and the external accounts.

15.4 2022: the global inflation shock

Inflation in 2022 was $+2.1$ above steady state, decomposed into **world commodity prices** $(+1.5)$ and **domestic cost-push** $(+1.4)$, partly offset elsewhere.

The near-equal split matters. It says roughly half of Uganda's 2022 inflation was imported, transmitted through the import-price terms of the Phillips curve (equation (8), coefficients γ_2 and γ_3), and roughly half was domestic. Uganda's estimated import-price pass-through of 0.30 — above South Africa's 0.23 — is what makes the imported share so large in an economy of this size.

15.5 2024: the position the forecast starts from

The projection's initial conditions are worth stating explicitly, since every number in Section 6 descends from them.

Table 22: The 2024 starting position and its attribution

Observable	2024	Model attribution
Output gap, %	-1.7	Weak domestic demand (-3.6), part offset by domestic prices ($+1.7$)
Public debt, % GDP	51.8	Domestic demand ($+5.5$) net of fiscal consolidation (-4.1)
Current account, % GDP	-9.2	External income and finance (-1.8) — weak remittances and other income
CPI inflation, %	3.3	Below the 4.9% anchor, consistent with the negative gap

Notes: Attributions are contributions to the deviation from steady state, from the cumulative group-wise decomposition. The 2024 output gap of -1.7 reported here is the model's inverted value; the observed series (Appendix (Appendix H) shows $+0.1$ for 2024, the difference being the model's own estimate of the trend against the HP-filtered observable.

Two features of this starting position drive the forecast. The **negative output gap** is why the projection opens with a further contraction in 2025 before recovering. The **debt overshoot of 1.8 points** is why the fiscal rule calls for a modest consolidation that carries debt down to 48.8% by 2030.

15.6 What the decomposition says about the projection

The decomposition runs through 2035, so it also reports what is carrying the forecast rather than only the history. The answer is: less and less.

The 2025 growth pickup is attributed to **domestic prices** ($+1.8$) and **domestic demand** ($+1.0$), and every group's contribution fades within two years. After roughly 2028 the contribution bars are visually empty.

This is the correct behaviour for a no-shock projection converging on its steady state, and it is a useful diagnostic in its own right: **if a group's contribution were still large at the end of the projection, the model would not have converged**, and the forecast would not be a forecast. The emptiness of the late bars in Figure 37 is the visual counterpart of the convergence statistic in Section 6.5.

15.7 The same exercise for Kenya and South Africa

The narrative above is Uganda's because Uganda is the case the model fits best and therefore the one whose attribution carries most weight. The decomposition is computed on the same basis for the other two (Figures 30 and 31), and two contrasts are worth drawing out.

Kenya's recent growth is attributed differently from Uganda's. In 2024 the largest contributions to Kenyan growth are Domestic demand (+10.0), Monetary policy (-8.5), Domestic prices (-1.7), against Domestic demand (-1.6), World commodity prices (-0.7), Foreign activity & rates (+0.7) for Uganda. The channels the model leans on are not the same, which is what one would expect of two economies whose estimated interest and inflation blocks differ as much as Section 5.6 showed.

South Africa's 2024 inflation is attributed to external and policy channels rather than to domestic ones. The decomposition gives External income & finance (-1.1), Monetary policy (-0.8), World commodity prices (+0.7), with the external and monetary terms pulling the same way and world commodity prices partly offsetting them — a different balance from either East African case, where domestic demand carries the largest single contribution.

Both readings carry the caveat of Section 13: contributions are large and offsetting, the ordering is fixed, and the method is order-dependent. They are a discipline on a narrative, not a measurement of one.

16. A scenario library

The single-shock responses of Sections 9 and 10 are diagnostic. Operational work needs named scenarios that correspond to identifiable risks, are reproducible, and can be re-run when the data or the estimate change. Seven are maintained, each specified as a TOML file and each editable as a starting point rather than fixed. **The specification below applies to all three countries**: the scenario engine is country-agnostic, the same file runs against any estimated country by changing one field, and the seven are offered as presets in the web application for each of them (Appendix G). What is country-specific is the *result*, and the worked results reported in this paper are Uganda's — so the sizes below are in standard deviations of each country's own estimated or calibrated shock, and the same nominal scenario is a different-sized event in each.

Table 23: Standard stress scenarios, defined once for all three countries

Scenario	Composition	Principal exposure
Oil price spike	World oil price $+2$ s.d., 2027–29	Inflation, fiscal cost of fuel, current account
Global slowdown	US and euro-area activity -1.5 s.d., 2027–28	Non-resource exports, investment
Tighter global financial conditions	US short rate $+1.5$ s.d. and country risk premium $+1$ s.d., 2027–28	External debt service, exchange rate, reserves
Drought / supply shock	Cost-push $+1.5$ s.d., 2027	Inflation, and the policy response to it
Remittance decline	Remittances -2 s.d., 2027–28	Household consumption, current account
Fiscal slippage	Structural deficit $+1.5$ s.d., 2027–29	Debt path, risk premium, crowding out
Exchange-rate pressure	Country risk premium $+2$ s.d., 2027	Foreign-currency debt revaluation, inflation

Notes: Shock sizes are in standard deviations of the corresponding structural shock. Where that standard deviation is calibrated rather than estimated (Appendix I) the scenario's *scale* inherits that assumption, though its *shape* does not. Each scenario is a starting point: the years, sizes and persistence are editable, and combinations are permitted.

16.1 Why these seven

The list is not generic. Each corresponds to a channel that Uganda's own estimated structure makes material, and several would be unimportant in a differently-structured economy.

- **Remittances** appear because they enter the consumption equation *directly* (equation (11), coefficient c_6) rather than only through the external accounts. At above 3% of GDP for Uganda this is a first-order consumption channel, and a remittance shock is not a scaled-down export shock.

- **Cost-push** appears because it is the largest estimated domestic shock in Uganda's data ($\sigma = 4.07$), and because its output cost operates mostly through the policy response rather than directly.
- **Risk premium** appears because its consequences are fiscal before they are real: it revalues foreign-currency debt and moves the exchange rate well before it moves the output gap.
- **Oil** appears despite a modest GDP effect, because its scale is the one the estimation demonstrably fails to identify (Section 5.7) and therefore the one most in need of explicit scenario treatment rather than reliance on estimated variances.

16.2 Reading a scenario honestly

Three disciplines apply to every scenario run in this framework, and all three are enforced in the tooling rather than left to the analyst.

Attribution is computed, not assumed. When a scenario contains more than one shock, each is also run alone and the interaction reported explicitly (Section 12.1). Where the interaction is small, the components may be read separately; where it is large — typically because the effective lower bound binds — they may not.

Anticipation is a choice with consequences. A pre-announced consolidation moves the economy before it happens, through the forward-looking terms in equations (11) and (13). An unanticipated one does not. The two are different scenarios and the specification records which is intended.

Scale governs whether linearity holds. Beyond roughly two standard deviations, or from a starting position near the lower bound, the model's nonlinearities — the bound itself, and the exponential risk-appetite terms of Section 3.6.3 — stop being negligible. Scenarios at that scale must be simulated in full rather than scaled from a stored impulse response.

17. What this model cannot do

Four limitations bound every result above. They are properties of the exercise, not defects to be worked around.

17.1 The foreign block is not identified from one country

With 24 annual observations and 40 shock variances, the non-negative least squares assigns explainable covariance to a subset of shocks and drives the collinear remainder to zero — 19 of 40 for Uganda and 25 for South Africa. It bites hardest on the world oil price, whose model variance is exactly zero against a data variance of 549 — the model declines to fit it rather than fitting it badly. The same few moments dominate the residual misses in South Africa.

The fix is structural. The foreign shocks drive exogenous processes whose data are identical for every country. Their standard deviations should be estimated once, directly off the foreign series as univariate autoregressive residuals, and held fixed while the NNLS allocates only domestic shocks. Until that is done, any stochastic exercise — fan charts especially — understates foreign-driven risk.

17.2 The external block has a level mismatch

The steady-state current account (-2.8% of GDP) sits above the sample mean (-5.4%) because the national accounts and the balance of payments disagree: $C+I+G$ in the national accounts implies a bigger trade deficit than the BoP records, and the model can satisfy only one. External *levels* in the projection carry a step at the 2024/25 handover; external *dynamics* do not.

17.3 The policy rate series ends in 2019

The monetary block is estimated on 20 rather than 24 observations, and on an operational framework that has since changed. The monetary results of Section 9.1 are the least well-anchored in this report.

17.4 The projection contains no off-model information

It knows nothing about Ugandan oil production coming on stream, and it is a no-shock conditional forecast rather than a judgemental one. It answers "where does the estimated structure take the economy from 2024", which is a narrower question than "what will happen".

18. The three economies compared

All three countries were estimated with identical code, an identical moment set and identical methodology, differing only in data and in a short country-specific calibration block. Differences in the results are therefore differences in the economies, not in their treatment.

They span a wide range. Uganda is low-income, fast-growing, aid- and remittance-supported, with a small government and largely concessional external debt. South Africa is upper-middle-income and slow-growing, with a large government, deep domestic financial markets and market-priced debt. Kenya sits between them on almost every structural ratio — a lower-middle-income economy with a moderate government, a mix of concessional and market borrowing, and growth close to Uganda's.

They do not fit equally well, and that is part of the result. Uganda and South Africa reach moment correlations of 0.783 and 0.774; Kenya reaches only 0.191. Kenya clears the identification threshold and its projection converges, so it is reported — but it is a materially weaker estimate than the other two, and Section 18.3 explains why in a way that turns out to be informative about the method rather than about Kenya.

18.1 The macroeconomic starting points

Table 24: Steady states and 2024 positions compared

	Uganda		South Africa		Kenya	
	SS	2024	SS	2024	SS	2024
Real GDP growth, %	4.9	5.9	1.3	0.5	4.7	4.6
CPI inflation, %	4.9	3.3	4.4	4.3	4.9	4.4
Short-term rate, %	9.4	—	5.2	8.3	10.1	15.3
Current account, % GDP	-2.8	-9.2	-0.6	-0.6	-2.7	-1.3
Revenue, % GDP	13.3	15.3	25.9	27.3	17.3	17.5
Expenditure, % GDP	17.4	19.7	29.8	33.0	22.0	23.2
Overall deficit, % GDP	4.1	4.4	3.9	5.7	4.7	5.7
Interest bill, % GDP	**2.9**	**3.4**	**6.5**	**5.2**	**5.6**	**5.2**
Public debt, % GDP	**50.0**	**51.8**	**75.0**	**76.0**	**60.0**	**67.3**
Private consumption, % GDP	72.2	66.1	63.5	64.6	76.2	77.3
Private investment, % GDP	18.7	17.6	14.8	12.3	15.6	11.7

Notes: Steady states are calibrated from each country's own 2010–2024 sample means plus its policy anchors: 5% inflation and a 50%-of-GDP debt anchor for Uganda, the 4.5% midpoint of the SARB's 3–6% band with a 75% anchor for South Africa, and 5% inflation with a 60% anchor for Kenya. Uganda has no 2024 short-rate observation (Section 4.2).

The interest burden orders the three countries, and not in the way the debt ratio alone would suggest. South Africa carries the heaviest bill at 6.5% of GDP on 75% debt, Kenya 5.6% on 60%, Uganda only 2.9% on 50%. Kenya's burden is nearly twice Uganda's on a debt ratio only ten points higher, because Kenya borrows substantially at market rates where Uganda borrows concessionally. *The composition of debt matters as much as its level*, and the model carries that distinction explicitly through the sovereign spread on foreign-currency borrowing.

Kenya's 2024 policy rate of 15.3% against a 10.1% steady state is the largest monetary tightening in the three samples, and its consequences run through the whole Kenyan projection.

18.2 Three economies, three transmission mechanisms

Table 25: Estimated behavioural parameters compared

Parameter	Meaning	Cal.	Uganda	S. Africa	Kenya
C3_L_CONS_GAP	IS real-rate elasticity	0.10	**0.58**	**0.02**	**0.59**
C1_L_CONS_GAP	Consumption persistence	0.25	0.45	0.46	0.51
C1_L_IMP_NOIL_GAP	Import persistence	0.15	**0.00**	**0.06**	**0.77**
C1_L_EXP_NNR_GAP	Export persistence	0.15	0.74	0.09	0.01
C1_DL_P_CONS	Inflation backward-looking	0.80	**0.63**	**0.90**	**0.91**
C2_DL_P_CONS	Import-price pass-through	0.30	0.32	0.14	0.11
C4_DL_P_CONS	Phillips slope on the gap	0.30	0.19	0.46	0.19
C1_RS	Policy-rate smoothing	0.15	**0.73**	**0.47**	**0.68**
C2_RS	Taylor weight on inflation	1.00	0.55	0.79	0.95
C3_RS	Taylor weight on the gap	0.20	0.04	0.47	0.02
C1_DEF_S_RAT	Structural-deficit persistence	0.30	0.08	0.13	0.25
C1_GEXP_CONS_RAT	Spending persistence	0.70	0.86	0.62	0.70

Notes: Bold marks parameters where the three countries diverge most. All three were estimated by the Simulated Method of Moments on the same moment set — 168 conditions each — over 2001–2024. Kenya's estimates rest on a substantially weaker overall fit (Section 18.3) and should carry correspondingly less weight.

18.2.1 Kenya isolates the South African anomaly

The single most useful thing Kenya contributes is a third reading of the interest-rate channel. **Kenya's IS real-rate elasticity is 0.59 — all but identical to Uganda's 0.58, and twenty-eight times South Africa's 0.02.**

With two countries the South African estimate was merely surprising: one could argue either that transmission is genuinely weak there, or that the estimate is contaminated. With three, the picture is sharper. Two African economies with different income levels, different financial depth and different monetary histories both estimate an elasticity near 0.5 — while the country with by far the deepest financial markets estimates essentially zero. That ordering is the opposite of what financial structure would predict.

This substantially strengthens the case made in Section 9 that South Africa's 0.02 is an identification failure rather than a finding. South Africa's policy rate is the most systematically endogenous of the three — a long-established inflation-targeting regime moving rates in reliable response to the cycle — which is precisely the condition under which a moment estimator without an instrument will attribute the covariance between rates and consumption to the rule rather than to the elasticity. Kenya, with a less rule-bound policy history, does not suffer it to the same degree.

18.2.2 Inflation: Kenya sides with South Africa

Kenyan inflation is as backward-looking as South Africa's (0.91 against 0.90) and both are far more inertial than Uganda's 0.63. Kenya's Phillips curve is as flat as Uganda's (0.19 against 0.19) where South Africa's is more than twice as steep (0.46). Kenya has the lowest import-price pass-through of the three (0.11 against Uganda's 0.32).

Kenya is therefore neither a larger Uganda nor a poorer South Africa: it combines South African inflation persistence with Ugandan Phillips-curve flatness and weaker external pass-through than either. Disinflation in this estimated structure is slow in Kenya — inertia is high — and the output gap is a weak lever for achieving it.

18.2.3 The policy rules order cleanly

Smoothing orders cleanly: Uganda 0.73, Kenya 0.68, South Africa 0.47. The Bank of Uganda is estimated to move rates in the smallest steps and the Reserve Bank in the largest, which matches what is known about the two regimes.

The weight on inflation does not order cleanly, and that is the more interesting result. It runs Kenya 0.95, South Africa 0.79, Uganda 0.55 — so the country with the longest-established inflation-targeting regime does not have the strongest estimated response to inflation. **No country's estimated rule satisfies the Taylor principle**: all three coefficients are below one, Kenya's closest at 0.95. Read literally that would mean none of the three stabilises inflation through the real rate, which is not credible for South Africa in particular. It is better read as the same identification problem the IS elasticity displays: a policy rate that responds systematically to the cycle is the hardest thing in the model for a moment estimator to separate from the cycle itself, and the three monetary shock scales are among those the estimation could not identify at all.

18.3 Why Kenya fits so much worse

Kenya's moment correlation of 0.191 sits far below Uganda's 0.783 and South Africa's 0.774, and the reason is specific rather than general.

Table 26: Kenya's six largest moment misses

Moment	Data	Model
Variance, oil import price inflation	675.4	9.8
Variance, world oil price gap	548.9	**0.0**
Variance, relative oil import price gap	501.3	9.5
Covariance, oil price gap $\times$ oil import inflation	331.1	**0.0**
Autocovariance, world oil price gap	251.8	**0.0**
Autocovariance, relative oil import price gap	228.4	5.1

Notes: **All six of Kenya's largest misses are oil moments**, and three have a model variance of exactly zero. The estimator has set the world oil price shock variance to the zero bound, so the model puts no variance whatever on those series. Uganda's misses are similarly dominated by commodity moments but less exclusively so.

Every one of Kenya's six largest misses is an oil moment, and three are exact zeros. This is the identification problem of Section 17 in its purest observed form. The estimator is not fitting Kenya's oil block badly — it is declining to fit it at all, and because oil moments carry large variances they dominate the correlation statistic.

Two implications follow, and they point in opposite directions.

Kenya's headline fit understates the quality of its domestic estimate. Excluding the oil moments raises it substantially, and the domestic parameters — the IS elasticity, the Phillips curve, the policy rule — are estimated on moments the model does fit.

But the headline fit is still the right number to report. A model that puts zero variance on the world oil price is not a model one should use for oil scenarios in an oil-importing economy, and reporting a fit that excludes the failure would conceal exactly the limitation a user needs to know about. Kenya's 0.191 is honest; it is simply not a summary of everything the estimate does.

18.4 The projections compared

Figure 41: The three economies: history and projection to 2035

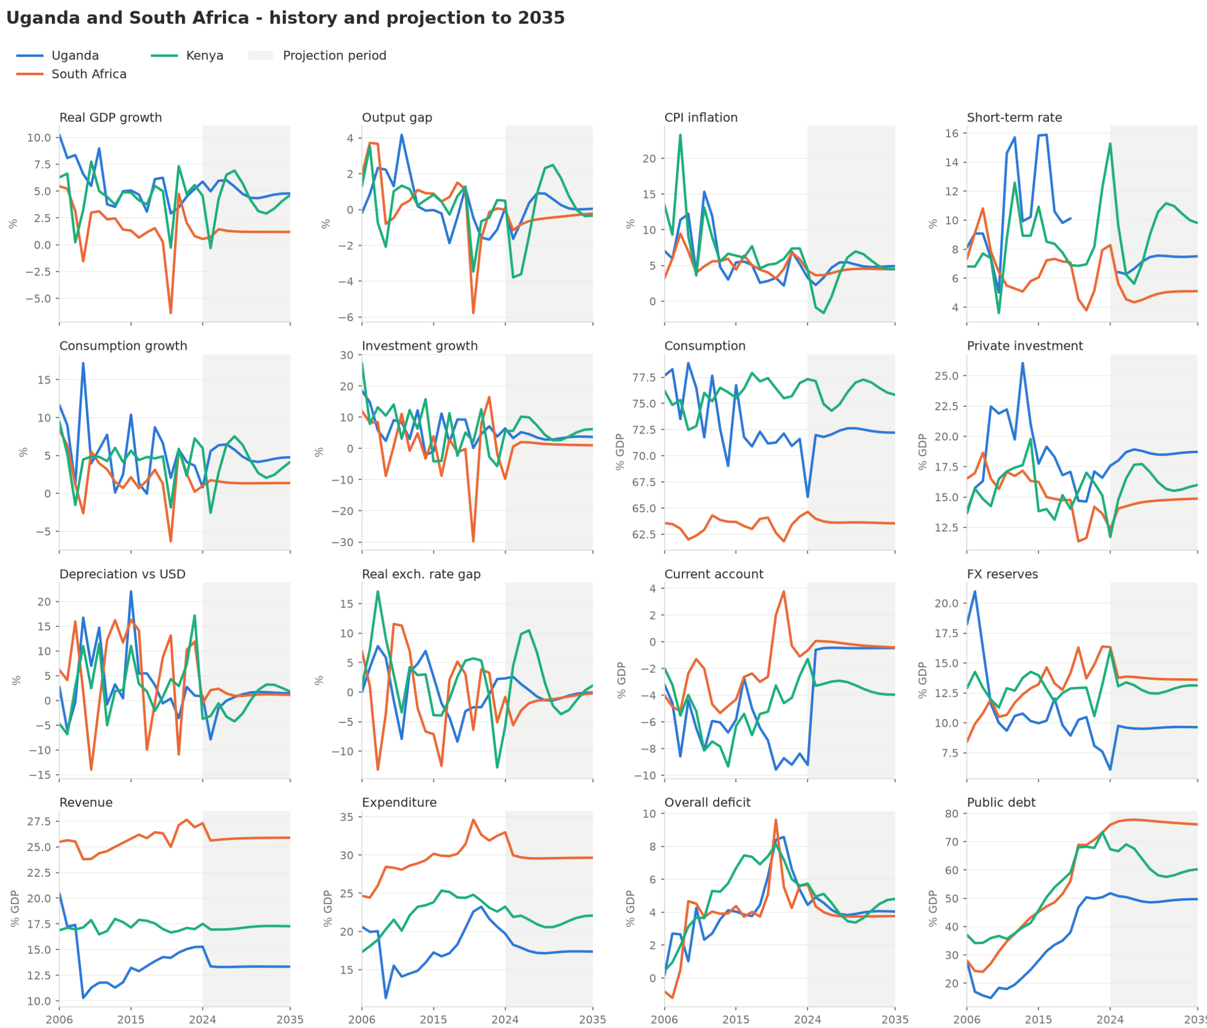

Notes: Blue is Uganda, orange South Africa, green Kenya. Data through 2024, model projection thereafter; the shaded band marks the projection period. Growth, inflation, rates and gaps in percent; ratios in percent of nominal GDP. Panels compare shapes and levels within a variable but amplitudes are not comparable across panels.

Table 27: Baseline projections compared, selected years

	Uganda		South Africa		Kenya	
	2030	2035	2030	2035	2030	2035
Real GDP growth, %	4.3	4.8	1.2	1.2	4.6	4.6
CPI inflation, %	5.2	4.9	4.5	4.5	4.8	4.7
Short-term rate, %	9.4	9.4	5.0	5.1	9.9	9.8
Current account, % GDP	-2.8	-2.8	-0.2	-0.4	-2.6	-2.6
Overall deficit, % GDP	3.9	4.0	3.7	3.7	4.7	4.7
Interest bill, % GDP	2.9	2.9	6.8	6.6	5.6	5.5
Public debt, % GDP	**48.8**	**49.8**	**77.2**	**76.2**	**59.9**	**59.5**

Notes: Unconditional projections with no future shocks, from model-consistent initial conditions. None contains off-model information: Uganda's excludes oil production, and neither of the others assumes any structural reform.

All three debt paths converge on their anchors, but from different directions and at different speeds. Uganda's falls from 51.8% to its 50% anchor. Kenya's falls further in absolute terms — from 67.3% to about 60% — because it starts further above target, and the fiscal rule responds proportionately. South Africa's converges from above but only slowly, because its growth–interest differential is the least favourable of the three.

Kenya's adjustment is the largest and the least comfortable. Closing a seven-point debt gap on a 4.7% deficit and a 5.6%-of-GDP interest bill is a more demanding path than Uganda's two-point gap on a 2.9% bill, even though the two economies grow at similar rates. The interest burden, not the growth rate, is what separates them.

18.5 What the comparison establishes

The method transfers, with a caveat about degree. Three economies spanning the African income distribution all estimate to convergent, coherent steady states. But only two reach a strong moment fit, and the third's weakness is traceable to a specific, diagnosable cause rather than to a general failure.

The parameters are country-specific and economically interpretable. They do not drift toward the generic calibration, and the pattern across three countries is richer than across two: Kenya sides with Uganda on the interest channel and with South Africa on inflation persistence, which no single pairwise comparison would have revealed.

Three observations are enough to isolate an outlier. South Africa's near-zero interest elasticity looked like a puzzle with two countries and looks like an identification failure with three. This is the clearest methodological payoff from the third case.

The same limitation binds all three, and hardest where it is most visible. Between 15 and 21 of 40 shock variances are identified in each country, and the unidentified set is dominated by the shared world-price block everywhere. Kenya simply shows it most starkly, because oil moments make up a larger share of what its correlation statistic is averaging over.

19. Conclusions

19.1 What was done

FINEX — the IMF's semi-structural model of internal and external balance — has been applied end to end to **Uganda, South Africa and Kenya**. Observable panels were built from national and international sources; each steady state was calibrated to that country's own sample means and policy anchors; 16 behavioural coefficients and 40 shock variances were estimated for each by Simulated Method of Moments on 166, 168 and 160 moment conditions respectively over 2001–2024; and the estimated models were used to produce baseline projections to 2035, impulse responses to all 83 structural shocks, a dated policy scenario with explicit attribution, and decompositions of each country's history since 2006. All three were run with identical code and an identical moment set.

19.2 What the model says about the two economies

Six findings are robust within the estimated framework.

The framework transfers, unevenly. All three project to convergent steady states despite spanning the African income distribution. But only Uganda and South Africa reach strong moment fits (0.783 and 0.774); Kenya reaches 0.191, and that weakness is diagnosable rather than general — all six of its largest moment misses are oil moments.

The third case isolates an outlier. Kenya's IS real-rate elasticity of 0.59 sits all but on top of Uganda's 0.58 and twenty-eight times above South Africa's 0.02. With two countries the South African estimate was a puzzle; with three it is identifiable as a probable identification failure, because the ordering is the opposite of what financial depth would predict.

Transmission mechanisms otherwise differ sharply and coherently. Ugandan inflation carries the least of its own past (0.63) where Kenyan and South African inflation are both strongly inertial (0.91 and 0.90), and South Africa's Phillips curve is the steepest of the three (0.46 against 0.19 for Uganda and 0.19 for Kenya), so a given output gap moves South African prices furthest even though its inflation is the most inertial. The Bank of Uganda smooths most heavily (0.73, against 0.68 and 0.47) and responds least strongly to inflation (0.55, against 0.95 and 0.79), while the Reserve Bank responds to the output gap an order of magnitude more strongly (0.47 against 0.04). A one-standard-deviation Ugandan tightening cuts growth 0.83 percentage points and a Kenyan one 0.59; South Africa's cuts it 0.20 on a shock 2.7 times larger.

Fiscal multipliers are nonetheless nearly the same, and the usual explanations for that do not survive three countries. The impact effect on growth is 0.62 in Uganda, 0.66 in Kenya and 0.57 in South Africa per standard deviation. It is tempting to attribute a small multiplier to import leakage or to monetary offset, and with two countries either story fits; with three, neither does. Kenya's imports adjust 207 times more slowly than Uganda's with no consequence for its multiplier, and the policy rate moves too little in every case to offset anything. **What the three share is the reversal**: growth rises on impact and falls back by almost the same amount the next year, leaving the level of output +0.05 %

higher in Uganda, +0.02 in Kenya and -0.02 in South Africa after five years. The impact multiplier and the cumulative multiplier differ by an order of magnitude, and it is the second that answers the question a finance ministry is asking.

The debt paths diverge on arithmetic, not behaviour. Uganda's ratio falls from 51.8% of GDP to 48.8% by 2030 and settles at its 50% anchor on a structural-deficit adjustment of half a point. South Africa's rises to 77.7% by 2027 before edging back toward its 75% anchor. The difference is a growth rate a quarter as fast and an interest bill more than twice as large — 6.5% of GDP against 2.9% — not a difference in fiscal rules. Uganda's favourable position rests on largely concessional external terms and would not survive a shift to market-rate borrowing.

Policy instruments are close to additive at ordinary shock sizes. Combining a 2027 monetary tightening with a 2030 fiscal consolidation produces an interaction term whose largest value across all observables and all years is 0.030. Within this range the two instruments may be analysed separately and added. That conclusion is explicitly conditional on shock size and on the effective lower bound not binding, and the interaction is computed for each scenario rather than assumed from this one.

19.3 What the exercise establishes about method

Three methodological results generalise beyond these two countries, and all three were learned from failures rather than successes.

Fit and stability are independent, and both must be checked. An unrestricted estimate of the investment persistence parameter scored a *better* moment objective than the restricted one and produced a projection that diverged to 252% growth. A good fit is not evidence of a usable model.

A single country's second moments cannot identify a shared foreign block. With 24 annual observations and 40 variances, 19 reach the zero bound. The world oil price gap ends with a model variance of exactly zero against a data variance of 549 — the model declines to fit it rather than fitting it badly. Because the foreign processes are common to both countries, the same moments are missed in each.

The model has a domain of validity, and it is narrower than its data requirements suggest. A complete observable panel and a respectable moment fit are not sufficient for a usable projection. FINEX's net-foreign-asset accumulation, for instance, is stabilising for a net debtor — the stock is deflated by nominal growth faster than it accumulates — and would compound rather than damp for a net creditor. Whether the framework suits a given economy is a structural question that must be settled before, not after, the estimation.

19.4 Priorities for further work

Four improvements would materially raise the quality of the results, in descending order of value.

1. **Estimate the foreign block off foreign data.** The shocks to US and euro-area activity, the world oil price and world commodity prices drive processes whose data are identical for every country.

Their standard deviations should be estimated once, directly, as univariate autoregressive residuals and held fixed while the non-negative least squares allocates only domestic shocks. This is the single highest-value change: it would repair the largest moment misses, make the cross-country fits comparable, and make fan charts meaningful.

2. **Close the policy-rate data gap.** The rate series ends in 2019, so the monetary block — the source of this report's strongest policy claims — is estimated on 20 observations from a superseded operational framework. Adding the Bank of Uganda's Central Bank Rate series would be a small task with a large return.

3. **Reconcile the national accounts and the balance of payments.** The model's steady-state current account sits three points above the sample mean because the two data sources disagree about the trade deficit, and the model can satisfy only one. This is why external levels carry a step at the 2024/25 handover.

4. **Run instrument-specific fiscal scenarios.** The aggregate multiplier reported here is less useful than a comparison of public consumption against public investment, which the model's structure supports and which would exploit the investment–growth nexus that distinguishes FINEX from conventional gap models.

19.5 A closing caution

The most important thing this report has to say about its own results is where they stop. The estimated models reproduce both countries' second moments well and project to coherent steady states, which licenses the dynamic analysis in Sections 9 to 13. They do not licence treating shock *scales* as estimated where they sit at the zero bound — between 15 and 21 of 40 across the three; they say nothing about the level accuracy of the external block; and they contain no off-model information, notably no Ugandan oil production, which is a first-order omission for that country's external and fiscal accounts from the late 2020s.

One estimate we have reported and used but do not believe. South Africa's IS real-rate elasticity of 0.02 is not credible as a structural parameter for an economy with its financial depth, and Section 18 sets out why it is more likely an identification failure — a highly endogenous policy rate contaminating the covariance the estimator reads — than a finding. We have left it in place rather than adjusting it, because a reported estimate that has been quietly overridden is worse than one that is flagged.

A model of this kind earns its place by making a coherent narrative possible — by letting an analyst say what drove an outcome and what would happen under an alternative — rather than by predicting well. Read that way, and with the four limitations of Section 17 kept in view, the results here are a usable foundation for policy analysis in both countries.

A. Using the model in a policy process

A model of this size is only useful if there is a disciplined way to use it. This section sets out the workflow the framework is built around, and the judgements that must be made by a person rather than by the code.

A.1 The forecast round

1. **Update the data and rebuild the panel.** The observable panel is regenerated from source rather than edited, so a data revision propagates consistently through every transform.

2. **Re-invert history.** This produces both the updated structural shocks and the new initial conditions. The shocks are worth inspecting in their own right: an unusually large recovered shock is often a data problem rather than an economic event.

3. **Run the unconditional baseline** and check it converges (Section 6.5). This is the reference, not the forecast.

4. **Impose judgement by conditioning**, not by editing outputs. Known policy measures, near-term official projections and off-model information such as oil production enter through conditioning paths (Section 6.6), so the rest of the accounts adjust coherently.

5. **Run the scenario suite** (Section 16) around the conditioned forecast to bound it.

6. **Decompose the result** (Section 13) and check the story is one you are willing to tell. If the decomposition attributes the forecast to a shock nobody can name, the forecast is not yet understood.

A.2 Judgements that cannot be automated

Four decisions in the workflow above are genuinely judgements, and the tooling deliberately surfaces rather than resolves them.

The inflation anchor. The projection converges to whatever is set here. Each country uses its own central bank's stated target: 5.0 % for Uganda, 5.0 % for Kenya and 4.5 % for South Africa, the midpoint of its 3–6 % band. **These appear in the steady-state table as 4.9, 4.9, 4.4 rather than as the targets themselves**, because the model carries inflation as $100 \times \log$ of the gross rate (Section 3.2); the difference is a transform, not a different number. For a country whose realised inflation has never been near its stated target, using the target produces a forecast that converges somewhere the economy has never been.

The debt anchor. This determines both the destination of the debt ratio and the sign and size of the fiscal adjustment the rule calls for. Setting it at the authorities' stated anchor and setting it at the sample mean are different exercises answering different questions. The three are set at 50, 60 and 75 % of GDP for Uganda, Kenya and South Africa, each the authorities' own stated anchor.

The instrument for a conditioning path. Conditioning inflation via the price shock and via the policy-rate shock are different economic statements that produce different paths for everything else (Section 6.6).

Whether a shock is anticipated. A pre-announced consolidation moves the economy before it lands; a surprise one does not. Both are legitimate; the specification must record which is meant.

A.3 Reading output responsibly

Three habits prevent the most common misreadings of this model's output, and all three are violations that would compile, run and produce plausible numbers.

Do not read a cyclical peak as a trend. The baseline's 6.1% growth in 2027 is the output gap closing from below, not Uganda's growth potential. The model's view of potential is the 4.9% steady state.

Do not read a single decomposition bar. Contributions are large and offsetting, and the ordering is fixed and order-dependent. The informative object is the balance within a year, not any one entry.

Do not treat a calibrated shock scale as an estimate. Roughly half the shock standard deviations reverted to calibrated values (Appendix I). Responses to those shocks have estimated shapes and assumed scales, and the distinction is recorded in every table in this report for exactly that reason.

B. Notation

Table 28: Symbols used in the model equations

Symbol	Meaning
Operators and decompositions	
$\overline{x}_t$	Trend component of x
$\widehat{x}_t$	Gap: deviation of x from its trend
Δx_t	First difference, $x_t - x_{t-1}$
$\mathbb{E}_t$	Model-consistent expectation formed at t
$\mathcal{S}(\cdot)$	Steady-state value of the argument
ε_t^z	Structural shock to equation z
x^{SS}	Steady-state (balanced growth path) value
Real quantities (superscript R denotes real)	
$Y_t,\ y_t^R$	Nominal GDP; real GDP in logs
$C_t^R,\ \widehat{c}_t^R$	Private consumption; its gap
$I_t^R,\ \widehat{\imath}_t^R$	Private investment; its gap
$G_t^R,\ \widehat{g}_t^R$	Government absorption; its gap
$X_t^{R,NNR}$	Non-natural-resource exports
$X_t^{R,NR}$	Natural-resource exports
$M_t^{R,NOIL}$	Non-oil imports
$M_t^{R,OIL}$	Oil imports
$K_t^R,\ K_t^{R,G}$	Private and public capital stocks
$\overline{a}_t^R$	Total factor productivity (quasi-TFP)
Prices and inflation	
$P_t^{\mathcal{X}},\ p_t^{\mathcal{X}}$	Deflator of component $\mathcal{X}$; its log
$p_t^{R,\mathcal{X}}$	Relative price of $\mathcal{X}$ against consumption
$\pi_t^{\mathcal{X}}$	Inflation of component $\mathcal{X}$
π_t^*	Inflation target
$RMC_t^{\mathcal{X}}$	Real marginal cost
$z_t,\ \widehat{z}_t$	Real exchange rate; its gap
s_t^{US}	Nominal exchange rate against the US dollar (log)
Interest rates and the premium	
i_t	Nominal policy rate
i^{ELB}	Effective lower bound
$r_t^R,\ \widehat{r}_t^R$	Real interest rate; its gap
$\gamma_t,\ \overline{\gamma}_t,\ \widehat{\gamma}_t$	UIP premium; its trend and gap
$\gamma_t^B,\ \gamma_t^{FXR},\ \gamma_t^{NFA^O}$	Risk-appetite terms in debt, reserves, net foreign assets
External accounts (superscript Y denotes % of GDP)	
$CA_t^Y,\ FA_t^Y$	Current and financial account
$FA_t^{O,Y}$	Endogenous private financial inflows
$NFA_t^{O,Y}$	Private net foreign assets
$FXR_t^Y,\ FXI_t^Y,\ FXA_t^Y$	Reserves; discretionary intervention; planned accumulation
REM_t^Y	Remittances
Fiscal	
$d_t^s,\ d_t^{s,*}$	Structural deficit; its target
$b_t,\ b_t^*,\ \widetilde{b}_t$	Public debt; its anchor; forward-looking deviation
$b_t^{lcy},\ b_t^{fcy}$	Local- and foreign-currency debt
GE_t^{R,I^G}	Public investment
$\tau_t^C,\ \tau_t^Y,\ \tau_t^M$	Consumption, income and import tax rates
$\tau^{FA^O},\ \tau^{FA^O,adm}$	Capital inflow tax; administrative capital-flow measures

Notes: Notation follows Berg et al. (2023) where possible. Model variable names in the code differ — `L_GDP_GAP` for $\widehat{y}_t^R$, `DL_P_CONS` for π_t^C, `RS` for i_t, `DEB_RAT` for b_t — and the mapping is recorded in the data dictionary (Appendix H). The convention throughout is that `L_` denotes $100\log$, `DL_` its first difference, `_GAP` a deviation from trend, `_BAR` a trend, `_RAT` a ratio to nominal GDP, and `RES_` a structural shock.

C. The real block in detail: equations 1–82

This appendix sets out the first eighty-two equations of FINEX exactly as they are numbered in Appendix A.IV of Berg et al. (2023), and explains what each group of them does. Together they are the entire real side of the model: the supply of output, the market-clearing conditions that tie the demand components to it, and the five demand blocks — consumption, investment, non-resource exports, resource exports, and imports split between oil and non-oil. Everything that follows in the model — prices, the monetary and fiscal blocks, the external financing conditions — acts on the economy through the equations below.

The numbering is the paper's own. Equation n here is equation n in the appendix and, equally, the nth equation in the Julia implementation described in Appendix D; the correspondence was checked across the whole range rather than assumed, and the mathematics printed here is generated from the same source the code solves, so the two cannot drift apart as the model is edited.

How to read these equations

One denominator, named once. Six of the equations below express a demand component as a share of GDP, and all six divide by the same thing: potential nominal expenditure at purchaser prices. Written out it runs to most of a page, and printing it six times would push those equations off the page while telling the reader nothing new after the first. It is therefore named $\overline{Y}^N$ here and stated once:

$$\overline{Y}^N = \left(1 + \frac{\tau_t^C}{100}\right) \cdot \overline{C}_t^R + P_t^{R,I} \cdot \overline{I}_t^R + \overline{P}_t^{R,X^{NNR}} \cdot \overline{X}_t^{R,NNR} + \overline{P}_t^{R,X^{NR}} \cdot \overline{X}_t^{R,NR} \tag{37}$$

$$notag \tag{38}$$

$$\tag{39}$$

$$- \overline{P}_t^{R,M^{NOIL}} \cdot \left(1 + \frac{\tau_t^{M^{NOIL}}}{100}\right) \cdot \overline{M}_t^{R,NOIL} - \overline{P}_t^{R,M^{OIL}} \cdot \left(1 + \frac{\tau_t^{M^{OIL}}}{100}\right) \cdot \overline{M}_t^{R,OIL} \tag{40}$$

$$notag \tag{41}$$

$$\tag{42}$$

$$+ \overline{P}_t^{R,G} \cdot \overline{G}_t^R$$

Each ratio definition below then reads as a share of $\overline{Y}^N$, which is also the clearest way to see the accounting identity behind them: the shares of consumption, investment, the two export categories and the two import categories sum to one by construction.

C.1 Supply: capital, potential output and productivity

FINEX's supply side is deliberately thin, but it is not a simple trend. Potential output grows with the accumulation of capital — private and public, entering separately — and with a residual productivity term. This is the investment–growth nexus that distinguishes FINEX from earlier gap–trend projection models: an investment surge does not merely raise demand today, it raises the economy's productive capacity thereafter, and public and private capital are allowed to do so at different rates.

Equation (1) is the capital accumulation identity, written so that investment is the residual of the capital stock and its depreciated predecessor; δ^{K^R} is the depreciation rate. Equations (2) and (3) convert the stock to logs and difference it.

$$I_t^R = K_t^R - \left(1 - \delta^{K^R}\right) \cdot K_{t-1}^R \tag{1}$$

$$k_t^R = 100 \cdot \log\left(K_t^R\right) \tag{2}$$

$$\Delta k_t^R = k_t^R - k_{t-1}^R \tag{3}$$

Equations (4) to (6) introduce output and the central decomposition: observed log output is potential plus the gap.

$$y_t^R = 100 \cdot \log\left(Y_t^R\right) \tag{4}$$

$$y_t^R = \overline{y}_t^R + \hat{y}_t^R \tag{5}$$

$$\overline{y}_t^R = 100 \cdot \log\left(\overline{Y}_t^R\right) \tag{6}$$

Equation (7) is the heart of the supply block and the single most consequential equation in this section. Potential output growth responds to the growth of private capital and of public capital, each measured as a deviation from its own steady-state growth rate, plus trend productivity growth. The two coefficients $c_1^{\Delta\overline{y}^R}$ and $c_2^{\Delta\overline{y}^R}$ are what make public investment a supply-side instrument rather than only a demand impulse: a government capital programme raises $\Delta k^{R,G}$ and, through this equation, the economy's potential growth rate. In the calibration used in this report public capital is productive only after a lag, so the effect appears in potential output some years after the spending occurs.

$$\Delta\overline{y}_t^R = c_1^{\Delta\overline{y}^R} \cdot \left(\Delta k_t^R - \Delta k^{R,SS}\right) + c_2^{\Delta\overline{y}^R} \cdot \left(\Delta k_t^{R,G} - \Delta k^{R,G,SS}\right) + \Delta\tilde{a}_t^R \tag{7}$$

Equation (8) is the trend-growth identity and (9) the observed-growth identity. Equations (10) to (13) carry the productivity residual: $\tilde{a}^R$ is what the paper calls *quasi* total factor productivity — quasi because it is not measured from a production function with factor shares, but backed out as whatever potential growth is left once the capital terms are accounted for. Its growth has a persistent component, equation (12), which mean-reverts to the steady-state rate $\Delta\tilde{a}^{R,SS}$ and receives its own shock.

$$\Delta\overline{y}_t^R = \overline{y}_t^R - \overline{y}_{t-1}^R \tag{8}$$

$$\Delta y_t^R = y_t^R - y_{t-1}^R \tag{9}$$

$$\tilde{a}_t^R = 100 \cdot \log\left(\tilde{A}_t^R\right) \tag{10}$$

$$\Delta\tilde{a}_t^R = \Delta\tilde{a}_t^{R,g} + \varepsilon_t^{\Delta\tilde{a}^R} \tag{11}$$

$$\Delta\tilde{a}_t^{R,g} = c_1^{\Delta\tilde{a}^{R,g}} \cdot \Delta\tilde{a}_{t-1}^{R,g} + \left(1 - c_1^{\Delta\tilde{a}^{R,g}}\right) \cdot \Delta\tilde{a}^{R,SS} + \varepsilon_t^{\Delta\tilde{a}^{R,g}} \tag{12}$$

$$\Delta\tilde{a}_t^R = \tilde{a}_t^R - \tilde{a}_{t-1}^R \tag{13}$$

C.2 Market clearing: the trend and the gap

The next two equations are the adding-up conditions, and they are where the supply block and the five demand blocks meet. Equation (14) requires potential output growth to equal the expenditure-share-weighted sum of the potential growth rates of the components, with the shares lagged one period. Equation (15) imposes the same weighting on the gaps: the output gap is the share-weighted sum of the consumption, investment, government, export and import gaps.

These are the equations that make FINEX a general-equilibrium system rather than a collection of independent gap processes. A shock that opens an investment gap shows up in the output gap here, in proportion to investment's share of GDP, and from there feeds back into every demand block through the income and capacity terms. Note that imports enter with a negative sign, so an import gap subtracts from the output gap, as the national accounts require.

$$\begin{aligned}\Delta\overline{y}_t^R = {} & \frac{\overline{C}_{t-1}^Y}{100} \cdot \Delta\overline{c}_t^R + \frac{\overline{I}_{t-1}^Y}{100} \cdot \Delta\overline{i}_t^R + \frac{\overline{X}_{t-1}^{NNR,Y}}{100} \cdot \Delta\overline{x}_t^{R,NNR} \\ & + \frac{\overline{X}_{t-1}^{NR,Y}}{100} \cdot \Delta\overline{x}_t^{R,NR} - \frac{\overline{M}_{t-1}^{NOIL,Y}}{100} \cdot \Delta\overline{m}_t^{R,NOIL} - \frac{\overline{M}_{t-1}^{OIL,Y}}{100} \cdot \Delta\overline{m}_t^{R,OIL} \\ & + \frac{\overline{G}_{t-1}^Y}{100} \cdot \Delta\overline{g}_t^R \end{aligned} \tag{14}$$

$$\begin{aligned}\hat{y}_t^R = {} & \frac{\overline{C}_t^Y}{100} \cdot \hat{c}_t^R + \frac{\overline{I}_t^Y}{100} \cdot \hat{i}_t^R + \frac{\overline{X}_t^{NNR,Y}}{100} \cdot \hat{x}_t^{R,NNR} \\ & + \frac{\overline{X}_t^{NR,Y}}{100} \cdot \hat{x}_t^{R,NR} - \frac{\overline{M}_t^{NOIL,Y}}{100} \cdot \hat{m}_t^{R,NOIL} - \frac{\overline{M}_t^{OIL,Y}}{100} \cdot \hat{m}_t^{R,OIL} \\ & + \frac{\overline{G}_t^Y}{100} \cdot \hat{g}_t^R \end{aligned} \tag{15}$$

C.3 Private consumption

Equation (16) is the consumption gap equation — FINEX's IS curve for households. It is hybrid: consumption today depends on consumption yesterday (habit, or rule-of-thumb behaviour) and on expected consumption tomorrow (the forward-looking Euler term), with the two weights summing to less than one so that the process is stationary. Beyond that inertia it responds to the real interest rate gap, to the real exchange rate, to the fiscal stance, and to remittances. The real rate term is the channel through which monetary policy reaches demand; its coefficient is one of the sixteen behavioural parameters estimated by the Simulated Method of Moments in Section 5.7.

The stability restriction discussed in Section 5.4 applies to this equation: the sum of the backward- and forward-looking weights must remain below one, or the consumption process ceases to mean-revert and the projection diverges even when the in-sample moment fit improves.

$$\begin{aligned}\hat{c}_t^R &= c_1^{\hat{c}^R} \cdot \hat{c}_{t-1}^R + c_2^{\hat{c}^R} \cdot \hat{c}_{t+1}^R - c_3^{\hat{c}^R} \cdot \hat{r}_t^R + c_4^{\hat{c}^R} \cdot \hat{y}_t^R \\ &- c_5^{\hat{c}^R} \cdot \left(\tau_t^C - \tau_{t+1}^C\right) + c_6^{\hat{c}^R} \cdot R\hat{E}M_t^Y + c_7^{\hat{c}^R} \cdot \widehat{GE}_t^{Tr,Y} + \varepsilon_t^{\hat{c}^R}\end{aligned} \tag{16}$$

$$\hat{c}_t^{R,E} = \hat{c}_{t+1}^R \tag{17}$$

Equation (18) governs potential consumption growth. It mean-reverts to a combination of potential output growth and the steady-state consumption ratio, which is what prevents the consumption trend from drifting away from the economy's capacity to produce over long horizons.

$$\begin{aligned}\Delta\overline{c}_t^R = c_1^{\Delta\overline{c}^R} \cdot \Delta\overline{c}_{t-1}^R + \left(1 - c_1^{\Delta\overline{c}^R}\right) \cdot \left(\Delta\overline{p}_t^{R,Y} + \Delta\overline{y}_t^R\right) \\ - c_2^{\Delta\overline{c}^R} \cdot \left(\overline{C}_t^Y - C^{Y,SS} + c_3^{\Delta\overline{c}^R} \cdot \left(\tau_t^C - \tau^{C,SS}\right) - c_4^{\Delta\overline{c}^R} \cdot \left(\overline{REM}_t^Y - REM^{Y,SS}\right) \\ - c_5^{\Delta\overline{c}^R} \cdot \left(\overline{GE}_t^{Tr,Y} - \overline{GE}^{Tr,Y,SS}\right)\right) + \varepsilon_t^{\Delta\overline{c}^R}\end{aligned} \tag{18}$$

Equations (19) to (23) are the identity boilerplate described above: the gap–trend decomposition, the logarithms of the level and the trend, and the two growth definitions.

$$c_t^R = \hat{c}_t^R + \overline{c}_t^R \tag{19}$$

$$c_t^R = 100 \cdot \log\left(C_t^R\right) \tag{20}$$

$$\overline{c}_t^R = 100 \cdot \log\left(\overline{C}_t^R\right) \tag{21}$$

$$\Delta c_t^R = c_t^R - c_{t-1}^R \tag{22}$$

$$\Delta\overline{c}_t^R = \overline{c}_t^R - \overline{c}_{t-1}^R \tag{23}$$

Equations (24) and (25) convert consumption into a share of GDP, and they are more substantive than they look. Both carry the consumption tax wedge $(1 + \tau^C/100)$, so the ratio is measured at purchaser prices while the volume is measured at basic prices. The denominator of (24) is the full expenditure identity written in trend terms — consumption, investment, government, exports less imports, each at its own relative price — so the potential ratios of all five components are mutually consistent by construction. Equation (25) does the same for the observed ratio using the relative price of consumption to the GDP deflator. A change in the consumption tax rate therefore moves the measured consumption share even when the underlying volume is unchanged, which is exactly the behaviour a fiscal simulation requires.

$$\overline{C}_t^Y = \frac{100 \cdot \left(1 + \frac{\tau_t^C}{100}\right) \cdot \overline{C}_t^R}{\overline{Y}^N} \tag{24}$$

$$C_t^Y = 100 \cdot \frac{P_t^C \cdot \left(1 + \frac{\tau_t^C}{100}\right)}{P_t^Y} \cdot \frac{C_t^R}{Y_t^R} \tag{25}$$

C.4 Private investment

The investment block mirrors the consumption block in structure, and the parallel is deliberate: the same hybrid gap equation, the same mean-reverting trend, the same identity boilerplate, the same tax-wedged ratio definitions. What differs is the set of driving variables.

Equation (26) makes the investment gap depend on its own lag and lead, on the real interest rate gap, on the output gap — the accelerator term, by which strong demand pulls in capacity — and on the relative price of investment goods, which is how an exchange rate depreciation raises the cost of imported capital equipment and discourages investment. In practice the accelerator and the real rate are the two channels that dominate the impulse responses in Section 9.

$$\begin{aligned}\hat{i}_t^R = {} & c_1^{\hat{i}^R} \cdot \hat{i}_{t-1}^R + c_2^{\hat{i}^R} \cdot \hat{i}_{t+1}^R - c_3^{\hat{i}^R} \cdot \hat{r}_t^R + c_4^{\hat{i}^R} \cdot \hat{c}_{t+1}^R + c_6^{\hat{i}^R} \cdot \hat{x}_{t+1}^{R,NNR} \\ & - c_5^{\hat{i}^R} \cdot \left(\hat{z}_t^{Mw} - \hat{p}_t^{R,I}\right) + \varepsilon_t^{\hat{i}^R}\end{aligned} \tag{26}$$

$$\hat{i}_t^{R,E} = \hat{i}_{t+1}^R \tag{27}$$

$$\begin{aligned}\Delta\bar{i}_t^R = {} & c_1^{\Delta\bar{i}^R} \cdot \Delta\bar{i}_{t-1}^R \\ & + \left(1 - c_1^{\Delta\bar{i}^R}\right) \cdot \left(\Delta\overline{p}_t^{R,Y} + \Delta\overline{y}_t^R - \Delta\overline{p}_t^{R,I}\right) - c_2^{\Delta\bar{i}^R} \cdot \left(\overline{I}_t^Y - I^{Y,SS}\right. \\ & + c_3^{\Delta\bar{i}^R} \cdot \left(\overline{r}_t^R - \left(\overline{r}^{R,US,SS} + \mathrm{RR\_PREM\_SS} + \Delta\overline{z}^{SS,US}\right)\right) + c_4^{\Delta\bar{i}^R} \cdot \left.\left(\tau_t^Y - \tau^{Y,SS}\right)\right) \\ & + \varepsilon_t^{\Delta\bar{i}^R}\end{aligned} \tag{28}$$

Equations (29) and (30) are the ratio definitions, carrying the relative price of investment $P^{R,I}$ in place of the consumption tax wedge, and (31) to (35) are the identities.

$$\overline{I}_t^Y = \frac{100 \cdot P_t^{R,I} \cdot \overline{I}_t^R}{\overline{Y}^N} \tag{29}$$

$$I_t^Y = 100 \cdot \frac{P_t^I}{P_t^Y} \cdot \frac{I_t^R}{Y_t^R} \tag{30}$$

$$i_t^R = \hat{i}_t^R + \bar{i}_t^R \tag{31}$$

$$\Delta i_t^R = i_t^R - i_{t-1}^R \tag{32}$$

$$i_t^R = 100 \cdot \log\left(I_t^R\right) \tag{33}$$

$$\Delta \bar{i}_t^R = \bar{i}_t^R - \bar{i}_{t-1}^R \tag{34}$$

$$\bar{i}_t^R = 100 \cdot \log\left(\bar{I}_t^R\right) \tag{35}$$

C.5 Non-resource exports

FINEX splits exports in two, and the split matters for the African applications in this report. Non-resource exports (NNR) behave like a conventional traded good: they respond to foreign demand and to competitiveness. Resource exports, treated in the next subsection, do not — they are driven by supply capacity and world prices, and respond weakly if at all to the real exchange rate. Collapsing the two would attribute to competitiveness what is really a commodity-price or production-capacity story.

Equations (36) and (37) add the two export ratios to a total. Equation (38) is the non-resource export gap: hybrid in its own lag and lead, driven by the foreign output gap — the external demand channel — and by the real effective exchange rate gap, which is the competitiveness channel.

$$X_t^Y = X_t^{NNR,Y} + X_t^{NR,Y} \tag{36}$$

$$\overline{X}_t^Y = \overline{X}_t^{NNR,Y} + \overline{X}_t^{NR,Y} \tag{37}$$

$$\begin{aligned}\hat{x}_t^{R,NNR} &= c_1^{\hat{x}^{R,NNR}} \cdot \hat{x}_{t-1}^{R,NNR} + c_2^{\hat{x}^{R,NNR}} \cdot \hat{x}_{t+1}^{R,NNR} + c_3^{\hat{x}^{R,NNR}} \cdot \left(\hat{z}_t^{Xw} - \hat{p}_t^{R,X^{NNR}}\right) \\ &\quad + c_4^{\hat{x}^{R,NNR}} \cdot \hat{y}_t^{R,*} + \varepsilon_t^{\hat{x}^{R,NNR}}\end{aligned} \tag{38}$$

$$\hat{x}_t^{R,NNR,E} = \hat{x}_{t+1}^{R,NNR} \tag{39}$$

Equations (40) and (41) govern the trend. Equation (41) introduces a device used throughout the trade blocks: a proxy for the equilibrium supply of non-resource exports, which accumulates the deviation of the relative export price from its steady state. The idea is that persistently favourable relative prices draw resources into the export sector and raise its potential scale, with a lag; this is what lets a sustained terms-of-trade improvement raise trend exports rather than only cyclical ones.

$$\begin{aligned}\Delta \overline{x}_t^{R,NNR} &= c_1^{\Delta \overline{x}^{R,NNR}} \cdot \Delta \overline{x}_{t-1}^{R,NNR} \\ &\quad + \left(1 - c_1^{\Delta \overline{x}^{R,NNR}}\right) \cdot \left(\Delta \overline{p}_t^{R,Y} + \Delta \overline{y}_t^R - \Delta \overline{p}_t^{R,X^{NNR}}\right) \\ &\quad - c_2^{\Delta \overline{x}^{R,NNR}} \cdot \left(\overline{X}_t^{NNR,Y} - \left(X^{NNR,Y,SS} + c_3^{\Delta \overline{x}^{R,NNR}} \cdot \overline{X}_t^{R,NNR,S}\right)\right) + \varepsilon_t^{\Delta \overline{x}^{R,NNR}}\end{aligned} \tag{40}$$

$$\begin{aligned}\overline{X}_t^{R,NNR,S} &= c_1^{\overline{X}^{R,NNR,S}} \cdot \overline{X}_{t-1}^{R,NNR,S} + c_2^{\overline{X}^{R,NNR,S}} \cdot \left(\Delta \overline{p}_t^{R,X^{NNR}} - \Delta \overline{p}^{R,X^{NNR},SS}\right) \\ &\quad + \left(1 - c_2^{\overline{X}^{R,NNR,S}}\right) \cdot \left(\Delta \overline{y}_t^{R,*} - \Delta \overline{y}^{R,*,SS}\right) + \varepsilon_t^{\overline{X}^{R,NNR,S}}\end{aligned} \tag{41}$$

Equations (42) to (48) are the ratio definitions and the identity boilerplate, following exactly the pattern set out for consumption.

$$\overline{X}_t^{NNR,Y} = \frac{100 \cdot \overline{P}_t^{R,X^{NNR}} \cdot \overline{X}_t^{R,NNR}}{\overline{Y}^N} \tag{42}$$

$$X_t^{NNR,Y} = 100 \cdot \frac{P_t^{X^{NNR}}}{P_t^Y} \cdot \frac{X_t^{R,NNR}}{Y_t^R} \tag{43}$$

$$x_t^{R,NNR} = \hat{x}_t^{R,NNR} + \overline{x}_t^{R,NNR} \tag{44}$$

$$\Delta x_t^{R,NNR} = x_t^{R,NNR} - x_{t-1}^{R,NNR} \tag{45}$$

$$x_t^{R,NNR} = 100 \cdot \log\left(X_t^{R,NNR}\right) \tag{46}$$

$$\Delta \overline{x}_t^{R,NNR} = \overline{x}_t^{R,NNR} - \overline{x}_{t-1}^{R,NNR} \tag{47}$$

$$\overline{x}_t^{R,NNR} = 100 \cdot \log\left(\overline{X}_t^{R,NNR}\right) \tag{48}$$

C.6 Natural-resource exports

The resource export block has the same skeleton but a different economic content. Equation (49) drives the resource export gap by the world resource price rather than by the real exchange rate, and equations (50) and (51) build the equilibrium supply proxy from the trend in the real resource price. The behavioural claim is that resource production expands when the real price of the resource is high relative to its own trend, subject to a long lag reflecting the investment cycle in extractive industries.

$$\hat{x}_t^{R,NR} = c_1^{\hat{x}^{R,NR}} \cdot \hat{x}_{t-1}^{R,NR} + c_2^{\hat{x}^{R,NR}} \cdot \hat{p}_t^{R,X^{NR}} + \varepsilon_t^{\hat{x}^{R,NR}} \tag{49}$$

$$\begin{aligned}\Delta \overline{x}_t^{R,NR} = {} & c_1^{\Delta \overline{x}^{R,NR}} \cdot \Delta \overline{x}_{t-1}^{R,NR} \\ & + \left(1 - c_1^{\Delta \overline{x}^{R,NR}}\right) \cdot \left(\Delta \overline{p}_t^{R,Y} - \Delta \overline{p}_t^{R,X^{NR}} + \Delta \overline{y}_t^R\right) - c_2^{\Delta \overline{x}^{R,NR}} \cdot \Big(\overline{X}_t^{NR,Y} \\ & - X^{NR,Y,SS} - c_4^{\Delta \overline{x}^{R,NR}} \cdot \overline{X}_t^{R,NR,S} + c_3^{\Delta \overline{x}^{R,NR}} \cdot \left(\tau_t^{NR} - \tau^{NR,SS}\right)\Big) + \varepsilon_t^{\Delta \overline{x}^{R,NR}}\end{aligned} \tag{50}$$

$$\overline{X}_t^{R,NR,S} = c_1^{\overline{X}^{R,NR,S}} \cdot \overline{X}_{t-1}^{R,NR,S} + \Delta \overline{p}_t^{R,X^{NR}} - \Delta \overline{p}^{R,X^{NR},SS} + \varepsilon_t^{\overline{X}^{R,NR,S}} \tag{51}$$

Equations (52) to (58) are the ratio and identity equations. Equations (59) and (60) are specific to this block and have no counterpart elsewhere: they compute the profit of resource producers as the margin between the resource export price and a cost share of it, expressed as a ratio to GDP, together with its equilibrium counterpart. This is the quantity a resource royalty or windfall tax acts on in the fiscal block, and it is why FINEX can represent a resource revenue regime without a separate sectoral model.

$$\overline{X}_t^{NR,Y} = \frac{100 \cdot \overline{P}_t^{R,X^{NR}} \cdot \overline{X}_t^{R,NR}}{\overline{Y}^N} \tag{52}$$

$$X_t^{NR,Y} = 100 \cdot \frac{P_t^{X^{NR}}}{P_t^Y} \cdot \frac{X_t^{R,NR}}{Y_t^R} \tag{53}$$

$$x_t^{R,NR} = \overline{x}_t^{R,NR} + \hat{x}_t^{R,NR} \tag{54}$$

$$\Delta x_t^{R,NR} = x_t^{R,NR} - x_{t-1}^{R,NR} \tag{55}$$

$$x_t^{R,NR} = 100 \cdot \log\left(X_t^{R,NR}\right) \tag{56}$$

$$\Delta \overline{x}_t^{R,NR} = \overline{x}_t^{R,NR} - \overline{x}_{t-1}^{R,NR} \tag{57}$$

$$\overline{x}_t^{R,NR} = 100 \cdot \log\left(\overline{X}_t^{R,NR}\right) \tag{58}$$

$$G_t^{NR,Y} = 100 \cdot \frac{P_t^{X^{NR}} - c_1^{G^{NR,Y}} \cdot P_t^{X^{NR},MAVG}}{P_t^Y} \cdot \frac{X_t^{R,NR}}{Y_t^R} \tag{59}$$

$$\overline{G}_t^{NR,Y} = \left(1 - c_1^{G^{NR,Y}}\right) \cdot \overline{X}_t^{NR,Y} \tag{60}$$

C.7 Non-oil imports

Imports are split between oil and non-oil for the same reason exports are split between resource and non-resource: the two respond to different things, and for a net oil importer the distinction is the difference between a terms-of-trade shock and a demand shock. Equation (61) adds the two ratios.

Equation (62) is the non-oil import gap. It is driven by domestic absorption — it is imports of goods a growing economy buys — and by the relative price of imports including the tariff wedge. Equation (65) is the demand proxy, the import counterpart of the export supply proxy of equation (41), accumulating the deviation of the relative import price from its trend.

$$M_t^Y = M_t^{NOIL,Y} + M_t^{OIL,Y} \tag{61}$$

$$\begin{aligned}\hat{m}_t^{R,NOIL} = {} & c_1^{\hat{m}^{R,NOIL}} \cdot \hat{m}_{t-1}^{R,NOIL} + c_2^{\hat{m}^{R,NOIL}} \cdot \hat{x}_t^{R,NNR} + c_3^{\hat{m}^{R,NOIL}} \cdot \hat{i}_t^R + c_4^{\hat{m}^{R,NOIL}} \cdot \hat{c}_t^R \\ & - c_5^{\hat{m}^{R,NOIL}} \cdot \hat{p}_t^{R,M^{NOIL}} - c_6^{\hat{m}^{R,NOIL}} \cdot \left(\hat{p}_t^{R,M^{NOIL}} - \hat{p}_t^{R,I}\right) + c_7^{\hat{m}^{R,NOIL}} \cdot \hat{g}_t^R \\ & - c_8^{\hat{m}^{R,NOIL}} \cdot \left(\tau_t^{M^{NOIL}} - \tau_{t+1}^{M^{NOIL}}\right) + \varepsilon_t^{\hat{m}^{R,NOIL}}\end{aligned} \tag{62}$$

$$\hat{m}_t^{R,NOIL,E} = \hat{m}_{t+1}^{R,NOIL} \tag{63}$$

$$\begin{aligned}\Delta\overline{m}_t^{R,NOIL} &= c_1^{\Delta\overline{m}^{R,NOIL}} \cdot \Delta\overline{m}_{t-1}^{R,NOIL} \\ &+ \left(1 - c_1^{\Delta\overline{m}^{R,NOIL}}\right) \cdot \left(\Delta\overline{p}_t^{R,Y} + \Delta\overline{y}_t^{R} - \Delta\overline{p}_t^{R,M^{NOIL}}\right) - c_2^{\Delta\overline{m}^{R,NOIL}} \cdot \Big(\\ &\overline{M}_t^{NOIL,Y} - M_t^{NOIL,Y,SS} - c_3^{\Delta\overline{m}^{R,NOIL}} \cdot \overline{M}_t^{R,NOIL,D} - c_4^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\overline{C}_t^{Y} - C^{Y,SS}\right) \\ &- c_5^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\overline{I}_t^{Y} - I^{Y,SS}\right) - c_6^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\overline{X}_t^{NNR,Y} - X^{NNR,Y,SS}\right) \\ &- c_7^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\overline{X}_t^{NR,Y} - X^{NR,Y,SS}\right) - c_8^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\overline{G}_t^{Y} - G^{Y,SS}\right) \\ &+ c_9^{\Delta\overline{m}^{R,NOIL}} \cdot \left(\tau_t^{M^{NOIL}} - \tau^{M^{NOIL},SS}\right)\Big) + \varepsilon_t^{\Delta\overline{m}^{R,NOIL}} \end{aligned} \tag{64}$$

$$\overline{M}_t^{R,NOIL,D} = c_1^{\overline{M}^{R,NOIL,D}} \cdot \overline{M}_{t-1}^{R,NOIL,D} - \Delta\overline{p}_t^{R,M^{NOIL}} + \Delta\overline{p}^{R,M^{NOIL},SS} + \varepsilon_t^{\overline{M}^{R,NOIL,D}} \tag{65}$$

Equations (66) and (67) deserve attention. Equation (66) is the potential import ratio, carrying the tariff wedge $(1 + \tau^{M^{NOIL}}/100)$ in the same way the consumption ratio carries the consumption tax; a tariff change therefore alters the measured import share directly. Equation (67) pins the steady-state import ratio to the steady-state shares of the absorption components, which is the restriction that keeps the trade balance from drifting without limit in long projections. Equations (68) to (73) complete the block with the observed ratio and the identities.

$$\overline{M}_t^{NOIL,Y} = \frac{100 \cdot \overline{P}_t^{R,M^{NOIL}} \cdot \left(1 + \frac{\tau_t^{M^{NOIL}}}{100}\right) \cdot \overline{M}_t^{R,NOIL}}{\overline{Y}^{N}} \tag{66}$$

$$M_t^{NOIL,Y,SS} = C^{Y,SS} + I^{Y,SS} + G^{Y,SS} + X^{NNR,Y,SS} + X^{NR,Y,SS} - M^{OIL,Y,SS} - 100 \tag{67}$$

$$M_t^{NOIL,Y} = 100 \cdot \frac{P_t^{M^{NOIL}} \cdot \left(1 + \frac{\tau_t^{M^{NOIL}}}{100}\right)}{P_t^{Y}} \cdot \frac{M_t^{R,NOIL}}{Y_t^{R}} \tag{68}$$

$$m_t^{R,NOIL} = \hat{m}_t^{R,NOIL} + \overline{m}_t^{R,NOIL} \tag{69}$$

$$\Delta m_t^{R,NOIL} = m_t^{R,NOIL} - m_{t-1}^{R,NOIL} \tag{70}$$

$$m_t^{R,NOIL} = 100 \cdot \log\left(M_t^{R,NOIL}\right) \tag{71}$$

$$\Delta\overline{m}_t^{R,NOIL} = \overline{m}_t^{R,NOIL} - \overline{m}_{t-1}^{R,NOIL} \tag{72}$$

$$\overline{m}_t^{R,NOIL} = 100 \cdot \log\left(\overline{M}_t^{R,NOIL}\right) \tag{73}$$

C.8 Oil imports

The oil import block, equations (74) to (82), repeats the non-oil structure with the world oil price in place of the general import price. Equation (74) makes the oil import gap respond to absorption and to the real oil price gap; equation (77) is the corresponding demand proxy, entering with a negative sign because a sustained rise in the real oil price reduces the equilibrium volume of oil imports — the substitution response that a pure price pass-through would miss.

This block is worth flagging for a reason that emerges only in estimation. In the three-country application of Section 5.7 it is the oil import equations that the data resist most strongly: for the oil-importing economies in the sample, the observed variance of oil import growth exceeds what the model can generate by an order of magnitude, and the moments involving equations (74) to (82) account for the majority of the total moment misfit. The structure below is the paper's; the difficulty is in matching it to African data where oil import volumes swing far more violently than a smooth gap process allows.

$$\begin{aligned}\hat{m}_t^{R,OIL} = {} & c_1^{\hat{m}^{R,OIL}} \cdot \hat{m}_{t-1}^{R,OIL} + c_2^{\hat{m}^{R,OIL}} \cdot \hat{x}_t^{R,NNR} + c_3^{\hat{m}^{R,OIL}} \cdot \hat{i}_t^R + c_4^{\hat{m}^{R,OIL}} \cdot \hat{c}_t^R \\ & - c_5^{\hat{m}^{R,OIL}} \cdot \hat{p}_t^{R,MOIL} - c_6^{\hat{m}^{R,OIL}} \cdot \left(\hat{p}_t^{R,MOIL} - \hat{p}_t^{R,I}\right) + c_7^{\hat{m}^{R,OIL}} \cdot \hat{g}_t^R \\ & - c_8^{\hat{m}^{R,OIL}} \cdot \left(\tau_t^{MOIL} - \tau_{t+1}^{MOIL}\right) + \varepsilon_t^{\hat{m}^{R,NOIL}}\end{aligned} \tag{74}$$

$$\hat{m}_t^{R,OIL,E} = \hat{m}_{t+1}^{R,OIL} \tag{75}$$

$$\begin{aligned}\Delta\overline{m}_t^{R,OIL} = {} & c_1^{\Delta\overline{m}^{R,OIL}} \cdot \Delta\overline{m}_{t-1}^{R,OIL} \\ & + \left(1 - c_1^{\Delta\overline{m}^{R,OIL}}\right) \cdot \left(\Delta\overline{p}_t^{R,Y} + \Delta\overline{y}_t^R - \Delta\overline{p}_t^{R,MOIL}\right) - c_2^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{M}_t^{OIL,Y}\right. \\ & - M^{OIL,Y,SS} - c_3^{\Delta\overline{m}^{R,OIL}} \cdot \overline{M}_t^{R,OIL,D} - c_4^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{C}_t^Y - C^{Y,SS}\right) \\ & - c_5^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{I}_t^Y - I^{Y,SS}\right) - c_6^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{X}_t^{NNR,Y} - X^{NNR,Y,SS}\right) \\ & - c_7^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{X}_t^{NR,Y} - X^{NR,Y,SS}\right) - c_8^{\Delta\overline{m}^{R,OIL}} \cdot \left(\overline{G}_t^Y - G^{Y,SS}\right) \\ & \left. + c_9^{\Delta\overline{m}^{R,OIL}} \cdot \left(\tau_t^{MOIL} - \tau^{MOIL,SS}\right)\right) + \varepsilon_t^{\Delta\overline{m}^{R,NOIL}}\end{aligned} \tag{76}$$

$$\overline{M}_t^{R,OIL,D} = c_1^{\overline{M}^{R,OIL,D}} \cdot \overline{M}_{t-1}^{R,OIL,D} - \Delta\overline{p}_t^{R,MOIL} + \Delta\overline{p}^{R,MOIL,SS} + \varepsilon_t^{\overline{M}^{R,NOIL,D}} \tag{77}$$

$$\overline{M}_t^{OIL,Y} = \frac{100 \cdot \overline{P}_t^{R,MOIL} \cdot \left(1 + \frac{\tau_t^{MOIL}}{100}\right) \cdot \overline{M}_t^{R,OIL}}{\overline{Y}^N} \tag{78}$$

$$M_t^{OIL,Y} = 100 \cdot \frac{P_t^{MOIL} \cdot \left(1 + \frac{\tau_t^{MOIL}}{100}\right)}{P_t^Y} \cdot \frac{M_t^{R,OIL}}{Y_t^R} \tag{79}$$

$$m_t^{R,OIL} = \hat{m}_t^{R,OIL} + \overline{m}_t^{R,OIL} \tag{80}$$

$$\Delta m_t^{R,OIL} = m_t^{R,OIL} - m_{t-1}^{R,OIL} \tag{81}$$

$$m_t^{R,OIL} = 100 \cdot \log\left(M_t^{R,OIL}\right) \tag{82}$$

What comes next

Equations 83 onwards, which are not reproduced here, carry the price block, the external balance and financing conditions, the monetary policy rule and the fiscal block. The real block above is where they all act: a policy rate enters through the real rate terms in equations (16) and (26), a tariff through equation (66), a resource royalty through equation (59), and public investment through the potential growth equation (7). Reading this section first makes the transmission channels in the simulations of Section 12 legible as movements in specific equations rather than as reduced-form responses.

D. The Julia code, mirrored to the model

The implementation is a Julia port built on `ModelBaseEcon`, `StateSpaceEcon` and `TimeSeriesEcon`.

D.1 Layout

```
model/FINEX_canonical.model model source (2,672 lines) - Sec 2
julia/src/FINEX.jl finex_model, solve_steadystate!, irf
julia/src/finex_generated.jl IT_FLOAT equations
julia/src/finex_it_managed.jl IT_MANAGED variant
julia/src/finex_fix_closed_part.jl FIX_CLOSED_PART variant

julia/scripts/country_common.jl shared: panel, steady state, inversion
julia/scripts/smm_country.jl SMM estimation - Sec 4
julia/scripts/simulate_country.jl projection - Sec 5
julia/scripts/irf_country.jl impulse responses, 83 shocks - Sec 6,7
julia/scripts/scenario_run.jl TOML-driven scenario engine - Sec 8
julia/scripts/decomp_country.jl shock decomposition - Sec 9

scripts/download_country_data.py data acquisition - Sec 3
scripts/build_country_observables.py observable construction - Sec 3
scripts/make_report_figures.py the figures in this report
```

`country_common.jl` is included by every downstream script, so projections, impulse responses, scenarios and the decomposition all come off *exactly the same model instance*.

D.2 The estimator, in outline

```julia
# Autocovariances are linear in the shock variances, so for any trial vector of
# behavioural coefficients the variances solve exactly by non-negative least
# squares; only the coefficients are searched.
function objective(theta)
    set_behavioural!(M, theta)
    H = irf_matrix(M, L) # h_j(l) for every shock j, lag l
    A = build_loadings(H, NPAD) # A_j(tau) = sum_l h_j(l) h_j(l+tau)'
    sigma2 = nnls(A, gamma_data) # Lawson-Hanson active set - exact
    return weighted_distance(A * sigma2, gamma_data)
end

# Not cosmetic: c1 + c2 >= 1 puts a root on or outside the unit circle and the
# model becomes explosive, while the SMM objective reports a finite value.
function stability_cap(name, lo, hi)
    startswith(String(name), "C1_") || return hi
    eq = String(name)[4:end]
    eq in ("L_CONS_GAP", "L_INV_GAP", "L_EXP_NNR_GAP") || return hi
    c2 = getproperty(M, Symbol("C2_", eq))
    return min(hi, max(lo + 0.05, 1 - c2 - 0.05))
end
```

D.3 The inversion, in outline

```
# Exogenise each observable and endogenise the structural shock in its OWN
# equation, period by period, so the counts balance. A series need only cover
# part of the window - which lets the policy rate, whose data end in 2019,
# still anchor the monetary block.
function run_inversion(pairs)
    p = Plan(M, RNG)
    for (v, s) in pairs, y in havey(v)
        exogenize!(p, Symbol(v), y2t(y))
        endogenize!(p, s, y2t(y))
    end
    data = zerodata(M, p)
    for (v, _) in pairs, y in havey(v)
        data[y2t(y), Symbol(v)] = D[v][y] - ss_level(v) # deviation from SS
    end
    return simulate(M, p, data; deviation = true, anticipate = false,
                    fctype = fcnatural)
end
```

Observables are added one at a time and any that makes the system singular is skipped and reported, so a single awkward series does not cost the rest.

D.4 Scenario specification

Scenarios are TOML. One file carries the calibration and the shocks, so a scenario and the calibration it assumed cannot drift apart.

```
[run]
fcast_end = 2035
variant = "IT_FLOAT" # or IT_MANAGED, FIX_CLOSED_PART

[calibration] # optional overrides
infl = 5.0
debt = 50.0

[[scenario]]
name = "Tightening 2027, consolidation 2030"
anticipate = false # agents do not see it coming
attribution = true # run each shock alone + report the interaction

  [[scenario.shock]]
  name = "RES_RS"
  years = [2027]
  size = 1.0
  unit = "sd" # or "raw"
  rho = 0.0 # 0 = one-off; >0 = geometric decay

  [[scenario.shock]]
  name = "RES_DEF_S_RAT"
  years = [2030]
  size = -1.0
  unit = "sd"

# Conditioning paths: exogenise a variable over future years and endogenise an
# instrument, so the model finds the shocks that deliver the requested path.
```

```
# The counts must balance or the system is under- or over-determined.
[[scenario.condition]]
variable = "L_RP_OIL_GAP"
years = [2026, 2027, 2028]
values = [25.0, 20.0, 15.0]
instrument = "RES_ABS_L_RP_OIL_GAP"
```

E. How to run it

E.1 Prerequisites

Julia 1.10+ and Python 3.10+.

```
cd julia
julia --project=. -e 'using Pkg; Pkg.instantiate()'
```

On machines where Windows Application Control blocks Julia's precompiled package images, every Julia call must pass `-pkgimages=no`; without it the scripts fail to start.

E.2 The full pipeline

```
# 1. data
python scripts/download_country_data.py UGA
python scripts/build_country_observables.py UGA

# 2. estimation (~45 min)
cd julia
FINEX_COUNTRY=uganda julia --project=. --pkgimages=no scripts/smm_country.jl

# 3. projection (~15 min)
FINEX_COUNTRY=uganda julia --project=. --pkgimages=no scripts/simulate_country.jl

# 4. impulse responses (~20 min)
FINEX_COUNTRY=uganda julia --project=. --pkgimages=no scripts/irf_country.jl

# 5. shock decomposition (~40 min)
FINEX_COUNTRY=uganda julia --project=. --pkgimages=no scripts/decomp_country.jl

# 6. a scenario
FINEX_SCENARIO=../scenarios/uganda_tightening_consolidation.toml \
FINEX_COUNTRY=uganda julia --project=. --pkgimages=no scripts/scenario_run.jl
cd ..

# 7. figures
python scripts/make_report_figures.py
```

E.3 Environment variables

Table 29: Environment variables recognised by the pipeline

Variable	Effect
`FINEX_COUNTRY`	Which country to run (default `uganda`)
`FINEX_CONFIG`	TOML overriding calibration, horizon, policy regime
`FINEX_OUTDIR`	Redirect output so concurrent runs do not collide
`FINEX_PARFILE`	Use a specific estimation vintage
`FINEX_SCENARIO`	The scenario TOML for `scenario_run.jl`
`SMM_TIME_LIMIT`	Seconds for the SMM outer search (default 3600)

Notes: All are optional. With none set, every script reproduces its command-line behaviour exactly.

E.4 The web application

An interactive front end provides a dashboard, projections with a horizon control, an impulse-response explorer over all 83 shocks with a live shock-size slider, a scenario builder, the decomposition, and Excel export of every view.

```
.\run.ps1 -Port 8080 -NoWait # returns immediately; survives terminal close
.\run.ps1 -Command status -Port 8080
.\run.ps1 -Command stop
```

E.5 Two things to check first if a run misbehaves

- **If the projection diverges**, check the stability restriction of Section 5.4 before anything else: an unrestricted investment persistence above 0.4 makes the model explosive and the SMM objective will not tell you.

- **If the projection shows a large, persistent real appreciation or depreciation**, compare the relative-price trends with the data. FINEX ships generic values that suit Uganda but are three to four points a year away from some countries, and the model unwinds that wedge as a real exchange-rate correction.

F. Solution method and numerical diagnostics

The model is solved with a stacked-time algorithm. This appendix records how, and what can go wrong — because in a model of this size the numerical properties are not separable from the economic ones.

F.1 The stacked-time solution

FINEX contains both lags and model-consistent expectations, so it cannot be solved recursively forward. The stacked-time method instead treats the entire simulation range as one large system: all variables at all dates are unknowns simultaneously, and the model's equations at every date are stacked into a single nonlinear system solved by Newton's method with a sparse Jacobian.

Three consequences follow that shape everything in this report.

- **Terminal conditions matter.** The system needs a boundary at the end of the range. Every simulation here uses *natural* terminal conditions — the model's own steady-state growth relations imposed at the end — and runs **60 years past the reporting horizon**, so the boundary sits far enough away not to contaminate 2025–2035.

- **Anticipation is a modelling choice implemented here.** Because the whole path solves at once, a shock dated 2030 can influence 2027 if agents are assumed to see it coming. Every scenario in this report is run *unanticipated* unless stated, matching the historical inversion.

- **Solutions can fail rather than degrade.** Newton's method either converges or does not. A failure is reported as a failure, not as an approximate answer — which is the desired behaviour for an infeasible conditioning path (Section 6.6).

F.2 Exogenising and endogenising

The inversion of Section 6.2 and the conditional forecasts of Section 6.6 both work by swapping which variables are unknowns. Exogenising a variable at a date fixes its value and removes it from the unknowns; endogenising a shock at a date adds that shock to the unknowns.

The counts must balance exactly. One endogenised shock-date per exogenised variable-date. With fewer, the system is over-determined and Newton's method fails; with more, it is under-determined and the Jacobian is singular. This is not a soft constraint that degrades gracefully — it is the difference between a solvable system and an unsolvable one, which is why the tooling checks it before submission rather than diagnosing it afterwards.

The inversion adds observables **one at a time** and drops any that makes the system singular, reporting which. For Uganda 25 of the 26 candidate pairs are retained; the procedure is run separately for each country and the retained set is not identical across them, because a series that is informative in one panel may be short or broken in another. This matters because a single awkward series — one with a structural break, or one whose shock is already pinned by another equation — would otherwise cost the entire inversion.

F.3 The nonlinearities, and where they bite

FINEX is not a linear model, and three features are responsible:

- **The effective lower bound** on the policy rate, equation (29).
- **Debt accumulation**, which is a product of stocks, interest rates and the exchange rate rather than a sum.
- **The risk-appetite terms** of equations (25)–(27), which are exponential.

Near the steady state all three are approximately linear, which is why the attribution exercise of Section 12.1 finds an interaction term of only 0.030. Far from it they are not, and three consequences follow for how results in this report must be read.

Impulse responses do not scale indefinitely. Doubling a shock does not exactly double the response once the bound binds or the risk terms steepen. Beyond roughly two standard deviations, scaling a stored response is no longer exact.

Shock contributions are not additive in general. This is why the decomposition of Section 13 uses cumulative group-wise re-simulation rather than a linear convolution: it is the only method that makes contributions sum to the total, which they do to 2×10^{-12}.

Starting conditions change the answer. The same shock from a worse starting position produces a larger response. A linear model cannot express this; it is the main reason FINEX is worth the extra computational cost for a country with meaningful debt or reserve vulnerabilities.

F.4 Diagnostics applied to every run

Four checks run on every simulation in this project. All four were added in response to a failure that they would have caught.

Table 30: Numerical and economic diagnostics

Check	Threshold	Why it exists
Steady state solves	Residual $< 10^{-4}$ on the country's own calibration	A fallback to the shipped generic calibration produces a model that is no longer about the country
Moment fit	$\text{corr} \geq 0.10$	An estimate can produce a complete parameter set at a fit near zero, where the parameters carry no information
Terminal convergence	$\lvert x_{2035} - x^{SS} \rvert \leq 25$	An unrestricted persistence estimate produced 252% growth and -135% debt while scoring a better objective
Path boundedness	$\lvert x_t - x^{SS} \rvert \leq 500$ anywhere	Catches divergence that has not yet reached the horizon

Notes: Thresholds are deliberately loose: they are designed to catch divergence and non-identification, not to certify accuracy. Uganda clears all four with a terminal gap of 1.24 and a moment fit of 0.783, and South Africa with 2.17 and 0.774. **Kenya converges (3.44) but its moment fit of 0.191 falls below the 0.40 threshold at which an estimate is called identified**, which is why every Kenyan result in this report carries that qualification.

The second and third checks are independent, and must both be applied. Fit does not imply stability: the unrestricted Uganda estimate scored a *better* moment objective than the restricted one and produced a divergent projection. Stability does not imply information either: a model can converge cleanly on a moment fit near zero, which is noise. Only a model that passes both is worth interpreting.

F.5 A practical note on reproducibility

Every run in this report is reproducible from the repository, and three conventions make that true rather than aspirational.

- **The configuration travels with the output.** Each run writes its own configuration file alongside its results, so a directory of numbers always states what produced it.
- **Scenarios are declarative.** A scenario is a TOML file specifying shocks, dates, sizes, persistence, regime and calibration in one place, so a scenario and the calibration it assumed cannot drift apart.
- **Estimation is seeded.** The stochastic elements — the bootstrap and any draw-based exercise — use fixed seeds, so a re-run reproduces the reported numbers exactly rather than approximately.

G. The web application

The model is delivered as a web application as well as a command-line pipeline, and this appendix shows what that application does. The motivation is stated in Section 2: a model that runs only on the machine of the person who built it has a reach of one. What follows is the same estimated model reached through a browser, by a user who need not install Julia, read the code, or know that either exists.

Nothing in the browser solves the model. The web process serves stored results and queues work; solving happens in separate worker processes, two of which can hold a model in memory at once on the machine used here. This is why every page distinguishes what it can answer immediately — anything that rescales a stored result — from what requires a run, and says how long that run will take. The same process serves Uganda, Kenya and South Africa, selected on the page.

Two conventions run through every screen. First, **results are never shown without their verdict**: the identification and convergence badges of Section 5 appear on each page that reports a number, because the two ways an estimate fails are invisible in the output itself. Second, **every view has an address**. A page, a tab within it, a selected question, a chosen run, a filtered list and a fully specified quantitative query are all expressible in the address bar, so a colleague receives a link to the result rather than instructions for reproducing the clicks. The screenshots in this appendix were themselves taken by visiting those addresses.

G.1 The dashboard

The page opens on a verdict rather than on a result. A reader who sees *identified* and *converges* knows the estimate behind the numbers cleared the same two tests this paper applies throughout, and a reader who does not see them knows to stop. This is deliberate: the two failure modes described in Section 5 — a parameter set that fits nothing, and a projection that runs away — are invisible in the output itself.

Figure 42: The opening page: the verdict, eight headline numbers, and what they are composed of

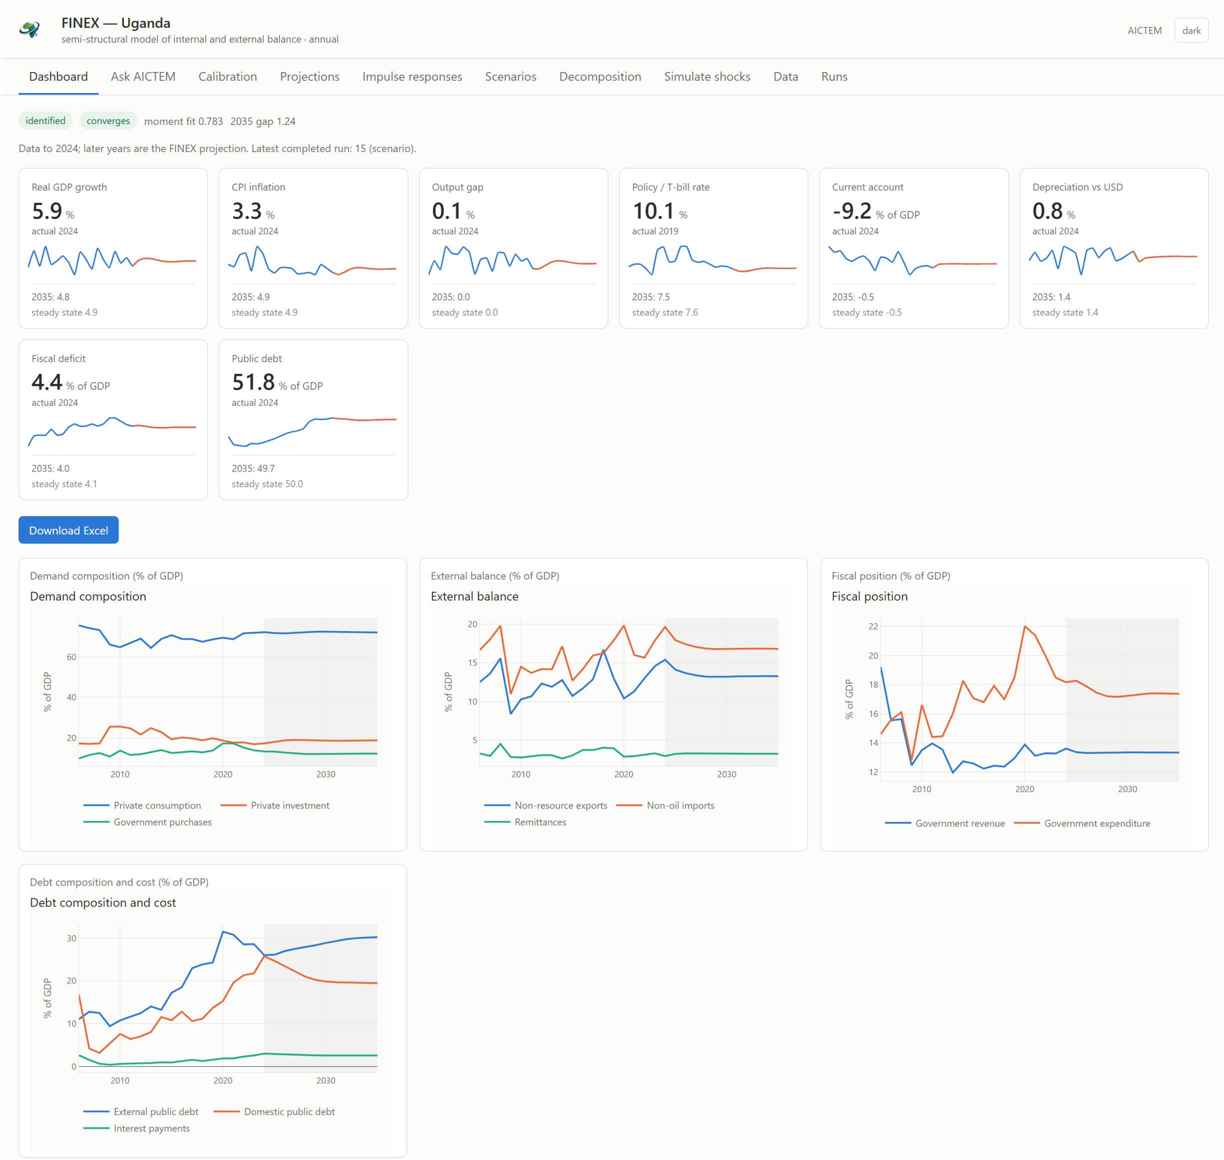

The badges on the top line are the gate of Section 5: *identified* reports a moment fit of 0.783 and *converges* that the projection settles 1.24 units from its steady state by 2035. Each tile carries the latest observation, the model's 2035 value and its steady state, so a number is never shown without the anchor it is moving towards. The panels below repeat no series that already has a tile; they carry what a single number cannot, which is a composition or a balance.

G.2 The question bank

The bank is filtered by topic, by originating agent, by status and by free text. **The status column is the honest part of the design.** A model asked 1,050 questions will answer all of them if nothing stops it; marking 496 as partial and dropping 120 outright is a statement about the boundary of the tool, made before a user reaches it rather than after.

Figure 43: The question bank: 930 questions, each marked by what the model can actually answer

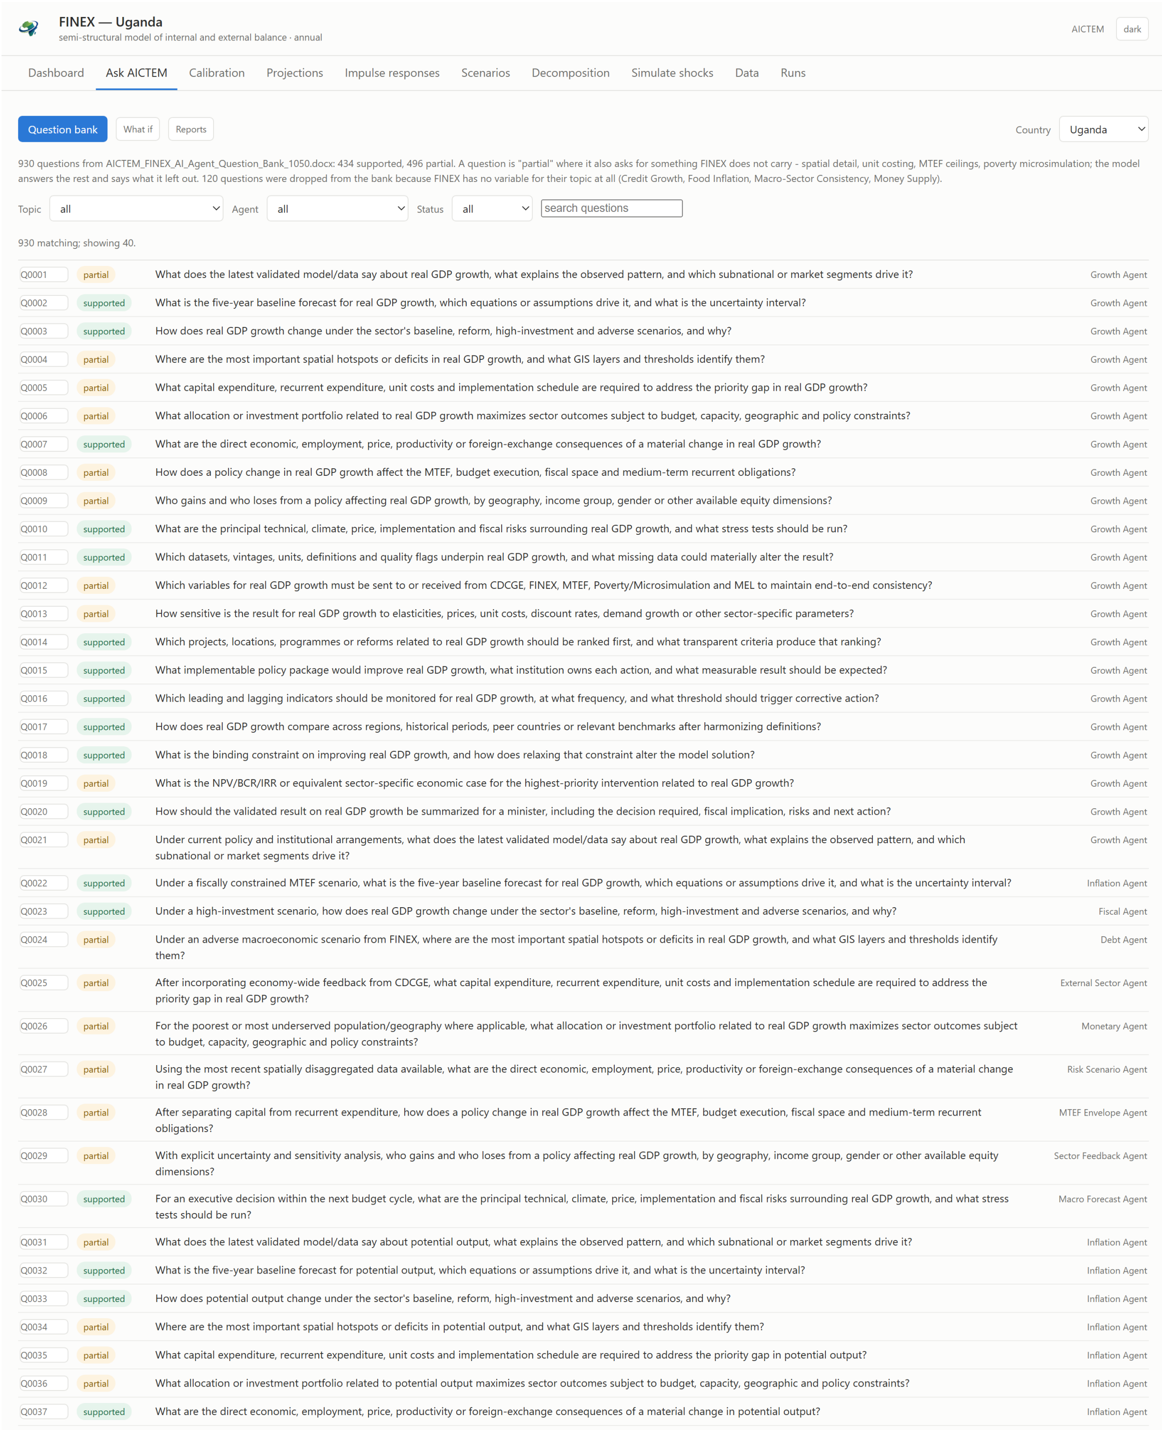

Each row carries a status. *Supported* means FINEX holds every variable the question needs; *partial* means the question also asks for something the model does not carry — spatial detail, unit costing, medium-term expenditure ceilings, poverty microsimulation — in which case the model answers the part it can and states what it left out. The 120 questions with no FINEX variable at all were removed from the bank rather than answered badly.

G.3 How a question is answered

Selecting a question identifier produces a structured answer rather than a paragraph. The order is fixed — headline, findings, attribution, validation, assumptions, charts — so a reader learns where to look once and then always knows. The charts carry their own notes, and the attribution chart repeats the warning that contributions are large, offsetting and order-dependent.

Figure 44: A question opened: executive answer, key findings, attribution, validation and provenance

The answer is assembled from the stored results, not from a fresh solve. *What is driving it* reads the group decomposition of Section 13 for the year in question. *Validation and confidence* repeats the gate. *Assumptions and provenance* names the estimation window, the projection convention and the analysis identifier, which is reproducible: the same identifier regenerates the same numbers.

G.4 Quantitative queries

The *What if* tab answers a numerical question without queueing a run. A user picks a preset or builds a combination of shocks, and the page scales the stored impulse responses and reports peak deviations and paths in under a second. **The honesty of this shortcut is stated where it is used, not in a manual**: every answer carries the caveat that says whether the scaling was still exact at the size requested.

Figure 45: A quantitative query: any combination of shocks, scaled and read off the stored impulse responses

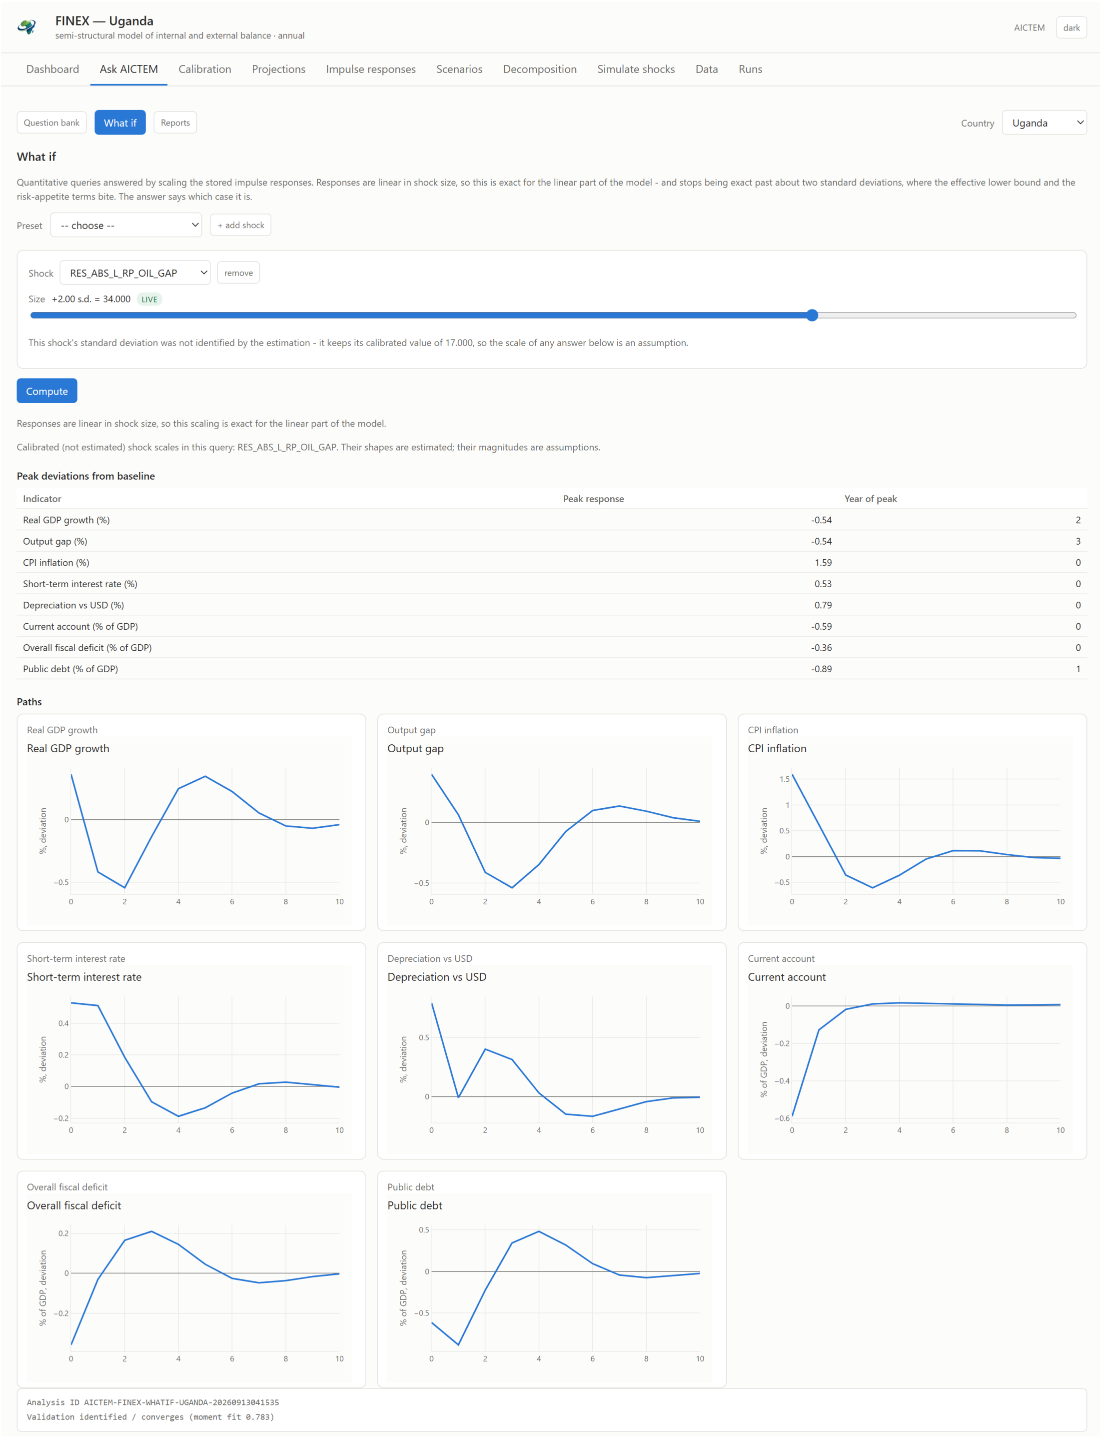

The answer is exact for the linear part of the model, because responses are linear in shock size; past roughly two standard deviations the effective lower bound and the risk-appetite terms bite, and the page says so. The note above the table repeats that this shock's scale is calibrated rather than estimated, so the size of every number below it is an assumption — the point made throughout Section 11.

G.5 Reports

The *Reports* tab turns one question into a document under a chosen profile: forecast, diagnostic, scenario, risk and resilience, or executive policy brief. The technical provenance block at the foot names the model, the estimation method and window, the scenario and the generation time, so a printed copy can be traced back to the run that produced it.

Figure 46: Report generation: five profiles over the same frozen evidence package

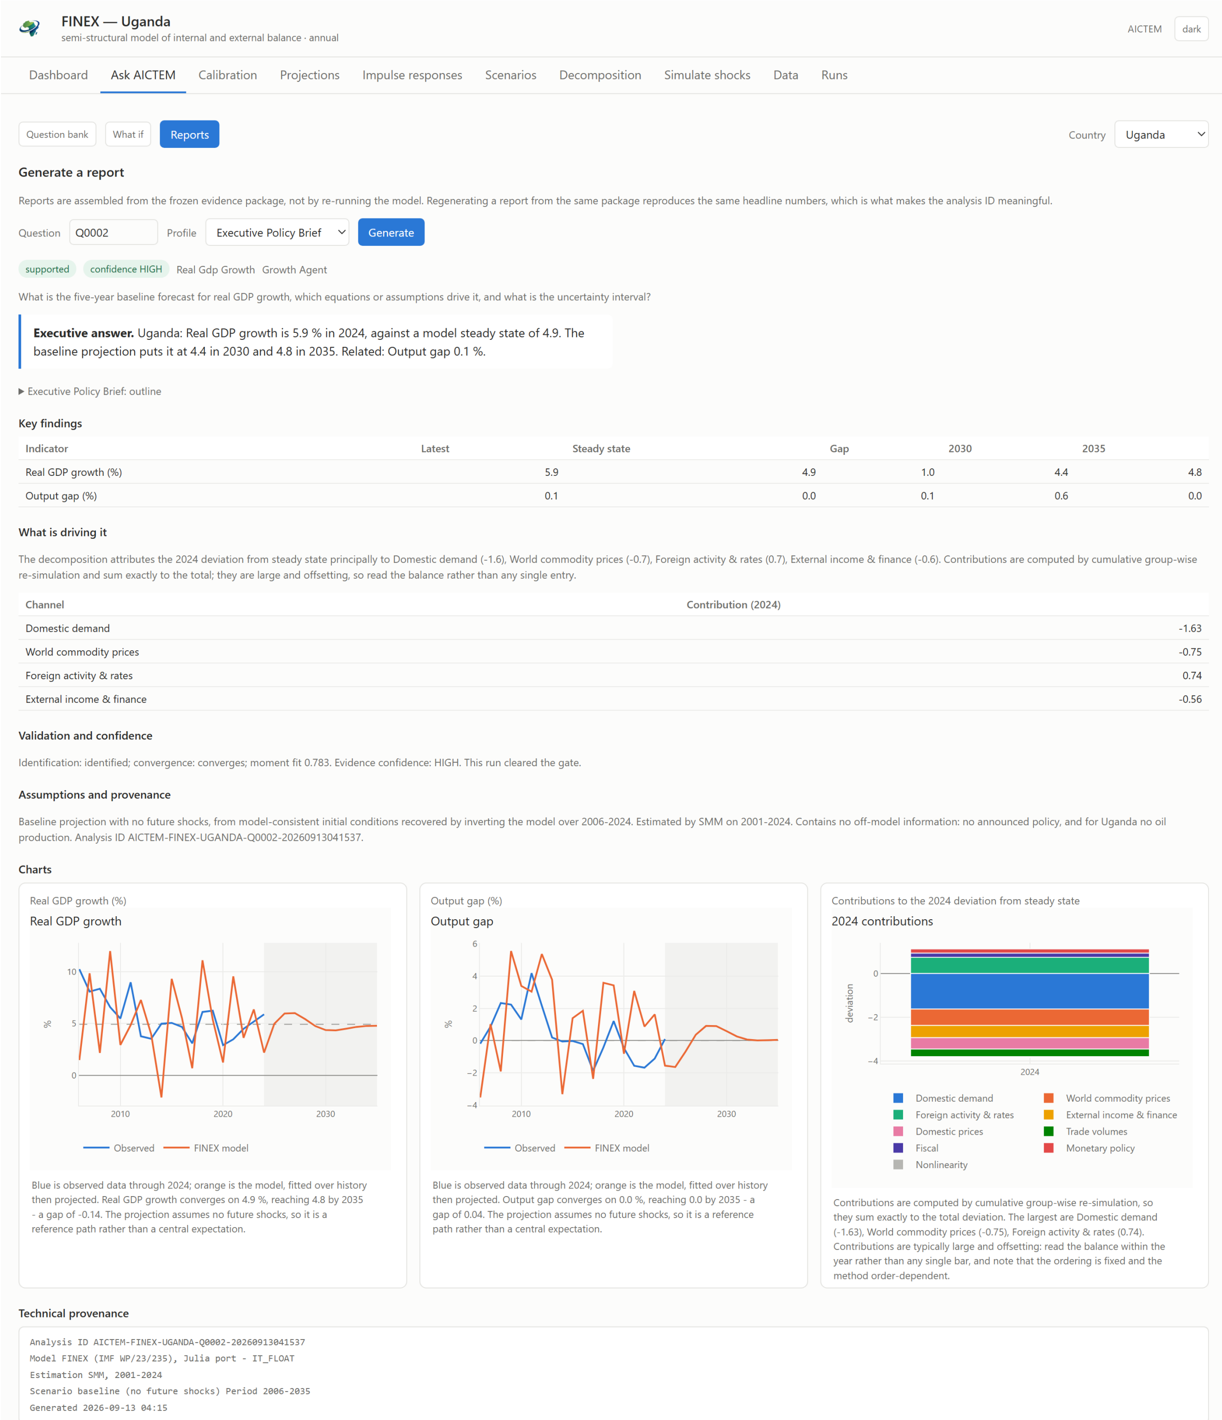

FINEX — Uganda
semi-structural model of internal and external balance · annual

AICTEM dark

Dashboard Ask AICTEM Calibration Projections Impulse responses Scenarios Decomposition Simulate shocks Data Runs

Question bank What if Reports Country Uganda

Generate a report

Reports are assembled from the frozen evidence package, not by re-running the model. Regenerating a report from the same package reproduces the same headline numbers, which is what makes the analysis ID meaningful.

Question Q0002 Profile Executive Policy Brief Generate

supported confidence HIGH Real Gdp Growth Growth Agent

What is the five-year baseline forecast for real GDP growth, which equations or assumptions drive it, and what is the uncertainty interval?

Executive answer. Uganda: Real GDP growth is 5.9 % in 2024, against a model steady state of 4.9. The baseline projection puts it at 4.4 in 2030 and 4.8 in 2035. Related: Output gap 0.1 %.

▸ Executive Policy Brief: outline

Key findings

Indicator	Latest	Steady state	Gap	2030	2035
Real GDP growth (%)	5.9	4.9	1.0	4.4	4.8
Output gap (%)	0.1	0.0	0.1	0.6	0.0

What is driving it

The decomposition attributes the 2024 deviation from steady state principally to Domestic demand (-1.6), World commodity prices (-0.7), Foreign activity & rates (0.7), External income & finance (-0.6). Contributions are computed by cumulative group-wise re-simulation and sum exactly to the total; they are large and offsetting, so read the balance rather than any single entry.

Channel	Contribution (2024)
Domestic demand	-1.63
World commodity prices	-0.75
Foreign activity & rates	0.74
External income & finance	-0.56

Validation and confidence

Identification: identified; convergence: converges; moment fit 0.783. Evidence confidence: HIGH. This run cleared the gate.

Assumptions and provenance

Baseline projection with no future shocks, from model-consistent initial conditions recovered by inverting the model over 2006-2024. Estimated by SMM on 2001-2024. Contains no off-model information: no announced policy, and for Uganda no oil production. Analysis ID AICTEM-FINEX-UGANDA-Q0002-20260913041537.

Charts

Real GDP growth (%)
Real GDP growth

Blue is observed data through 2024; orange is the model, fitted over history then projected. Real GDP growth converges on 4.9 %, reaching 4.8 by 2035 - a gap of -0.14. The projection assumes no future shocks, so it is a reference path rather than a central expectation.

Output gap (%)
Output gap

Blue is observed data through 2024; orange is the model, fitted over history then projected. Output gap converges on 0.0 %, reaching 0.0 by 2035 - a gap of 0.04. The projection assumes no future shocks, so it is a reference path rather than a central expectation.

Contributions to the 2024 deviation from steady state
2024 contributions

Contributions are computed by cumulative group-wise re-simulation, so they sum exactly to the total deviation. The largest are Domestic demand (-1.63), World commodity prices (-0.75), Foreign activity & rates (0.74). Contributions are typically large and offsetting: read the balance within the year rather than any single bar, and note that the ordering is fixed and the method order-dependent.

Technical provenance

```
Analysis ID AICTEM-FINEX-UGANDA-Q0002-20260913041537
Model FINEX (IMF WP/23/235), Julia port - IT_FLOAT
Estimation SMM, 2001-2024
Scenario baseline (no future shocks) Period 2006-2035
Generated 2026-09-13 04:15
```

The profiles differ in what they foreground, not in what they compute. Because the report is assembled from a stored evidence package rather than by re-solving, regenerating it from the same package reproduces the same headline numbers — which is what makes the analysis identifier at the foot of the page worth quoting in correspondence.

G.6 Projections

Every headline observable is drawn on the same convention, which makes the fit inspectable rather than asserted. Where the orange line tracks the blue one the model is reproducing history; where it does not — the real exchange-rate gap is the clearest case — the reader sees the discrepancy instead of reading a summary statistic that averages it away. The horizon is adjustable and the whole table exports to a spreadsheet.

Figure 47: The baseline projection: the first twelve of sixteen headline observables, model against data

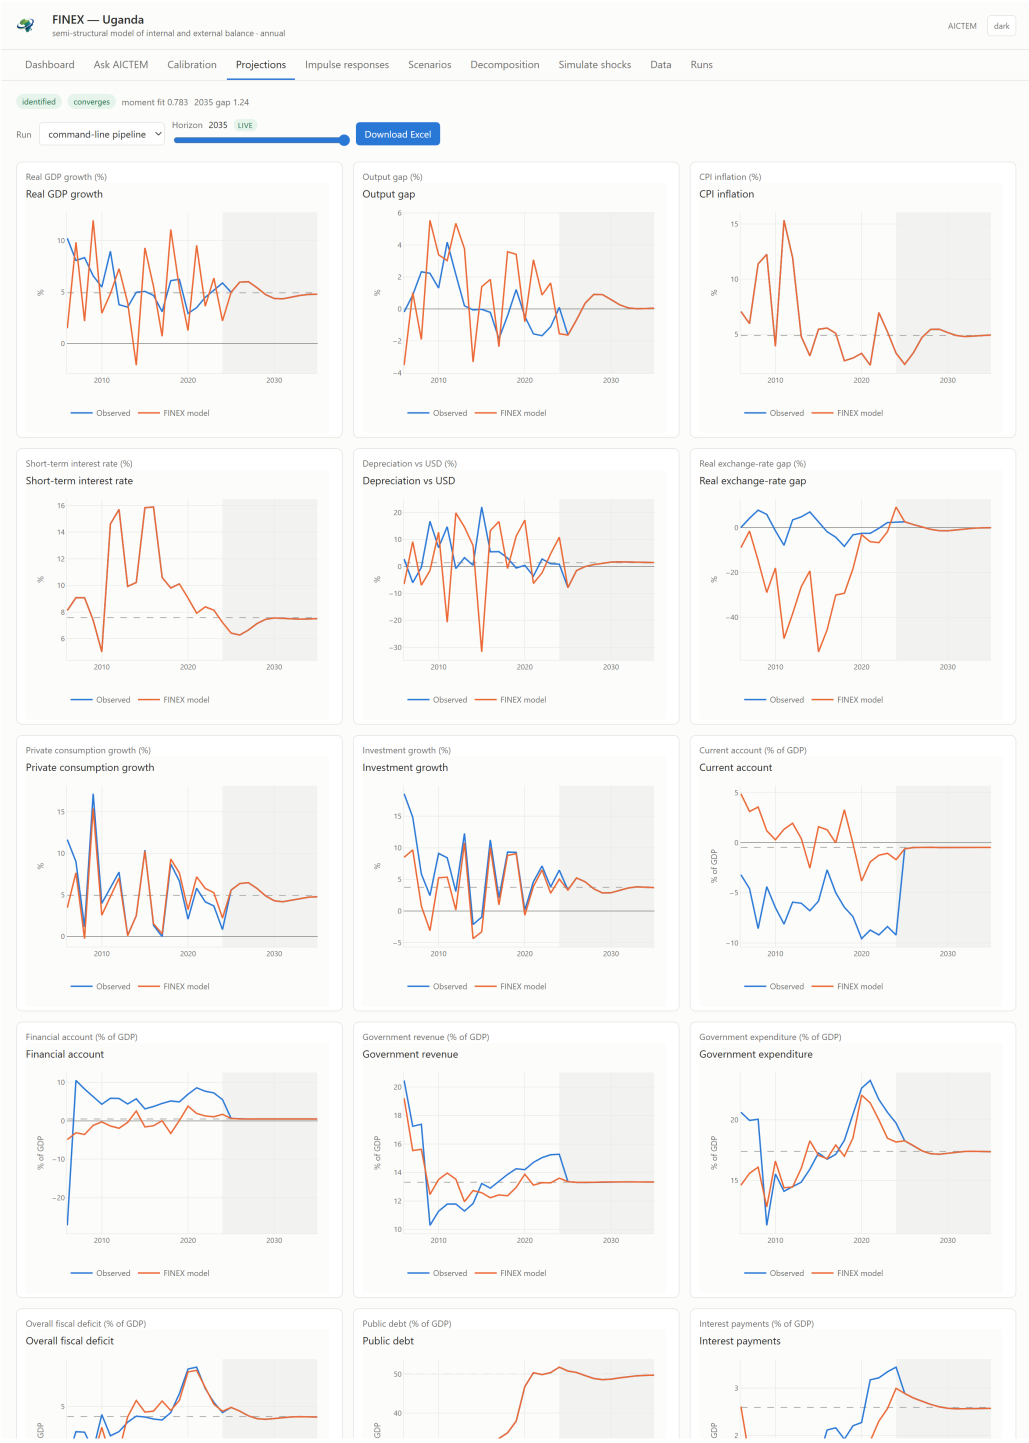

Blue is the observed panel through 2024 and orange is the model, fitted over history and projected beyond it; the shaded band is the projection and the dashed line the steady state. The run selector names which set of results is on the page — here the command-line pipeline, which is the baseline this paper reports. The page continues below the crop with the remaining four observables.

G.7 Impulse responses

Any of the 83 structural shocks can be selected and every headline observable responds on one screen. **The page distinguishes the shape of a response from its size.** The shape comes from the estimated model; the size comes from a standard deviation which, for 64 of the 83 shocks, the estimation could not identify — and the page says which case the reader is looking at rather than leaving both to look alike.

Figure 48: Impulse responses to a monetary tightening, scaled live

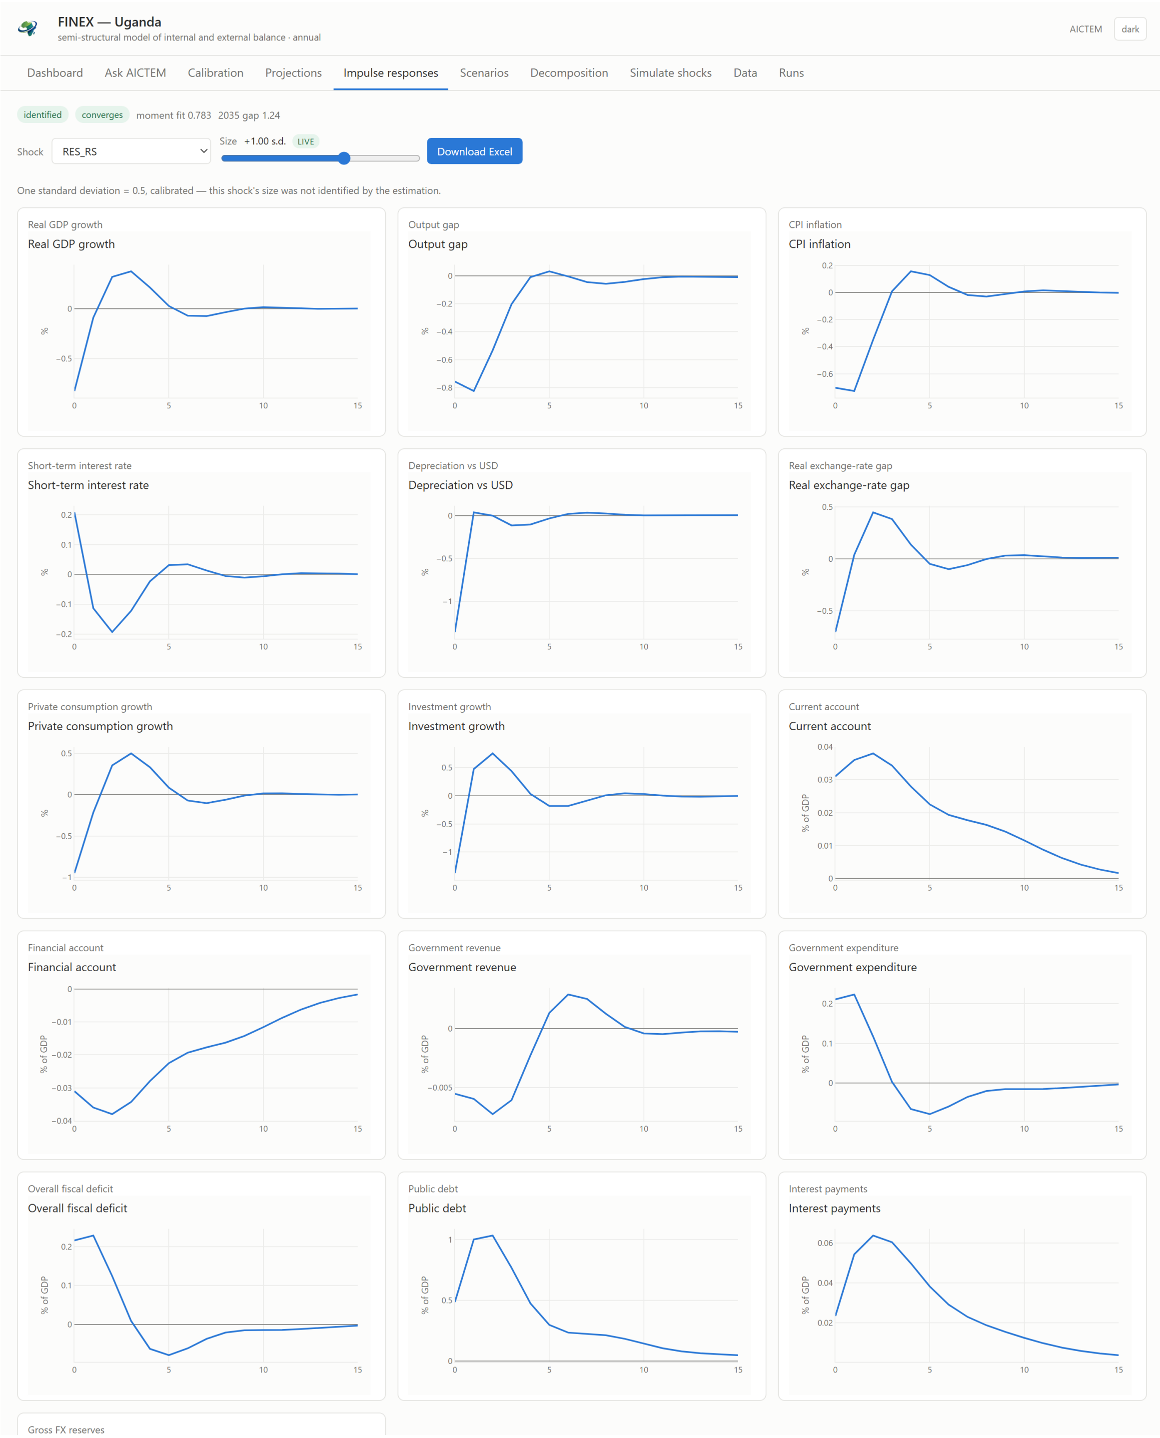

The size slider rescales the stored response instantly, which is exact while the model is linear. Past the threshold printed on the page the application says the scaling can no longer be trusted and asks for a queued run instead. The note beneath the selector states that this shock's standard deviation was calibrated, not estimated.

G.8 Scenarios and attribution

Scenarios are solved by a worker process and stored, so the page shows a result rather than computing one. **The attribution is the part worth having.** A combined scenario that moves inflation by two points says little on its own; the same scenario decomposed into each shock's own contribution plus an explicit interaction term says where the two points came from, and whether the model's nonlinearities were active.

Figure 49: A scenario: baseline, scenario, each shock alone, and the interaction between them

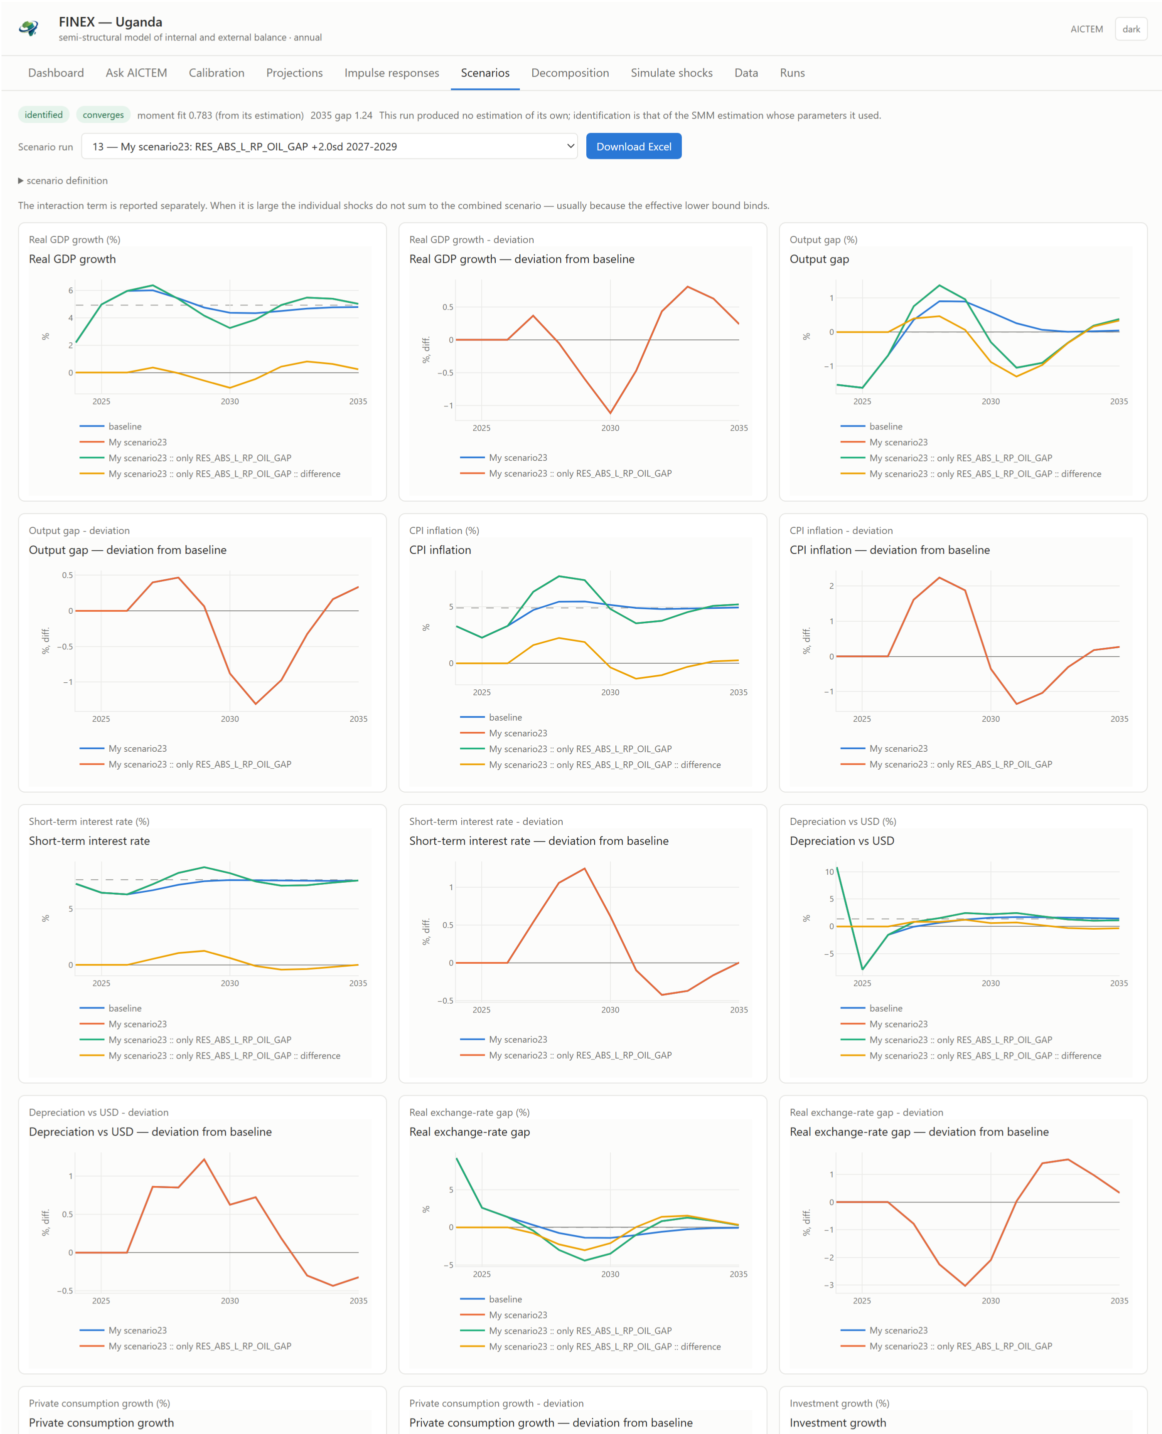

A world oil price two standard deviations above baseline for 2027–2029. The left column is the level, the right the deviation from baseline. Where a scenario carries more than one shock, each is also run alone and the residual is reported as an interaction: when it is large the shocks do not add up, usually because the effective lower bound binds.

G.9 Building a scenario

The builder is where a user composes an experiment. The preset shown combines a 2027 monetary tightening with a 2030 fiscal consolidation, the policy mix discussed in Section 9. Rows carry the same warning as elsewhere when a shock's size was calibrated rather than estimated. Pressing *Run* *scenario* queues the job, and the results appear beneath when the worker finishes.

Figure 50: Building a scenario: dated shocks, sizes in standard deviations, and persistence

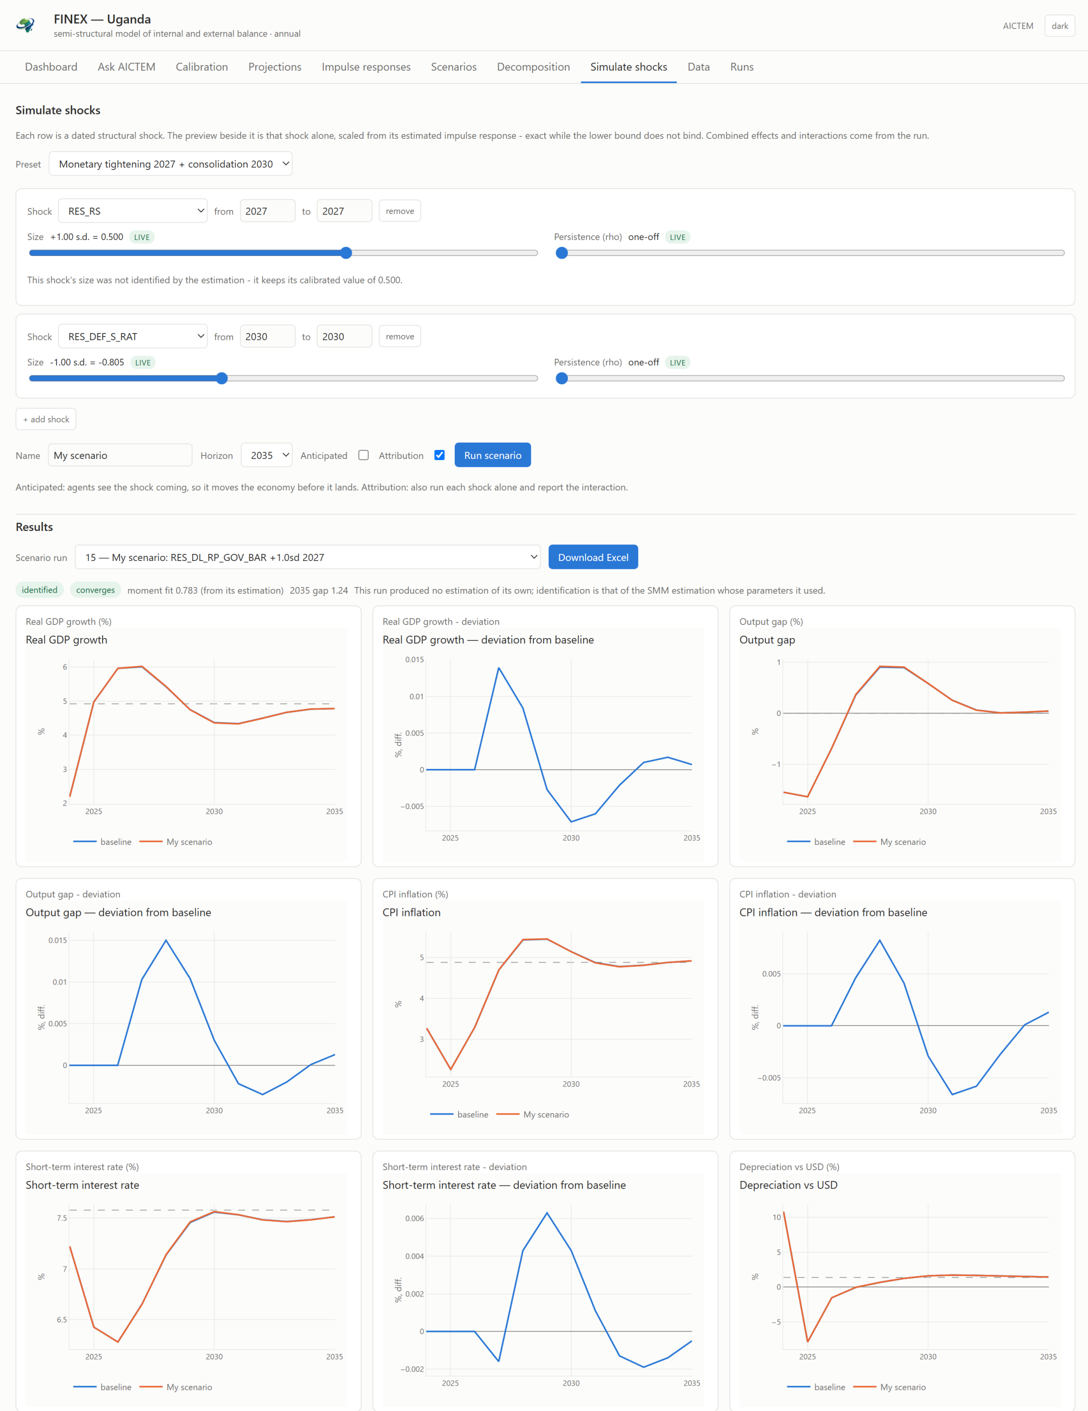

Each row is one structural shock with a start year, an end year, a size in standard deviations and a persistence. *Anticipated* decides whether agents see the shock coming, which moves the economy before it lands. *Attribution* asks the worker to run each shock alone as well, which is what produces the decomposition in Figure 49.

G.10 Historical decomposition

The decomposition is available for every observable through the selector, not only for the ones this paper prints. This is the page a user reaches for when asking why a particular year looked as it did, and it is also the page most easily over-read — which is why the caveat travels with the chart in the application exactly as it does in Section 13.

Figure 51: The historical decomposition, by shock group

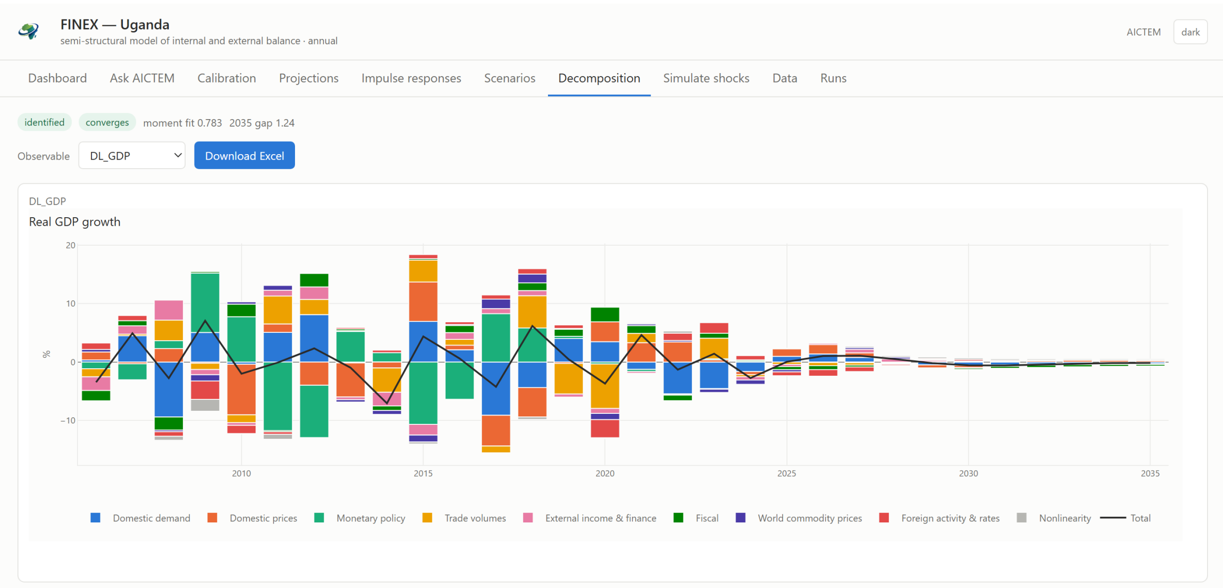

Each bar is one year, split into the eight groups of Section 13, with the white line the total deviation from steady state. Contributions are computed by cumulative group-wise re-simulation and sum exactly to the total. They are large and offsetting and the group ordering is fixed, so the balance within a year is the reading, not any single bar.

G.11 Calibration, and what a change costs

The distinction this page draws is the one that keeps the application honest under interaction. A user dragging a slider has no way of knowing whether the number that moved was computed or interpolated, so the page labels every control before it is touched. Changing the inflation target or the debt anchor queues a run and says what it will cost in time; changing a shock size does not.

Figure 52: Calibration: what changes instantly and what requires the model to be re-solved

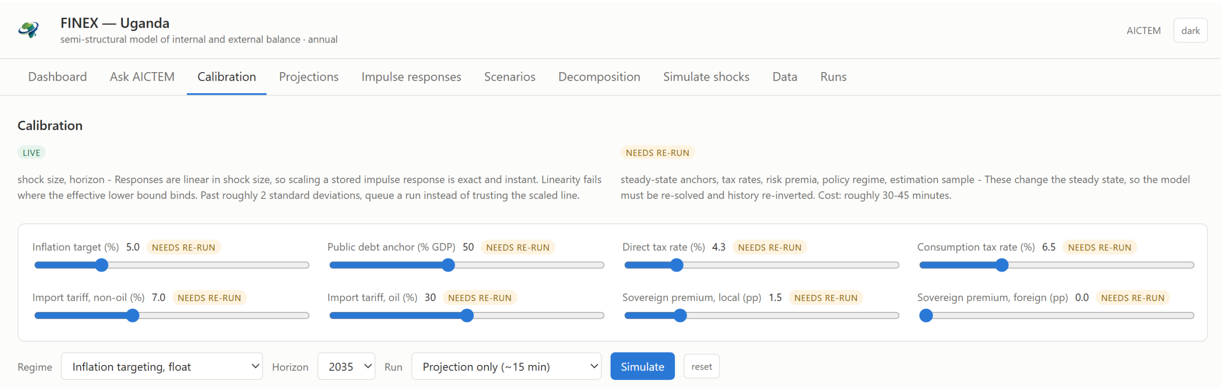

The left column of the legend is live: shock sizes and horizons rescale stored results and are exact while the model is linear. The right column is not: steady-state anchors, tax rates, risk premia, the policy regime and the estimation sample change the steady state, so the model must be re-solved and history re-inverted, which takes roughly 30 to 45 minutes. Each control carries its own badge.

G.12 Data and provenance

Provenance is part of the interface rather than an annex to it. A series built by splicing the International Monetary Fund's consumer price index onto the World Bank's, or by deflating a balance-of-payments aggregate, is a modelling choice, and a reader evaluating a result needs to see it without leaving the application. The same table underlies Table 31.

Figure 53: The observable panel and where every series came from

FINEX — Uganda
semi-structural model of internal and external balance · annual

AICTEM dark

Dashboard | Ask AICTEM | Calibration | Projections | Impulse responses | Scenarios | Decomposition | Simulate shocks | Data | Runs

84 observables built for Uganda.

observable	block	definition	source
DL_GDP	real	Real GDP growth	World Bank WDI national accounts (calendar year, LCU)
L_GDP_GAP	real	Output gap	World Bank WDI national accounts (calendar year, LCU)
DL_CONS	real	Private consumption growth	World Bank WDI national accounts (calendar year, LCU)
L_CONS_GAP	real	Consumption gap	World Bank WDI national accounts (calendar year, LCU)
CONS_RAT	real	Private consumption, % GDP	World Bank WDI national accounts (calendar year, LCU)
DL_INV	real	Investment growth	World Bank WDI national accounts (calendar year, LCU)
L_INV_GAP	real	Investment gap	World Bank WDI national accounts (calendar year, LCU)
INV_TOT_RAT	real	Total gross capital formation, % GDP	World Bank WDI national accounts (calendar year, LCU)
INV_RAT	real	Private investment (capital formation less the public development budget), % GDP	UBOS/World Bank WDI national accounts (calendar year, LCU)
KG_PUB_RAT	real	Public capital stock, % GDP	IMF Investment and Capital Stock Dataset (public/private split)
KG_PRIV_RAT	real	Private capital stock, % GDP	IMF Investment and Capital Stock Dataset (public/private split)
L_GOV_GAP	real	Government consumption gap	World Bank WDI national accounts (calendar year, LCU)
DL_EXP_NR	external	Natural-resource (gold + coffee) export volume growth	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual) + IMF metals/coffee prices
L_EXP_NR_GAP	external	NR export volume gap	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual) + IMF metals/coffee prices
EXP_NR_RAT	external	NR exports, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
DL_EXP_NNR	external	Non-NR export volume growth	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual) + NA export deflator
L_EXP_NNR_GAP	external	Non-NR export volume gap	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
EXP_NNR_RAT	external	Non-NR exports, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
DL_IMP_OIL	external	Oil import volume growth	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual) + Brent
L_IMP_OIL_GAP	external	Oil import volume gap	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
IMP_OIL_RAT	external	Oil imports, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
DL_IMP_NOIL	external	Non-oil import volume growth	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual) + NA import deflator
L_IMP_NOIL_GAP	external	Non-oil import volume gap	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
IMP_NOIL_RAT	external	Non-oil imports, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
DL_P_CONS	prices	Consumer price inflation	IMF CPI spliced onto World Bank CPI
DL_P_GDP	prices	GDP deflator inflation	World Bank WDI national accounts (calendar year, LCU)
DL_P_INV	prices	Investment deflator inflation	World Bank WDI national accounts (calendar year, LCU)
L_RP_INV_GAP	prices	Relative investment price gap	World Bank WDI national accounts (calendar year, LCU)
DL_P_GOV	prices	Government consumption deflator inflation	World Bank WDI national accounts (calendar year, LCU)
L_RP_GOV_GAP	prices	Relative government price gap	World Bank WDI national accounts (calendar year, LCU)
DL_P_EXP_NNR	prices	Other export price inflation (LCY)	World Bank WDI national accounts (calendar year, LCU)
L_RP_EXP_NNR_GAP	prices	Relative other export price gap	World Bank WDI national accounts (calendar year, LCU)
DL_P_EXP_NR	prices	NR export price inflation (shillings)	IMF metals/coffee prices x NER
L_RP_EXP_NR_GAP	prices	Relative NR export price gap	commodity prices x NER / CPI
DL_P_IMP_NOIL	prices	Non-oil import price inflation (LCY)	World Bank WDI national accounts (calendar year, LCU)
L_RP_IMP_NOIL_GAP	prices	Relative non-oil import price gap	World Bank WDI national accounts (calendar year, LCU)
DL_P_IMP_OIL	prices	Oil import price inflation (LCY)	Brent x NER
L_RP_IMP_OIL_GAP	prices	Relative oil import price gap	Brent x NER / CPI
RS	monetary	Short-term nominal interest rate	IMF monetary and financial statistics, interest rates
DL_NER	monetary	Nominal depreciation vs USD	World Bank GEM average exchange rate
L_REER_GAP	monetary	Real effective exchange-rate gap	World Bank REER
RR_GAP	monetary	Real interest rate gap	T-bill yield less CPI inflation
CA_RAT	external	Current account balance, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
FA_RAT	external	Net financial-account inflows, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
OFA_RAT	external	Net private financial inflows (FDI + portfolio + other investment), % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
ONFA_RAT	external	Banking-system net foreign assets, % GDP	IMF monetary statistics
OCA_RAT	external	Other current account (income + transfers ex-remittances), % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
REMIT_RAT	external	Remittance inflows, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
REMIT_RAT_GAP	external	Remittances gap, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
FX_RES_RAT	external	Gross FX reserves, % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
FXI_RAT	external	Reserve accumulation (FX intervention proxy), % GDP	Bank of Uganda, Balance of Payments BPM6 (quarterly, USD mn -> annual)
RR_PREM_GAP	external	Country risk-premium gap (UIP residual)	IMF monetary and financial statistics, interest rates, FRED, exchange rate

Eighty-four observables for Uganda, each with the block it belongs to, its definition and its source. Sources are named to the institution and the operation: a series spliced from two providers, or converted from quarterly balance-of-payments data to annual, says so.

G.13 The run ledger

Every result in the application is traceable to a row here, and failed runs are not deleted. **A ledger that records only successes is a worse record than no ledger**: run 1 failed on every attempt and says so, and runs 2 and 11 are marked as diverging even though they completed. The queue is bounded by the machine, with two workers able to hold a model in memory and one for lighter work.

Figure 54: The run ledger: every job, its verdict, and its failures

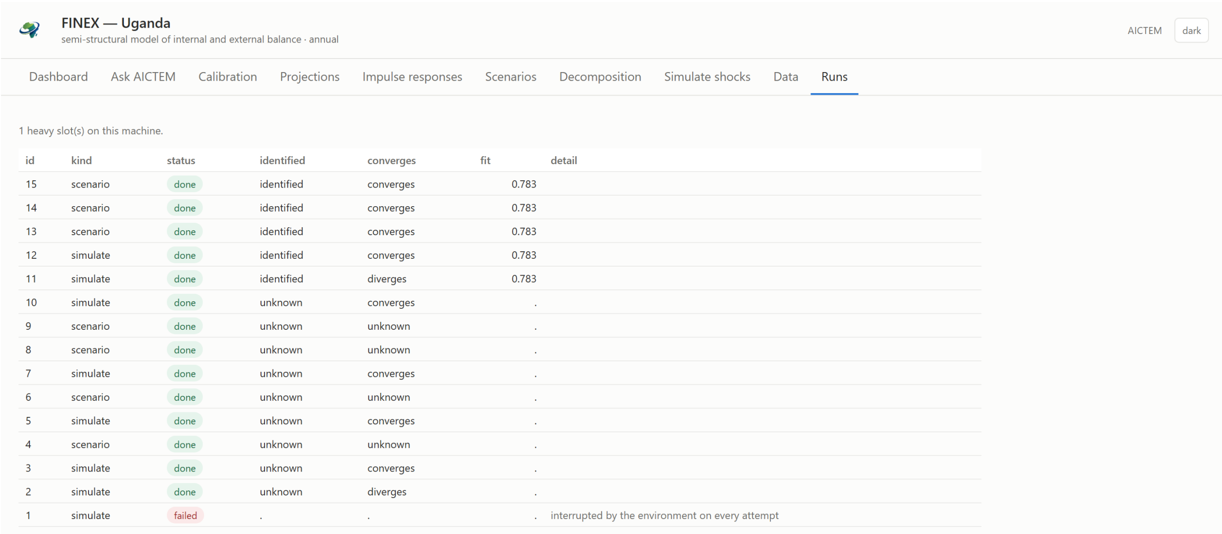

id	kind	status	identified	converges	fit	detail
15	scenario	done	identified	converges	0.783	
14	scenario	done	identified	converges	0.783	
13	scenario	done	identified	converges	0.783	
12	simulate	done	identified	converges	0.783	
11	simulate	done	identified	diverges	0.783	
10	simulate	done	unknown	converges	.	
9	scenario	done	unknown	unknown	.	
8	scenario	done	unknown	unknown	.	
7	simulate	done	unknown	converges	.	
6	scenario	done	unknown	unknown	.	
5	simulate	done	unknown	converges	.	
4	scenario	done	unknown	unknown	.	
3	simulate	done	unknown	converges	.	
2	simulate	done	unknown	diverges	.	
1	simulate	failed	.	.	.	interrupted by the environment on every attempt

One row per job, with kind, status, the two gate verdicts and the moment fit. Rows that predate the gate show *unknown* rather than a flattering default. Failures are kept and their reason is printed.

G.14 What the application is not

It is not a replacement for judgement, and it is not a forecasting service. The baseline it reports assumes no future shocks, which makes it a reference path rather than a central expectation, and the page says so wherever that path is drawn. It carries no off-model information: an announced policy that has not been entered as a shock is not in the projection. Its answers inherit every limitation set out in Section 11 — above all that 64 of 83 shock scales are assumptions rather than estimates, and that the monetary and exchange-rate blocks, which are the ones a central bank would most want to interrogate, are exactly the blocks the data could not pin down. **The application makes these limits visible at the point of use, which is the most an interface can honestly do about them.**

H. The observable panel

Every series built for each of the three countries, with coverage over the estimation window. Series whose name matches a FINEX model variable enter the estimation; the remainder are retained for diagnosis and cross-checking. Transforms follow the model's own definitions (Section 4). **The observation counts are the reason the three estimates are not equally anchored**: a moment can discipline a parameter only where the series behind it exists.

Table 31: The observable panel, all three countries

		Uganda		Kenya		South Africa	
Observable	Description	Obs	Mean	Obs	Mean	Obs	Mean
DL_GDP	Real GDP growth	24	5.8	24	4.4	24	2.1
L_GDP_GAP	Output gap	24	-0.2	24	-0.2	24	0.1
DL_CONS	Private consumption growth	24	5.2	24	4.2	24	2.6
L_CONS_GAP	Consumption gap	24	-0.4	24	-0.3	24	-0.0
CONS_RAT	Private consumption, % GDP	24	73.8	24	76.2	24	63.4
DL_INV	Investment growth	24	7.1	24	6.1	24	2.2
L_INV_GAP	Investment gap	24	-0.7	24	-0.8	24	-0.2
INV_RAT	Private investment (capital formation less the public development budget),	24	17.1	24	14.7	24	15.0
L_GOV_GAP	Government consumption gap	24	-0.5	24	-0.8	24	0.2
DL_EXP_NR	Natural-resource (gold + coffee) export volume growth	23	12.6	19	-0.4	24	2.8
L_EXP_NR_GAP	NR export volume gap	24	0.0	21	-0.3	24	0.9
EXP_NR_RAT	NR exports, % GDP	24	3.0	21	7.2	24	11.5
DL_EXP_NNR	Non-NR export volume growth	23	10.8	19	7.8	24	0.4
L_EXP_NNR_GAP	Non-NR export volume gap	24	0.0	21	0.1	24	-0.4
EXP_NNR_RAT	Non-NR exports, % GDP	24	12.9	21	11.4	24	16.5
DL_IMP_OIL	Oil import volume growth	23	8.9	19	3.0	24	3.0
L_IMP_OIL_GAP	Oil import volume gap	24	0.0	21	-1.4	24	-4.6
IMP_OIL_RAT	Oil imports, % GDP	24	2.8	21	4.9	24	4.9
DL_IMP_NOIL	Non-oil import volume growth	23	9.3	19	6.7	24	3.1
L_IMP_NOIL_GAP	Non-oil import volume gap	24	-0.0	21	-0.1	24	-0.1
IMP_NOIL_RAT	Non-oil imports, % GDP	24	23.4	21	22.1	24	22.3
DL_P_CONS	Consumer price inflation	24	5.8	24	8.0	24	5.1
DL_P_GDP	GDP deflator inflation	24	7.0	24	7.4	24	6.0
DL_P_INV	Investment deflator inflation	24	6.3	24	5.4	24	5.6
L_RP_INV_GAP	Relative investment price gap	24	0.1	24	0.4	24	0.2
DL_P_GOV	Government consumption deflator inflation	24	4.8	24	5.9	24	6.2
L_RP_GOV_GAP	Relative government price gap	24	0.8	24	1.0	24	0.0
DL_P_EXP_NNR	Other export price inflation (LCY)	24	5.4	24	5.6	24	7.6
L_RP_EXP_NNR_GAP	Relative other export price gap	24	-0.5	24	0.3	24	0.8
DL_P_EXP_NR	NR export price inflation (shillings)	24	8.2	24	7.0	24	8.6
L_RP_EXP_NR_GAP	Relative NR export price gap	24	-2.2	24	0.1	24	0.2
DL_P_IMP_NOIL	Non-oil import price inflation (LCY)	24	7.5	24	4.2	24	6.0
L_RP_IMP_NOIL_GAP	Relative non-oil import price gap	24	0.5	24	0.4	24	0.9
DL_P_IMP_OIL	Oil import price inflation (LCY)	24	7.7	24	6.6	24	8.3
L_RP_IMP_OIL_GAP	Relative oil import price gap	24	2.0	24	1.7	24	1.7
RS	Short-term nominal interest rate	19	10.7	24	8.3	24	7.2
DL_NER	Nominal depreciation vs USD	24	3.4	24	2.4	24	4.0
L_REER_GAP	Real effective exchange-rate gap	24	-0.6	24	-0.7	24	-0.2

Observable	Description	Obs	Mean	Obs	Mean	Obs	Mean
RR_GAP	Real interest rate gap	19	-0.1	24	-0.3	24	-0.2
CA_RAT	Current account balance, % GDP	24	-5.4	24	-4.3	24	-2.1
FA_RAT	Net financial-account inflows, % GDP	24	4.4	24	4.0	24	2.0
OFA_RAT	Net private financial inflows (FDI + portfolio + other investment), % GDP	24	3.7	24	4.9	24	2.6
ONFA_RAT	Banking-system net foreign assets, % GDP	24	12.7	24	7.9	24	14.5
OCA_RAT	Other current account (income + transfers ex-remittances), % GDP	24	1.1	24	2.5	24	-3.1
REMIT_RAT	Remittance inflows, % GDP	24	3.7	24	2.2	24	0.2
REMIT_RAT_GAP	Remittances gap, % GDP	24	-0.0	24	-0.1	24	-0.0
FX_RES_RAT	Gross FX reserves, % GDP	24	12.1	24	12.9	24	11.4
FXI_RAT	Reserve accumulation (FX intervention proxy), % GDP	23	0.8	24	1.3	24	0.7
RR_PREM_GAP	Country risk-premium gap (UIP residual)	19	0.3	24	-0.2	24	0.3
GREV_RAT	Revenue including grants, % GDP	24	15.0	24	16.8	24	24.9
GREV_OTH_RAT	Non-tax revenue and grants, % GDP	24	3.3	24	5.0	24	7.8
GEXP_CONS_RAT	Government consumption (national accounts), % GDP	24	10.7	24	13.7	24	18.2
GEXP_CONS_FIS_RAT	Government consumption, cash budget basis (wages + other current), % GDP	24	9.5	23	16.6	24	21.7
GEXP_INV_RAT	Public (development) investment, % GDP	24	6.9	23	4.6	24	1.8
GOV_RAT	Government purchases of goods and services, % GDP	24	17.7	24	17.8	24	20.0
GEXP_INT_RAT	Interest payments, % GDP	24	1.8	24	2.6	24	3.2
GEXP_OTH_RAT	Net lending and arrears clearance, % GDP	24	0.7	21	1.0	24	2.8
GEXP_RAT	Total expenditure and net lending, % GDP	24	19.0	24	21.0	24	28.3
DEF_O_RAT	Overall deficit (incl. grants), % GDP	24	4.0	24	4.2	24	3.3
DEF_S_RAT	Structural (cyclically adjusted) deficit, % GDP	24	4.0	24	4.2	24	3.3
DEF_PRIM_RAT	Primary deficit, % GDP	24	2.2	24	1.6	24	0.1
FIN_FCY_RAT	Net external financing, % GDP	24	2.3	24	1.8	24	1.1
DEB_RAT	General government gross debt, % GDP	24	35.2	24	47.8	24	44.1
DEB_FCY_RAT	External (foreign-currency) public debt, % GDP	19	18.9	24	26.4	24	15.1
DEB_LCY_RAT	Domestic-currency public debt, % GDP	19	12.3	24	21.4	24	29.0
DEB_FCY_SHARE	Foreign-currency share of public debt, %	19	61.9	24	56.6	24	32.7
US_L_GDP_GAP	US output gap	24	-0.8	24	-0.8	24	-0.8
US_DL_P_CONS	US CPI inflation	24	2.5	24	2.5	24	2.5
US_RS	US 3-month T-bill rate	24	1.6	24	1.6	24	1.6
EZ_L_GDP_GAP	Euro-area output gap	24	0.1	24	0.1	24	0.1
EZ_DL_P_CONS	Euro-area HICP inflation	24	2.1	24	2.1	24	2.1
EZ_L_RER_GAP	Euro/dollar real exchange-rate gap	24	-0.4	24	-0.4	24	-0.4
L_RP_OIL_GAP	Real world oil price gap	24	0.9	24	0.9	24	0.9
L_RP_NR_GAP	Real natural-resource (metals) price gap	24	-0.5	24	-0.7	24	-0.7
INV_TOT_RAT	Total gross capital formation, % GDP	24	24.0	24	19.5	24	16.8

Observable	Description	Obs	Mean	Obs	Mean	Obs	Mean
KG_PUB_RAT	Public capital stock, % GDP	19	21.3	19	36.4	19	46.6
KG_PRIV_RAT	Private capital stock, % GDP	19	76.7	19	89.9	19	141.0
GEXP_TRF_RAT	Subsidies, grants and social benefits, % GDP	24	2.5	23	1.8	24	6.3
DEB_PRIV_RAT	Private non-financial sector debt, % GDP	24	9.8	23	25.9	24	70.8
CREDIT_RAT	Domestic credit to the private sector, % GDP	14	13.6	23	28.6	24	116.5
NFA_RAT	Net international investment position, % GDP	24	-45.2	17	-34.7	24	-0.8
L_TOT_GAP	Commodity terms-of-trade gap	24	-2.2	24	-1.1	24	0.1
L_RP_CX_GAP	Real commodity export price gap	24	-3.1	24	-1.3	24	0.8
L_RP_CM_GAP	Real commodity import price gap	24	0.7	24	0.9	24	0.7

Notes: Observation counts and means over 2001–2024, the estimation window. A dash marks a series the country's panel does not contain; a count below the window length marks one that is present but short. **The counts are the reason the three estimates are not equally anchored**: a moment can only discipline a parameter where the series behind it exists.

I. The structural shocks

FINEX carries 83 structural shocks. 40 are estimable from a single country's second moments; the number actually identified is Uganda 19, South Africa 15, Kenya 16. The rest revert to their calibrated values, and the remaining 43 drive exogenous foreign and trend processes and are calibrated throughout. **The three countries do not identify the same 40**, which is why the table reports each separately rather than marking a shock estimated or calibrated once.

The "Src" columns are the ones to read before using any impulse response. A response to a shock whose scale was estimated is a statement about that country's data; a response to a shock whose scale was calibrated is a statement about the model's structure at an assumed scale. Both are useful; they are not the same claim, and the same shock can be one in Uganda and the other in Kenya.

Table 32: The 83 structural shocks, with estimated or calibrated scale, all three countries

Shock	Block	Uganda σ	Src	S. Afr. σ	Src	Kenya σ	Src
RES_L_CONS_GAP	Demand - consumption	0.8	cal.	1.081	est.	0.8	cal.
RES_L_INV_GAP	Demand - investment	1.67	est.	2.5	cal.	3.129	est.
RES_FIN_FCY_RAT	External - official financing	1.404	est.	0.6	cal.	0.5713	est.
RES_OCA_RAT	External - other current account	0.5	cal.	0.5	cal.	0.5	cal.
RES_OCA_RAT_BAR	External - other current account	0.1	cal.	0.1	cal.	0.1	cal.
RES_REMIT_RAT_BAR	External - remittances	0.2	cal.	0.2	cal.	0.2	cal.
RES_REMIT_RAT_GAP	External - remittances	0.5	cal.	0.5	cal.	0.5	cal.
RES_FXA_RAT	External - reserves	0.05	cal.	0.05	cal.	0.05	cal.
RES_FXI_RAT	External - reserves	0.1	cal.	0.1	cal.	0.1	cal.
RES_FX_RES_RAT	External - reserves	0.001	cal.	0.001	cal.	0.001	cal.
RES_FX_RES_RAT_TAR	External - reserves	0.03	cal.	0.03	cal.	0.03	cal.
RES_DEB_FCY_RAT	Fiscal - debt	0.1	cal.	0.1	cal.	0.1	cal.
RES_DEB_FCY_SHARE_TAR	Fiscal - debt	0.5	cal.	0.5	cal.	0.5	cal.
RES_DEB_LCY_RAT	Fiscal - debt	1	cal.	1	cal.	1	cal.
RES_DEB_RAT_TAR	Fiscal - debt	1	cal.	1	cal.	1	cal.
RES_DEF_S_RAT	Fiscal - deficit	0.805	est.	2	cal.	2	cal.
RES_GEXP_CONS_RAT	Fiscal - expenditure	0.2139	est.	0.0337	est.	0.1565	est.
RES_GEXP_OTH_RAT	Fiscal - expenditure	1	cal.	1	cal.	1	cal.
RES_GEXP_TRF_RAT	Fiscal - expenditure	1	cal.	1	cal.	1	cal.
RES_GEXP_TRF_RAT_BAR	Fiscal - expenditure	0.05	cal.	0.05	cal.	0.05	cal.
RES_GREV_OTH_RAT	Fiscal - revenue	0.5131	est.	1	cal.	1	cal.
RES_GREV_OTH_RAT_BAR	Fiscal - revenue	0.1	cal.	0.1	cal.	0.1	cal.
RES_ABS_US_DL_GDP_BAR	Foreign - United States	0.5	cal.	0.5	cal.	0.5	cal.
RES_ABS_US_DL_P_CONS	Foreign - United States	1.388	est.	1.355	est.	0.838	est.
RES_ABS_US_L_GDP_GAP	Foreign - United States	1.361	est.	1.438	est.	1.695	est.
RES_ABS_US_RR_BAR	Foreign - United States	0.2	cal.	0.2	cal.	0.2	cal.
RES_ABS_US_RS	Foreign - United States	1.405	est.	1.119	est.	1.258	est.
RES_ABS_EZ_DL_GDP_BAR	Foreign - euro area	0.2	cal.	0.2	cal.	0.2	cal.
RES_ABS_EZ_DL_P_CONS	Foreign - euro area	1.566	est.	1.58	est.	1.574	est.
RES_ABS_EZ_DL_RER_BAR	Foreign - euro area	1	cal.	1	cal.	1	cal.
RES_ABS_EZ_L_GDP_GAP	Foreign - euro area	1.424	est.	1.37	est.	1.442	est.
RES_ABS_EZ_L_RER_BAR	Foreign - euro area	0.5	cal.	0.5	cal.	0.5	cal.
RES_ABS_EZ_L_RER_GAP	Foreign - euro area	5.129	est.	4.93	est.	5.106	est.
RES_RS	Monetary policy	0.5	cal.	1.373	est.	0.5	cal.
RES_RS_GOV_FCY	Monetary policy	0.3	cal.	0.3	cal.	0.3	cal.
RES_RS_GOV_LCY	Monetary policy	0.4	cal.	0.4	cal.	0.4	cal.
RES_RS_ONFA	Monetary policy	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_CONS_BAR	Other	0.2	cal.	0.2	cal.	0.2	cal.
RES_DL_EXP_NNR_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.
RES_DL_EXP_NR_BAR	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_GTFP_BAR	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_IMP_NOIL_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.
RES_DL_IMP_OIL_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.
RES_DL_INV_BAR	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_RP_EXP_NNR_BAR	Other	0.2	cal.	0.2	cal.	0.2	cal.
RES_DL_RP_EXP_NR_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.
RES_DL_RP_GOV_BAR	Other	0.2	cal.	0.2	cal.	0.2	cal.
RES_DL_RP_IMP_NOIL_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.
RES_DL_RP_IMP_OIL_BAR	Other	0.4	cal.	0.4	cal.	0.4	cal.

Shock	Block	Uganda σ	Src	S. Afr. σ	Src	Kenya σ	Src
RES_DL_RP_INV_BAR	Other	0.2	cal.	0.2	cal.	0.2	cal.
RES_EXO_OFA_RAT_BAR	Other	0.1	cal.	0.1	cal.	0.1	cal.
RES_EXP_NNR_SUP_BAR	Other	0.1	cal.	0.1	cal.	0.1	cal.
RES_EXP_NR_SUP_BAR	Other	1	cal.	1	cal.	1	cal.
RES_IMP_NOIL_DEM_BAR	Other	0.1	cal.	0.1	cal.	0.1	cal.
RES_IMP_OIL_DEM_BAR	Other	0.1	cal.	0.1	cal.	0.1	cal.
RES_L_TFP_BAR	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_TAU_C	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_TAU_M_NOIL	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_TAU_M_OIL	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_TAU_NR	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_TAU_OFA	Other	0.1	cal.	0.1	cal.	0.1	cal.
RES_TAU_Y	Other	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_P_CONS	Prices	4.361	est.	0.9915	est.	1	cal.
RES_DL_P_CONS_TAR	Prices	0.5	cal.	0.5	cal.	0.5	cal.
RES_DL_P_EXP_NNR	Prices	4	cal.	4	cal.	4	cal.
RES_DL_P_EXP_NR	Prices	13.84	est.	3	cal.	8.79	est.
RES_DL_P_GOV	Prices	4.216	est.	1	cal.	2.969	est.
RES_DL_P_IMP_NOIL	Prices	2.945	est.	2	cal.	2.938	est.
RES_DL_P_IMP_OIL	Prices	2	cal.	4.535	est.	2	cal.
RES_DL_P_INV	Prices	2	cal.	2.3	est.	5.376	est.
RES_RR_PREM_BAR	Risk premium	0.5	cal.	0.5	cal.	0.5	cal.
RES_RR_PREM_GAP	Risk premium	0.7	cal.	0.7	cal.	0.7	cal.
RES_RR_PREM_GOV_FCY	Risk premium	0.1	cal.	0.9508	est.	2.134	est.
RES_RR_PREM_GOV_LCY	Risk premium	2.719	est.	0.25	cal.	0.25	cal.
RES_RR_PREM_ONFA	Risk premium	0.2	cal.	0.2	cal.	0.2	cal.
RES_L_EXP_NNR_GAP	Trade - exports	3	cal.	3	cal.	1.223	est.
RES_L_EXP_NR_GAP	Trade - exports	15.24	est.	3	cal.	8.831	est.
RES_L_IMP_NOIL_GAP	Trade - imports	3.56	est.	3	cal.	3	cal.
RES_L_IMP_OIL_GAP	Trade - imports	7.179	est.	8.678	est.	2	cal.
RES_ABS_L_RP_NR_GAP	World commodity price	17	cal.	17	cal.	17	cal.
RES_ABS_L_RP_OIL_GAP	World oil price	17	cal.	10.86	est.	17	cal.
RES_ABS_DL_RP_NR_BAR	World relative-price trend	6	cal.	6	cal.	6	cal.
RES_ABS_DL_RP_OIL_BAR	World relative-price trend	6	cal.	6	cal.	6	cal.

Notes: Columns run Uganda, South Africa, Kenya. “Src” records whether the standard deviation was identified by the estimation (*est.*) or reverted to its calibrated value after the estimate reached the zero bound (*cal.*). Of the 40 estimable variances the number identified is Uganda 19, South Africa 15, Kenya 16 — so the majority of every impulse response in this report is scaled by an assumption, and the three countries do not even assume in the same places.

J. Parameters in full

All 56 estimated quantities: 16 behavioural coefficients searched by Nelder–Mead and 40 shock variances solved by non-negative least squares at each outer evaluation (Section 5.3).

A variance at the zero bound is reported as such rather than as an estimate of zero. The distinction matters: the estimator is saying it cannot separate that shock's contribution from the others, not that the shock is absent from the economy.

Table 33: All estimated parameters: the shared calibration and each country's estimate

Parameter	Block	Cal.	Uganda	S. Africa	Kenya	Identification
C1_L_CONS_GAP	behavioural	0.25	**0.4455**	**0.4605**	**0.5121**	searched
C3_L_CONS_GAP	behavioural	0.1	**0.582**	**0.0213**	**0.5902**	searched
C1_L_INV_GAP	behavioural	0.15	0.1523	**0.164**	0.1562	searched
C1_L_EXP_NNR_GAP	behavioural	0.15	**0.7425**	**0.0897**	**0.0061**	searched
C1_L_IMP_NOIL_GAP	behavioural	0.15	**0.0037**	**0.0595**	**0.7668**	searched
C1_DL_P_CONS	behavioural	0.8	**0.634**	**0.9032**	**0.9092**	searched
C2_DL_P_CONS	behavioural	0.3	**0.3207**	**0.1414**	**0.1064**	searched
C4_DL_P_CONS	behavioural	0.3	**0.1857**	**0.462**	**0.1888**	searched
C1_RS	behavioural	0.15	**0.7327**	**0.4699**	**0.6802**	searched
C2_RS	behavioural	1	**0.5535**	**0.7857**	0.9511	searched
C3_RS	behavioural	0.2	**0.0443**	**0.4709**	**0.0185**	searched
C1_OFA_RAT	behavioural	1	1.016	1.022	1.005	searched
C1_OCA_RAT	behavioural	0.2	**0.107**	**0.1322**	**0.0457**	searched
C1_REMIT_RAT_GAP	behavioural	0.5	0.501	**0.4685**	0.511	searched
C1_DEF_S_RAT	behavioural	0.3	**0.0816**	**0.1344**	**0.2451**	searched
C1_GEXP_CONS_RAT	behavioural	0.7	**0.8635**	**0.618**	0.6969	searched
std_RES_L_CONS_GAP	shock	0.8	**0**	**1.081**	**0**	at zero bound
std_RES_L_INV_GAP	shock	2.5	**1.67**	**0**	**3.129**	identified
std_RES_L_EXP_NNR_GAP	shock	3	**0**	**0**	**1.223**	at zero bound
std_RES_L_EXP_NR_GAP	shock	3	**15.24**	**0**	**8.831**	identified
std_RES_L_IMP_NOIL_GAP	shock	3	**3.56**	**0**	**0**	identified
std_RES_L_IMP_OIL_GAP	shock	2	**7.179**	**8.678**	**0**	identified
std_RES_DL_P_CONS	shock	1	**4.361**	0.9915	**0**	identified
std_RES_DL_P_INV	shock	2	**0**	**2.3**	**5.376**	at zero bound
std_RES_DL_P_EXP_NNR	shock	4	**0**	**0**	**0**	at zero bound
std_RES_DL_P_EXP_NR	shock	3	**13.84**	**0**	**8.79**	identified
std_RES_DL_P_IMP_NOIL	shock	2	**2.945**	**0**	**2.938**	identified
std_RES_DL_P_IMP_OIL	shock	2	**0**	**4.535**	**0**	at zero bound
std_RES_DL_P_GOV	shock	1	**4.216**	**0**	**2.969**	identified
std_RES_REMIT_RAT_GAP	shock	0.5	**0**	**0**	**0**	at zero bound
std_RES_OCA_RAT	shock	0.5	**0**	**0**	**0**	at zero bound
std_RES_RR_PREM_GAP	shock	0.7	**0**	**0**	**0**	at zero bound
std_RES_TAU_OFA	shock	0.1	**0**	**0**	**0**	at zero bound
std_RES_RS	shock	0.5	**0**	**1.373**	**0**	at zero bound
std_RES_RS_GOV_LCY	shock	0.4	**0**	**0**	**0**	at zero bound
std_RES_RR_PREM_GOV_LCY	shock	0.25	**2.719**	**0**	**0**	weak (collinear)
std_RES_RS_GOV_FCY	shock	0.3	**0**	**0**	**0**	at zero bound
std_RES_RR_PREM_GOV_FCY	shock	0.1	**0**	**0.9508**	**2.134**	at zero bound
std_RES_RS_ONFA	shock	0.5	**0**	**0**	**0**	at zero bound
std_RES_RR_PREM_ONFA	shock	0.2	**0**	**0**	**0**	at zero bound
std_RES_DEF_S_RAT	shock	2	**0.805**	**0**	**0**	identified
std_RES_FIN_FCY_RAT	shock	0.6	**1.404**	**0**	0.5713	identified
std_RES_DEB_FCY_RAT	shock	0.1	**0**	**0**	**0**	at zero bound
std_RES_DEB_LCY_RAT	shock	1	**0**	**0**	**0**	at zero bound
std_RES_GREV_OTH_RAT	shock	1	**0.5131**	**0**	**0**	identified
std_RES_GEXP_CONS_RAT	shock	1	**0.2139**	**0.0337**	**0.1565**	identified
std_RES_GEXP_OTH_RAT	shock	1	**0**	**0**	**0**	at zero bound
std_RES_GEXP_TRF_RAT	shock	1	**0**	**0**	**0**	at zero bound
std_RES_ABS_US_RS	shock	0.5	**1.405**	**1.119**	**1.258**	identified
std_RES_ABS_US_L_GDP_GAP	shock	1	**1.361**	**1.438**	**1.695**	identified
std_RES_ABS_US_DL_P_CONS	shock	1	**1.388**	**1.355**	**0.838**	identified
std_RES_ABS_EZ_L_GDP_GAP	shock	0.2	**1.424**	**1.37**	**1.442**	identified
std_RES_ABS_EZ_DL_P_CONS	shock	0.3	**1.566**	**1.58**	**1.574**	identified
std_RES_ABS_EZ_L_RER_GAP	shock	1	**5.129**	**4.93**	**5.106**	identified
std_RES_ABS_L_RP_OIL_GAP	shock	17	**0**	**10.86**	**0**	at zero bound
std_RES_ABS_L_RP_NR_GAP	shock	17	**0**	**0**	**0**	at zero bound

Notes: Bold marks an estimate that moved more than 5% from the calibration the search started from, which is the same in all three countries. "Identification" records how the parameter was treated: `searched` for the behavioural coefficients in the Nelder–Mead outer loop, and the non-negative least-squares status for the shock variances; it is reported for Uganda and is the same in kind for the other two. A variance at the zero bound is reported as such rather than as an estimate of zero.

K. The steady states

The level each observable converges to in each country, and therefore the destination of every projection in this report. Calibrated from 2010–2024 sample means plus each country's own inflation target and public-debt anchor, with effective tax rates solved to reproduce its observed revenue ratio (Section 6.2).

Because the projection converges to these values, an error here is an error in every forecast. The two most consequential are the inflation anchor — which the projection converges to regardless of the data — and the debt anchor, which determines the direction and size of the fiscal adjustment the rule calls for.

Table 34: Steady state of the Uganda-calibrated model

Variable	Value	Variable	Value
CA_RAT	-0.462	GEXP_RAT	17.396
CONS_RAT	72.160	GEXP_TRF_RAT	1.870
DEB_FCY_RAT	30.600	GOV_RAT	12.113
DEB_LCY_RAT	19.400	GREV_OTH_RAT	1.790
DEB_RAT	50.000	GREV_RAT	13.320
DEF_O_RAT	4.076	IMP_NOIL_RAT	16.683
DEF_S_RAT	4.076	IMP_OIL_RAT	2.910
DL_CONS	4.915	INV_RAT	18.700
DL_EXP_NNR	4.515	L_CONS_GAP	0.000
DL_EXP_NR	4.915	L_EXP_NNR_GAP	0.000
DL_GDP	4.920	L_EXP_NR_GAP	0.000
DL_IMP_NOIL	3.415	L_GDP_GAP	0.000
DL_IMP_OIL	3.415	L_GOV_GAP	0.000
DL_INV	3.735	L_IMP_NOIL_GAP	0.000
DL_NER	1.379	L_IMP_OIL_GAP	0.000
DL_P_CONS	4.879	L_INV_GAP	0.000
DL_P_EXP_NNR	5.279	L_REER_GAP	0.000
DL_P_EXP_NR	4.879	L_RP_EXP_NNR_GAP	0.000
DL_P_GDP	4.874	L_RP_EXP_NR_GAP	0.000
DL_P_GOV	5.009	L_RP_GOV_GAP	0.000
DL_P_IMP_NOIL	6.379	L_RP_IMP_NOIL_GAP	0.000
DL_P_IMP_OIL	6.379	L_RP_IMP_OIL_GAP	0.000
DL_P_INV	6.059	L_RP_INV_GAP	0.000
EXP_NNR_RAT	13.220	L_RP_NR_GAP	0.000
EXP_NR_RAT	3.400	L_RP_OIL_GAP	0.000
EZ_DL_P_CONS	1.900	OCA_RAT	-0.570
EZ_L_GDP_GAP	0.000	OFA_RAT	-1.456
EZ_L_RER_GAP	0.000	ONFA_RAT	19.003
FA_RAT	0.462	REMIT_RAT	3.230
FIN_FCY_RAT	2.345	REMIT_RAT_GAP	0.000
FXI_RAT	0.000	RR_GAP	0.000
FX_RES_RAT	9.620	RR_PREM_GAP	0.000
GEXP_CONS_RAT	9.090	RS	7.579
GEXP_INT_RAT	2.593	US_DL_P_CONS	2.500
GEXP_INV_RAT	3.023	US_L_GDP_GAP	0.000
GEXP_OTH_RAT	0.820	US_RS	3.500

Notes: The level each observable converges to. Growth rates and interest rates in percent per year; ratios in percent of nominal GDP; gaps are zero in steady state by construction and are omitted where they carry no information. Calibrated from 2010–2024 sample means plus Uganda's inflation target and public-debt anchor, with effective tax rates solved to reproduce the observed revenue ratio.

L. Peak impulse responses

The peak response of eight headline observables to every one of the 83 structural shocks. This is the quickest way to see which shocks matter for what, and it is the table to consult before designing a scenario: a shock with a negligible peak response will not produce an interesting scenario however it is dated or scaled.

Table 35: Peak response to every structural shock, all three countries

	Uganda		Kenya		South Africa	
Shock	GDP gr.	Infl.	GDP gr.	Infl.	GDP gr.	Infl.
RES_ABS_DL_RP_NR_BAR	+0.04†	-0.12†	+0.07†	-0.12†	+0.17†	-0.25†
RES_ABS_DL_RP_OIL_BAR	-0.14†	+0.25†	-0.08†	+0.15†	-0.08†	+0.23†
RES_ABS_EZ_DL_GDP_BAR	+0.00†	-0.00†	+0.00†	-0.00†	+0.00†	-0.00†
RES_ABS_EZ_DL_P_CONS	-0.00	-0.00	+0.00	+0.00	-0.00	-0.00
RES_ABS_EZ_DL_RER_BAR	-0.02†	+0.03†	-0.02†	+0.02†	-0.01†	+0.02†
RES_ABS_EZ_L_GDP_GAP	+0.23	+0.06	+0.19	+0.08	+0.13	-0.04
RES_ABS_EZ_L_RER_BAR	-0.01†	+0.02†	-0.01†	+0.01†	-0.01†	+0.01†
RES_ABS_EZ_L_RER_GAP	+0.18	-0.54	+0.07	-0.22	+0.01	-0.22
RES_ABS_L_RP_NR_GAP	+0.16†	-0.07†	+0.29†	-0.05†	+0.71†	-0.15†
RES_ABS_L_RP_OIL_GAP	-0.27†	+0.80†	-0.24†	+0.76†	-0.04	+0.45
RES_ABS_US_DL_GDP_BAR	+0.02†	-0.02†	+0.02†	-0.01†	+0.01†	-0.01†
RES_ABS_US_DL_P_CONS	+0.05	-0.11	+0.03	-0.03	+0.06	-0.06
RES_ABS_US_L_GDP_GAP	+0.56	+0.14	+0.56	+0.24	+0.34	-0.10
RES_ABS_US_RR_BAR	-0.02†	+0.01†	-0.03†	+0.02†	-0.01†	+0.01†
RES_ABS_US_RS	+0.26	+0.74	+0.13	+0.36	-0.07	+0.29
RES_DEB_FCY_RAT	-0.01†	-0.02†	-0.01†	-0.01†	-0.01†	+0.01†
RES_DEB_FCY_SHARE_TAR	-0.03†	-0.06†	+0.02†	-0.05†	-0.01†	-0.04†
RES_DEB_LCY_RAT	-0.15†	-0.08†	-0.13†	-0.13†	-0.11†	+0.02†
RES_DEB_RAT_TAR	+0.15†	-0.04†	+0.16†	+0.06†	+0.14†	-0.04†
RES_DEF_S_RAT	+0.50	-0.02	+1.32†	-0.13†	+1.14†	+0.03†
RES_DL_CONS_BAR	-0.20†	-0.11†	-0.24†	-0.10†	-0.06†	-0.02†
RES_DL_EXP_NNR_BAR	-0.06†	-0.03†	-0.06†	-0.02†	-0.03†	-0.01†
RES_DL_EXP_NR_BAR	-0.02†	-0.01†	-0.04†	-0.01†	-0.03†	-0.01†
RES_DL_GTFP_BAR	+0.65†	+0.13†	+0.65†	+0.10†	+0.64†	+0.01†
RES_DL_IMP_NOIL_BAR	+0.08†	+0.04†	+0.10†	+0.03†	+0.02†	+0.01†
RES_DL_IMP_OIL_BAR	+0.01†	+0.01†	+0.02†	+0.01†	+0.01†	+0.00†
RES_DL_INV_BAR	+0.09†	-0.06†	+0.11†	-0.04†	+0.02†	-0.01†
RES_DL_P_CONS	-1.03	+2.93	+0.24†	+0.75†	-0.04	+0.67
RES_DL_P_CONS_TAR	+0.60†	+1.02†	+0.39†	+0.86†	+0.16†	+0.55†
RES_DL_P_EXP_NNR	-0.34†	-0.29†	+0.09†	-0.19†	-0.07†	-0.14†
RES_DL_P_EXP_NR	+0.05	-0.16	+0.12	-0.10	+0.10†	-0.07†
RES_DL_P_GOV	-0.21	-0.05	-0.13	-0.07	-0.04†	+0.00†
RES_DL_P_IMP_NOIL	-0.32	+0.72	-0.08	+0.24	-0.06†	+0.21†
RES_DL_P_IMP_OIL	-0.04†	+0.09†	-0.03†	+0.07†	-0.05	+0.18
RES_DL_P_INV	-0.18†	+0.38†	-0.36	+0.89	-0.07	+0.35
RES_DL_RP_EXP_NNR_BAR	+0.01†	-0.01†	+0.01†	-0.01†	+0.01†	-0.01†
RES_DL_RP_EXP_NR_BAR	+0.00†	+0.00†	+0.01†	+0.00†	-0.01†	+0.00†
RES_DL_RP_GOV_BAR	+0.01†	+0.01†	+0.01†	+0.01†	+0.01†	-0.00†
RES_DL_RP_IMP_NOIL_BAR	+0.02†	-0.01†	+0.02†	-0.01†	+0.02†	-0.00†
RES_DL_RP_IMP_OIL_BAR	-0.00†	-0.00†	+0.00†	-0.00†	+0.01†	-0.00†
RES_DL_RP_INV_BAR	-0.02†	-0.03†	-0.03†	-0.03†	-0.02†	-0.02†
RES_EXO_OFA_RAT_BAR	+0.04†	-0.05†	+0.04†	-0.03†	+0.02†	-0.01†
RES_EXP_NNR_SUP_BAR	+0.00†	-0.01†	+0.00†	-0.00†	+0.00†	-0.01†
RES_EXP_NR_SUP_BAR	-0.00†	+0.00†	-0.00†	+0.00†	-0.01†	-0.00†
RES_FIN_FCY_RAT	+0.13	+0.20	+0.02	+0.06	+0.02†	+0.05†
RES_FXA_RAT	-0.00†	-0.00†	-0.00†	-0.00†	-0.00†	-0.00†
RES_FXI_RAT	-0.00†	-0.01†	-0.00†	-0.01†	-0.00†	-0.01†
RES_FX_RES_RAT	-0.00†	-0.00†	+0.00†	-0.00†	-0.00†	-0.00†
RES_FX_RES_RAT_TAR	+0.00†	+0.01†	-0.00†	+0.01†	-0.00†	+0.00†
RES_GEXP_CONS_RAT	-0.07	-0.01	-0.05	+0.01	-0.00	-0.00
RES_GEXP_OTH_RAT	-0.87†	-0.17†	-0.86†	-0.34†	-0.73†	+0.03†
RES_GEXP_TRF_RAT	-0.50†	+0.34†	-0.44†	+0.25†	+0.21†	+0.30†
RES_GEXP_TRF_RAT_BAR	-0.04†	-0.02†	-0.05†	-0.02†	-0.01†	-0.01†
RES_GREV_OTH_RAT	+0.15	+0.14	+0.19†	+0.32†	+0.17†	-0.05†
RES_GREV_OTH_RAT_BAR	-0.08†	-0.03†	-0.10†	-0.03†	-0.05†	-0.00†
RES_IMP_NOIL_DEM_BAR	-0.00†	+0.00†	-0.00†	+0.00†	-0.00†	+0.00†
RES_IMP_OIL_DEM_BAR	+0.00†	+0.00†	-0.00†	+0.00†	+0.00†	+0.00†
RES_L_CONS_GAP	+0.76†	+0.23†	+0.82†	+0.28†	+0.74	+0.35
RES_L_EXP_NNR_GAP	+0.59†	+0.08†	-0.09	+0.01	+0.27†	-0.08†

Shock	GDP gr.	Infl.	GDP gr.	Infl.	GDP gr.	Infl.
RES_L_EXP_NR_GAP	+0.43	-0.09	+0.39	-0.04	+0.31†	-0.05†
RES_L_IMP_NOIL_GAP	+0.47	+0.18	-0.43†	+0.31†	-0.35†	+0.08†
RES_L_IMP_OIL_GAP	-0.16	+0.06	-0.07†	+0.01†	-0.37	+0.09
RES_L_INV_GAP	+0.45	+0.15	+0.75	+0.31	+0.45†	+0.16†
RES_L_TFP_BAR	+0.65†	+0.10†	+0.65†	+0.07†	+0.63†	-0.01†
RES_OCA_RAT	+0.06†	-0.18†	+0.04†	-0.09†	+0.02†	-0.09†
RES_OCA_RAT_BAR	+0.02†	-0.01†	+0.02†	-0.01†	+0.01†	-0.01†
RES_REMIT_RAT_BAR	-0.07†	-0.04†	-0.08†	-0.05†	-0.02†	-0.01†
RES_REMIT_RAT_GAP	+0.40†	-0.11†	+0.50†	+0.17†	+0.30†	+0.05†
RES_RR_PREM_BAR	-0.05†	-0.03†	-0.03†	+0.02†	-0.02†	+0.02†
RES_RR_PREM_GAP	-0.06†	+0.22†	-0.04†	+0.11†	-0.02†	+0.11†
RES_RR_PREM_GOV_FCY	-0.02†	+0.02†	-0.53	+0.42	+0.25	+0.08
RES_RR_PREM_GOV_LCY	-0.19	+0.20	+0.04†	+0.02†	+0.04†	+0.01†
RES_RR_PREM_ONFA	+0.01†	+0.01†	-0.01†	-0.01†	+0.00†	+0.00†
RES_RS	-0.83†	-0.73†	-0.59†	-0.54†	-0.20	-0.42
RES_RS_GOV_FCY	-0.06†	+0.02†	-0.06†	-0.01†	-0.05†	+0.01†
RES_RS_GOV_LCY	-0.05†	-0.01†	-0.09†	-0.02†	-0.12†	+0.00†
RES_RS_ONFA	+0.01†	+0.01†	-0.02†	-0.01†	+0.00†	+0.00†
RES_TAU_C	+0.29†	+0.11†	+0.33†	+0.10†	+0.13†	-0.02†
RES_TAU_M_NOIL	+0.04†	-0.01†	+0.05†	+0.01†	+0.02†	-0.01†
RES_TAU_M_OIL	+0.01†	-0.00†	+0.01†	-0.00†	+0.01†	-0.00†
RES_TAU_NR	+0.01†	-0.00†	+0.01†	+0.00†	+0.01†	-0.00†
RES_TAU_OFA	-0.04†	+0.05†	-0.04†	+0.03†	-0.02†	+0.01†
RES_TAU_Y	+0.55†	+0.25†	+0.53†	+0.20†	+0.20†	+0.02†

Notes: Peak absolute response over 15 years to a one-standard-deviation shock, in percentage points. A dagger marks a shock whose standard deviation reverted to its calibrated value: the shape of that response is estimated, its size assumed. Because the standard deviation differs by country, entries are comparable within a country but not across one.

M. Shock decomposition by year

The numbers behind Section 13 and the narrative of Section 15: each group's contribution to four key observables, year by year, over 2006–2024.

Read rows, not cells. Contributions are large relative to the total and heavily offsetting, so a single entry taken alone will mislead; the informative object is the balance of contributions within a year.

Table 36: Historical shock decomposition by year and channel

Variable	Year	DD	DP	MP	TV	EIF	FIS	WCP	FAR	NL	Total
L_GDP_GAP	2006	-0.1	1.3	-1.0	-1.2	-2.4	-1.4	0.5	0.9	-0.1	-3.5
	2007	3.8	0.9	-3.5	-1.0	-1.3	0.0	0.5	1.6	-0.0	1.0
	2008	-5.4	3.2	-2.3	1.8	1.8	-1.6	0.3	0.8	-0.6	-1.9
	2009	-1.1	3.7	7.0	1.0	0.7	-1.1	-0.7	-1.9	-2.1	5.5
	2010	-2.0	-4.1	13.9	-0.4	0.2	0.7	-0.3	-3.0	-1.6	3.4
	2011	1.8	-1.8	3.1	3.4	1.0	0.0	0.4	-3.2	-1.8	3.0
	2012	8.6	-4.8	-4.7	5.1	3.1	1.8	0.5	-3.0	-1.1	5.4
	2013	7.3	-9.9	0.3	4.9	2.8	1.5	0.2	-2.7	-0.6	3.8
	2014	6.7	-10.5	1.4	0.7	0.8	0.5	-0.4	-2.2	-0.4	-3.3
	2015	13.5	-3.6	-8.7	4.0	-0.8	0.4	-1.6	-1.5	-0.4	1.4
	2016	14.8	-2.0	-14.5	4.5	0.4	1.2	-1.6	-1.0	-0.1	1.8
	2017	6.0	-6.0	-6.9	3.2	1.1	0.8	-0.3	-0.2	-0.2	-2.4
	2018	3.0	-10.0	-1.9	7.7	1.8	1.7	1.0	0.7	-0.5	3.6
	2019	6.6	-9.9	-1.7	2.6	1.3	2.4	1.1	1.3	-0.4	3.4
	2020	8.5	-6.7	-1.7	-3.8	0.5	4.1	0.1	-1.3	-0.6	-0.8
	2021	7.0	-3.9	-1.6	-1.7	0.1	4.2	0.5	-1.1	-0.4	3.1
	2022	2.1	-0.8	-1.3	-2.0	-0.0	2.0	0.8	0.1	0.0	0.9
	2023	-1.5	-0.5	-0.8	0.9	-0.3	1.8	0.4	1.7	0.0	1.6
	2024	-2.9	-1.1	-0.5	0.5	-0.8	1.3	-0.3	2.2	-0.0	-1.6
DL_P_CONS	2006	0.2	4.3	-1.0	0.3	-7.1	4.5	0.7	0.4	-0.2	2.2
	2007	1.7	1.1	-3.2	0.3	-1.1	1.2	0.6	0.4	0.0	1.1
	2008	-1.0	7.5	-1.8	0.1	3.0	-1.5	1.2	-1.1	0.0	6.5
	2009	-0.4	8.1	7.2	-0.0	-2.2	-1.9	-1.1	-1.3	-1.0	7.4
	2010	0.1	-6.2	13.2	-0.2	-3.3	-1.0	-0.8	-1.6	-1.1	-1.0
	2011	1.7	10.1	1.5	-0.4	-1.1	-0.5	1.1	-1.7	-0.2	10.4
	2012	4.4	4.0	-6.8	0.1	4.5	0.6	1.6	-1.3	0.1	7.1
	2013	4.2	-7.4	-1.0	0.2	3.1	0.4	1.2	-1.0	0.1	-0.1
	2014	3.3	-7.2	1.3	0.7	0.3	0.3	0.5	-0.9	-0.1	-1.8
	2015	4.5	11.3	-8.1	-0.2	-2.7	-1.8	-2.4	-0.9	0.8	0.6
	2016	4.9	12.9	-13.3	-0.4	-0.1	-0.4	-3.0	-0.8	0.8	0.7
	2017	1.9	3.4	-4.7	-0.1	0.5	0.5	-1.0	-0.5	0.3	0.2
	2018	-0.6	-5.5	1.1	-0.7	2.1	0.1	1.4	-0.0	-0.2	-2.3
	2019	0.5	-6.8	0.7	0.9	1.3	0.4	1.2	0.1	-0.3	-2.1
	2020	2.8	-3.1	-0.3	1.8	-0.6	0.0	-1.4	-0.3	-0.5	-1.6
	2021	3.0	-1.5	-0.7	-0.4	-0.2	-0.7	-0.2	-1.8	-0.2	-2.7
	2022	0.5	2.8	-0.5	-1.2	0.6	-0.6	1.8	-1.4	-0.0	2.1
	2023	-1.9	0.1	-0.0	-0.6	0.6	-0.2	1.0	1.3	0.1	0.3
	2024	-2.5	-3.3	0.2	0.5	0.4	0.7	0.1	2.1	0.1	-1.6
DEB_RAT	2006	-0.1	-0.7	0.7	1.4	-2.3	-19.3	-0.9	-1.1	-0.0	-22.2
	2007	-3.6	-0.5	2.9	1.1	4.4	-33.9	-1.3	-2.5	0.2	-33.0
	2008	3.1	-4.8	3.9	-2.7	1.1	-33.2	-1.0	-3.0	2.4	-34.3
	2009	1.6	-14.1	-2.1	-1.2	-3.4	-29.0	0.2	1.9	10.9	-35.2
	2010	2.3	-4.2	-12.2	-0.3	-0.9	-27.6	0.2	4.0	7.1	-31.6
	2011	-0.8	-6.3	-11.5	-5.3	0.2	-25.1	-0.6	4.8	12.5	-32.0
	2012	-8.1	-2.0	-2.4	-7.9	0.6	-23.8	-1.0	5.5	8.7	-30.5
	2013	-10.2	6.9	-0.3	-7.9	-3.2	-23.1	-1.2	5.8	5.2	-27.9
	2014	-11.4	13.5	-2.6	-3.9	-2.7	-24.5	-0.5	5.4	1.4	-25.2
	2015	-18.0	5.9	3.2	-7.4	-1.9	-17.4	1.5	4.9	7.3	-22.0
	2016	-22.2	0.9	13.3	-8.0	0.1	-14.3	3.1	4.2	4.2	-18.7
	2017	-17.1	4.7	14.6	-6.4	-1.7	-17.0	2.3	3.0	1.2	-16.4
	2018	-12.8	12.5	9.4	-12.3	-1.6	-14.1	0.4	1.5	2.0	-15.0
	2019	-14.0	16.4	6.5	-6.3	-2.5	-12.5	-0.9	0.1	1.2	-12.0
	2020	-16.2	15.5	5.9	2.4	-2.7	-10.4	-0.2	2.8	-0.3	-3.2
	2021	-15.8	13.2	5.6	1.2	-1.0	-4.1	-0.4	1.4	0.3	0.3
	2022	-11.3	8.8	5.0	1.7	-0.1	-2.8	-1.1	-0.3	-0.1	-0.1
	2023	-6.0	7.2	4.0	-0.6	-0.1	-2.0	-1.3	-1.1	0.2	0.4
	2024	-2.1	7.5	3.1	0.0	0.7	-4.4	-0.6	-2.7	0.2	1.8
CA_RAT	2006	-0.1	-0.0	0.0	-1.2	5.7	0.4	-0.0	0.4	0.1	5.4
	2007	-0.3	-0.0	0.2	-1.0	3.7	0.0	0.2	0.8	0.1	3.6

Variable	Year	DD	DP	MP	TV	EIF	FIS	WCP	FAR	NL	Total
	2008	-0.0	-0.1	0.1	-0.2	3.7	0.1	-0.3	0.5	0.3	4.0
	2009	-0.2	-0.3	-0.2	1.6	1.3	-0.1	0.3	-0.9	0.3	1.7
	2010	-0.2	-0.2	-0.5	-0.0	2.5	-0.4	0.2	-1.3	0.8	0.8
	2011	-0.4	-0.3	-0.2	1.6	2.3	-0.2	-0.1	-1.4	0.5	1.8
	2012	-0.7	-0.4	-0.1	3.2	2.2	-0.5	-0.3	-1.4	0.3	2.4
	2013	-1.0	-0.3	-0.2	3.8	0.7	-0.5	-0.4	-1.4	0.2	0.9
	2014	-1.1	-0.1	-0.2	0.7	0.2	-0.4	-0.4	-1.2	0.4	-2.1
	2015	-1.2	-0.2	0.3	3.2	1.0	-0.5	0.3	-0.8	-0.0	2.1
	2016	-1.6	-0.4	0.5	3.8	0.5	-0.7	0.3	-0.6	-0.1	1.8
	2017	-1.4	-0.3	0.4	3.0	-0.6	-0.4	0.1	-0.3	0.1	0.5
	2018	-1.0	-0.2	0.4	5.8	-0.9	-0.6	-0.2	0.1	0.4	3.8
	2019	-1.2	0.0	0.4	1.8	-0.6	-0.8	-0.0	0.5	0.3	0.4
	2020	-1.5	0.1	0.4	-2.3	0.5	-1.1	0.6	-0.8	0.8	-3.4
	2021	-1.2	0.2	0.3	1.0	-0.6	-1.2	0.1	-0.3	0.2	-1.4
	2022	-1.0	0.2	0.3	0.9	-0.7	-0.8	-0.3	0.2	0.5	-0.8
	2023	-0.6	0.2	0.3	1.3	-1.5	-0.7	-0.1	0.4	0.1	-0.6
	2024	-0.5	0.2	0.3	0.3	-1.6	-0.5	-0.1	0.5	0.1	-1.2

Notes: Column codes: DD = Domestic demand; DP = Domestic prices; MP = Monetary policy; TV = Trade volumes; EIF = External income & finance; FIS = Fiscal; WCP = World commodity prices; FAR = Foreign activity & rates; NL = Non-linearity. Contributions of each shock group to the observable's deviation from its steady state, computed by cumulative group-wise re-simulation so that contributions sum exactly to the total. Contributions are large relative to the total and heavily offsetting; that is a property of the exact method, and means individual entries are more uncertain than the total. The ordering is fixed and order-dependent. "Nonlinear" is a residual, not a channel.

N. Data revisions and vintages

Every result in this report is tied to a specific data vintage. The panel is rebuilt from source rather than edited, so a revision to any underlying series propagates consistently through every transform, every moment, and every projection that depends on it.

Three consequences are worth recording for anyone reproducing or updating this work.

- **National accounts revisions move the steady state.** The great ratios are calibrated from 2010–2024 sample means, so a rebasing of Uganda's national accounts changes the destination of every projection, not merely its starting point.

- **Balance-of-payments revisions move the external block only.** Because the balance of payments enters through its own observables and the identity binds each period, a BoP revision does not propagate into the fiscal or monetary blocks except through the exchange rate.

- **Fiscal revisions propagate furthest.** Revenue and expenditure feed the tax calibration, which feeds government purchases, which feeds $C + I + G$, which feeds the steady-state trade deficit. A revision here is the one most likely to require re-checking results that appear unrelated.

The estimation sample is 2001–2024 and the inversion window 2006–2024. The earlier start for estimation exploits the additional observations available for second moments; the later start for inversion reflects the debt currency split, which is only reported from 2006 and before which the HIPC and MDRI relief make the gross-debt and external-debt measures non-comparable.

References

Bank of Canada (2024). StateSpaceEcon.jl, ModelBaseEcon.jl and TimeSeriesEcon.jl. Open-source Julia packages. https://github.com/bankofcanada/StateSpaceEcon.jl.

Berg, A., Hul, Y., Karam, P., Remo, A., and Rodriguez Guzman, D. (2023). FINEX — a new workhorse model for macroeconomic forecasting and policy analysis. IMF Working Paper WP/23/235, International Monetary Fund, Washington, DC.

Bezanson, J., Edelman, A., Karpinski, S., and Shah, V. B. (2017). Julia: A fresh approach to numerical computing.

Coleman, C., Lyon, S., Maliar, L., and Maliar, S. (2021). Matlab, python, julia: What to choose in economics? *Computational Economics*, 58:1263–1288.

Hodrick, R. J. and Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, 29(1):1–16.

Independent Evaluation Office of the IMF (2021). Fiscal adjustment, fiscal reforms and growth. Background paper, growth and adjustment in IMF-supported programs, International Monetary Fund, Washington, DC. https://ieo.imf.org/en/-/media/IEO/Files/evaluations/completed/09-09-2021-growth-and-adjustment-in-imf-supported-programs/adg-6-fiscal-adjustment-fiscal-reforms-and-growth.pdf.

International Monetary Fund (2015). Making public investment more efficient. Imf policy paper, International Monetary Fund, Washington, DC. https://www.imf.org/external/np/pp/eng/2015/061115.pdf.

International Monetary Fund (2016). Remittances and macroeconomic volatility in african countries. Imf working paper, International Monetary Fund, Washington, DC. https://www.imf.org/en/Publications/WP/Issues/2016/12/31/Remittances-and-Macroeconomic-Volatility-in-African-Countries-42752.

International Monetary Fund (2023). Integrated policy framework — principles for the use of foreign exchange intervention. Imf policy paper, International Monetary Fund, Washington, DC. https://www.imf.org/-/media/Files/Publications/PP/2023/English/PPEA2023061.pdf.

International Monetary Fund (2026). The fiscal multipliers narrative of sub-saharan africa. IMF Working Paper WP/26/012, International Monetary Fund, Washington, DC. https://www.imf.org/en/publications/wp/issues/2026/01/23/the-fiscal-multipliers-narrative-of-sub-saharan-africa-573440.

Kabundi, A. and Mlachila, M. (2025). Oil price passthrough to consumer price inflation in south africa: The role of the inflation environment. *Journal of Applied Economics*.

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*, 17(3):1217–1241.

Lawson, C. L. and Hanson, R. J. (1974). Solving least squares problems. *Prentice-Hall Series in Automatic Computation*.

Mishra, P. and Montiel, P. J. (2012). How effective is monetary transmission in low-income countries? a survey of the empirical evidence. Imf working paper, International Monetary Fund, Washington, DC.

Mishra, P., Montiel, P. J., and Spilimbergo, A. (2010). Monetary transmission in low-income countries. IMF Working Paper WP/10/223, International Monetary Fund, Washington, DC. https://www.imf.org/external/pubs/ft/wp/2010/wp10223.pdf.

Nelder, J. A. and Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, 7(4):308–313.

Nwosa, P. I. (2019). Implications of oil price shocks on net oil-importing african countries. *Heliyon*. https://pmc.ncbi.nlm.nih.gov/articles/PMC6728737/.

Ruge-Murcia, F. (2012). Estimating nonlinear DSGE models by the simulated method of moments: With an application to business cycles. *Journal of Economic Dynamics and Control*.

Woldu, G. T. (2023). Fiscal multipliers and structural economic characteristics: Evidence from countries in sub-saharan africa. *The World Economy*.