

SCOPE

What v1 is, what it isn't, and what's next.

For this audience, an honest scope statement is worth more than a feature list. Here is exactly what the first generation ships, and the frontier it is built to grow into.

What v1 ships

- **Modular composition** — four bases (Closed, Small Open Economy, Currency-Union Member, plus a declared two-country base), nine drop-in modules, twelve archetypes.
 - **82 pre-calibrated economies** across every IMF region.
 - **A validated solver core** — Klein (2000) generalized-Schur and Sims (2002) gensys, numerically validated to 1e-8 against closed-form solutions, with explicit Blanchard–Kahn determinacy handling.
 - **Forecasting & analysis** — impulse responses, central-bank-style fan charts, theoretical moments, variance decomposition, and a diagnostics panel.
 - **Reproducibility** — SHA-256 provenance hashing, immutable vintages, deterministic vintage diff, bit-exact across machines.
 - **Six export formats** — HTML, Quarto, LaTeX, Excel, PDF, CSV — all from one canonical run-state.
 - **CSV data ingestion** with column-to-variable mapping and history overlay.
 - **Research-grade documentation** — a 14-chapter manual and a full Engine SOP.
-

What v1 does not yet do

Stated plainly, so you can decide whether v1 fits your work:

- The shipping dashboard model runs a reduced-form simulator with calibrated proxies for moments, variance decomposition, and fan-band widths (see Engine SOP). The full QZ/gensys engine is validated and being wired into the live solve path.
- No heterogeneous-agent (HANK) models — representative agent only.
- No Bayesian estimation or Kalman filtering — calibrate-only; parameters are not fit to data.
- No climate / NGFS scenarios, regime-switching, or behavioural-expectations variants.
- No multi-country general equilibrium (the two-country base is declared, not implemented).
- No optimal-policy / welfare solver, and no live data-API integration (CSV upload only).

Roadmap – the frontier the architecture is built for

The modular design means each item below is *configuration on the existing platform*, not a re-build:

DIRECTION	WHAT IT UNLOCKS
Full state-space solve on the live path	Analytical moments (Lyapunov), exact FEVD, model-implied fan bands.
Bayesian estimation (SMC + Metropolis–Hastings)	Fit parameters to data, posterior bands.
Kalman filter	Latent-variable extraction, e.g. output-gap nowcasts.
HANK via Sequence-Space Jacobians	Distributional and heterogeneous-agent analysis.
NGFS climate scenarios	Carbon-pricing, transition- and physical-risk blocks.
Identification diagnostics	Iskrev, Komunjer–Ng, Qu–Tkachenko checks.

The durable bet. The architecture — declarative composition, a canonical run-state, and provenance hashing — is designed so the solver and the science can deepen without re-platforming. v1 is the foundation, deliberately shipped honest.