An open replication of the Bank of England structural VAR*

Vahid Ahmadi†

November 2026

Abstract

I present an open-source Python replication of the eight-variable Bayesian structural VAR that the Bank of England uses to disentangle domestic and global drivers of the UK business cycle (Brignone and Piffer, 2025). The model combines a Minnesota normal-inverse-Wishart prior with zero and sign restrictions imposed via the algorithm of Arias et al. (2018), identifying world demand, world energy, world supply, UK demand, UK supply and UK monetary policy shocks, with two unidentified residual shocks. The replication serves as the baseline and conditioning member of PolicyEngine Macro: it produces forecasts with credible bands, structural readings of the latest shocks, and forecast-revision narratives, but does not score reforms. Validation against the published paper is reported in two groups, because the checks are not of one kind. Checks that *cannot fail* — zero restrictions holding on every accepted draw, sign restrictions holding on every accepted draw, $BB' = \Sigma$, variance decompositions summing to one within 10^{-10}, the historical decomposition reconstructing the data within 10^{-8} — are properties of the Arias et al. (2018) construction, of the accept/reject rule, or algebraic identities; they check the implementation, not the replication, and their residuals are solver tolerances rather than accuracy. Checks that can fail are the substantive ones: the unrestricted oil-price response to a world supply shock is positive (+5.83 on impact) as in the paper's Figure 2, and, computed on the paper's own definition (posterior mean of the per-draw group share of total four-quarter-ahead variance), the identified global shocks explain 38.7 per cent of UK GDP forecast-error variance against roughly 40 per cent in the paper — a match — and 41.5 per cent of UK CPI against roughly 50, about eight and a half points short, though the 68 per cent posterior band of some ±15 points is wider than that gap. An earlier version of this paper reported 42.1 and 49.5 per cent and read both as matches; the CPI agreement came from renormalising a sum of posterior medians. An out-of-sample look from the 2024Q2 data edge shows the model's GDP forecast (year-on-year medians of 1.3–2.0 per cent) bracketing the ONS outturns, while CPI inflation (medians of 2.3–2.8 per cent) undershoots the 2025 re-acceleration to 3.8 per cent — an error shared by the Bank's own August 2024 projections — before the outturn converges back onto the model's path by 2026Q1; over seven evaluated quarters the root-mean-squared forecast errors are 0.32 percentage points for both GDP growth and CPI inflation. That is one forecast origin: across 50 rolling origins the model shows no accuracy gain over a random walk *with*

*Acknowledgements: I thank Max Ghenis for his great suggestions and helpful discussions. This paper and the accompanying codebase were prepared with the assistance of Claude (Anthropic); all remaining errors are my own. The replication is an independent open-source project and is not an official product of, or endorsed by, the Bank of England.

†Research Associate at PolicyEngine. Email: vahid@policyengine.org.

drift on any variable at any horizon that survives a false-discovery-rate adjustment of the 64 tests involved (minimum $q = 0.39$). The model is served through hosted MCP tools, making a central-bank-grade structural VAR callable from any AI workflow.

Current-vintage note (29 July 2026). The live forecast (5,600 draws, 385 accepted, effective sample size 176.3) is estimated on 1992Q1–2025Q1 with observations through 2026Q1 and covers 2026Q2–2029Q2. Median year-on-year GDP growth ranges from 1.14 to 1.76 per cent and CPI inflation from 2.54 to 3.10 per cent. The frozen 2024Q2 experiment below remains the out-of-sample evaluation.

Keywords: structural VAR, sign restrictions, Bayesian estimation, UK economy, replication, open source

1 Introduction

In July 2025 the Bank of England published, for the first time, a full technical description of the Bayesian structural vector autoregression (SVAR) it uses in the monetary policy round to separate domestic from global drivers of the UK business cycle (Brignone and Piffer, 2025). The model is small — eight quarterly variables, six identified structural shocks — but it carries substantial weight in the Bank's process: it supplies impulse responses, forecast-error variance decompositions, real-time structural shock readings, and a round-by-round structural narrative for revisions to the Monetary Policy Report forecast. No official replication package accompanies the paper.

This paper documents an independent, open-source Python replication of that model, developed as part of *PolicyEngine Macro*, PolicyEngine's suite of open macroeconomic models for the UK and US (Ghenis and PolicyEngine contributors, 2025).[1] Within the suite the SVAR plays a specific role: it is the *baseline and conditioning member*. It reads the current state of the economy in structural-shock terms, produces unconditional forecasts with credible bands, and generates revision narratives that other models in the suite condition on. It deliberately does *not* score policy reforms — reform scoring is the province of the suite's overlapping-generations, macroeconometric and microsimulation members — because a sign-restricted SVAR identifies shocks, not policy rules.

The contribution is threefold. First, *replication as a public good*: the Bank released no code or data, and its two world aggregates are internal, unpublished series. Reconstructing the pipeline — data assembly, Minnesota-prior Bayesian estimation, zero-and-sign-restriction identification following Arias et al. (2018), and the standard outputs — makes the Bank's analytical machinery inspectable and extensible by anyone. Second, *validation with numbers rather than adjectives*: Section 6 compares the replication against the paper's published figures quantitatively, reporting exact satisfaction of every identifying restriction, the sign of key unrestricted responses, and variance-decomposition shares side by side with the paper's, together with machine-precision internal-consistency checks. Third, *operationalisation*: the model runs as a set of hosted Model Context Protocol (MCP) tools — `forecast_uk`, `latest_shocks`, `model_summary` — so that forecasts and structural readings are callable from AI agents and ordinary scripts alike.

Section 2 places the model in the Bayesian-VAR, sign-restriction and central-bank-suite literatures. Section 3 sets out the model, the prior and posterior algebra, the identification algorithm with its importance weights, and the mathematics of the impulse-response, variance-decomposition, historical-decomposition and forecasting outputs. Section 4 identifies the source paper and delimits the replication scope. Section 5 describes the implementation. Section 6 reports the calibration and validation results together with the replication's core structural figures. Section 8 shows a forecast from the 2024Q2 data edge and confronts it, quarter by quarter, with subsequent ONS outturns and the Bank's own contemporaneous and later Monetary Policy Report narratives. Section 9 concludes with limitations. Appendices document the full dataset construction, the prior hyperparameters, and the sampling diagnostics.

[1] Code: https://github.com/PolicyEngine/boe-var-model; hosted interface: https://policyengine-macro.vercel.app/svar/.

2 Related literature

The model sits at the intersection of three literatures: Bayesian VAR estimation for forecasting, set identification of structural shocks by sign and zero restrictions, and the use of empirical suites inside central banks. This section situates the replication in each.

2.1 Bayesian VARs for macroeconomic forecasting

The VAR entered empirical macroeconomics with Sims (1980) as a deliberately atheoretical alternative to large simultaneous-equation models, but unrestricted VARs are profligate with parameters: with $k = 8$ variables and $p = 4$ lags the system of Section 3 has $8 \times (32+1+6) = 312$ conditional-mean coefficients against 126 quarterly observations. The Minnesota prior of Doan et al. (1984) and Litterman (1986) made the VAR usable at this scale by shrinking each equation towards an independent random walk, with tightness decaying in the lag order. Kadiyala and Karlsson (1997) systematised the conjugate normal-inverse-Wishart (NIW) implementation that preserves a closed-form posterior, and Sims and Zha (1998) developed the general dummy-observation technology — including the sum-of-coefficients and single-unit-root priors — through which such priors are imposed by augmenting the data. Bańbura et al. (2010) showed that with shrinkage scaled appropriately in the cross-section, Bayesian VARs remain competitive forecasters even with dozens or hundreds of variables, a result that underpins the modern central-bank practice of running medium and large BVARs alongside structural models. Giannone et al. (2015) closed the loop by treating the prior hyperparameters themselves as parameters to be chosen by maximising the closed-form marginal likelihood — the approach the replication implements as an option (Section 3.2) alongside fixed default tightness. The Covid-19 observations pose a well-known outlier problem for all of these systems; the replication follows the dummy-variable spirit of the pandemic prior of Cascaldi-Garcia (2022), absorbing 2020Q1–2021Q2 with one exogenous dummy per quarter.

2.2 Sign and zero restrictions in structural VARs

Identification by sign restrictions descends from Faust (1998), Canova and De Nicoló (2002) and, canonically, Uhlig (2005), who replaced zero-style exclusion restrictions with inequality restrictions on impulse responses and accepted that the structural impact matrix is then only *set*-identified: many rotations Q of the Cholesky factor are consistent with the data and the restrictions. Rubio-Ramírez et al. (2010) provided the general theory of identification for SVARs together with efficient algorithms for drawing orthogonal matrices from the Haar measure. Arias et al. (2018) is the immediate methodological source for this model: it gives an exact algorithm for drawing from the posterior of structural parameters under *combined* zero and sign restrictions, constructing each column of Q in the null space of its zero restrictions and correcting, via importance weights, for the fact that the recursive construction does not sample uniformly from the restricted set. Section 3.3.1 states the algorithm as implemented, and Section 5 reports the weights' effective sample size.

The sign-restriction approach has attracted well-known critiques. Fry and Pagan (2011) emphasise that reporting pointwise medians across accepted draws mixes different structural models ("median is not a model"); the replication follows the source paper in reporting pointwise medians with 68 and 90 per cent bands, so this caveat applies here exactly as it applies to the Bank's own outputs. Baumeister and Hamilton (2015) show that the seemingly agnostic

uniform (Haar) prior over rotations is informative about impulse responses, and argue for explicit priors on structural parameters; the importance-weighting step of Arias et al. (2018), which the replication implements, addresses the related but distinct problem of sampling uniformly *conditional* on the zero restrictions, not the deeper prior-sensitivity point. Antolín-Díaz and Rubio-Ramírez (2018) add narrative restrictions on the signs of specific historical shocks — a natural extension for this model given the clear 2008 and 2022 episodes in its estimated shock series — and Chan et al. (2025) develop the permutation-invariance results that the source paper and the replication use to reduce rejections without distorting the posterior. Kilian and Lütkepohl (2017) provide the standard book-length treatment of the identification, inference and decomposition machinery used throughout Section 3.

2.3 VARs in central-bank forecasting suites

Central banks rarely rely on a single model. The Bank of England's forecasting platform is organised around a central organising model — COMPASS, a New Keynesian DSGE — surrounded by a large "suite" of statistical models used for data-consistent forecasting, cross-checking and risk assessment, of which Bayesian VARs are a prominent part (Burgess et al., 2013). The SVAR replicated here is a modern member of that ecosystem: as Brignone and Piffer (2025) describe, it is run every forecast round to give the Monetary Policy Committee a *structural* reading of incoming data — which shocks moved the forecast — rather than a policy simulation. Its companion methodology paper (Brignone and Piffer, 2026) generalises the forecast-revision decomposition of Section 3.5. The conditional-forecasting tradition of Waggoner and Zha (1999), in which VAR forecasts are conditioned on assumed paths for some variables (market interest-rate paths, energy futures), is the natural next step for the replication and is discussed as an extension in Section 9; the machinery of Section 3.5 already contains the required moving-average representation. Within PolicyEngine Macro the SVAR plays the same role in miniature: it is the baseline and conditioning member of an open suite (Ghenis and PolicyEngine contributors, 2025), alongside reform-scoring models, mirroring at open-source scale the division of labour that Burgess et al. (2013) describe inside the Bank.

3 The model

3.1 Reduced form

Let y_t be the $k \times 1$ vector ($k = 8$) of quarterly variables listed in Table 1. The reduced form is a VAR(p) in (mostly log) levels with a constant and six Covid-19 dummies:

$$y_t \;=\; c \;+\; \sum_{l=1}^{p} \Pi_l \, y_{t-l} \;+\; \sum_{j=1}^{6} \delta_j \, d_{j,t} \;+\; u_t, \qquad u_t \sim \mathcal{N}(0, \Sigma), \tag{1}$$

where each dummy $d_{j,t}$ equals one in exactly one of the quarters 2020Q1–2021Q2 and zero otherwise, absorbing the extreme pandemic observations in the spirit of the pandemic prior of Cascaldi-Garcia (2022). The source paper does not state the lag length; the replication uses $p = 4$ as a configurable default.

Reduced-form innovations map to orthonormal structural shocks through an impact matrix B:

$$u_t = B \, \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, I_k), \qquad \Sigma = BB'. \tag{2}$$

Table 1: Variables, transformations and ordering (source paper Table 1)

#	Variable	Enters model as	Shown in figures as
1	Real world GDP (UK-trade-weighted)	$100 \cdot \log$	YoY growth
2	World CPI (UK-trade-weighted)	$100 \cdot \log$	YoY growth
3	Real oil price in sterling	$100 \cdot \log$	YoY growth
4	Bank Rate	level (%)	level
5	Sterling effective exchange rate (ERI)	$100 \cdot \log$	level
6	UK CPI, seasonally adjusted (CPISA)	$100 \cdot \log$	YoY growth
7	UK CPI Energy	$100 \cdot \log$	YoY growth
8	Real UK GDP	$100 \cdot \log$	YoY growth

B is parameterised through the Cholesky factor of Σ and an orthogonal rotation:

$$B = \mathrm{chol}(\Sigma)\, Q, \qquad Q'Q = I_k. \tag{3}$$

Any orthogonal Q delivers the same reduced-form fit; identification consists of restricting the admissible set of Q.

3.2 Prior and posterior

Estimation is Bayesian in the tradition of Doan et al. (1984) and Litterman (1986), implemented in the conjugate normal-inverse-Wishart (NIW) form of Kadiyala and Karlsson (1997) and Giannone et al. (2015), with the prior imposed through dummy observations as in Sims and Zha (1998). Stack the system as $Y = X\mathcal{B} + U$, where row t of X is $x_t = (y'_{t-1}, \ldots, y'_{t-p}, 1, d'_t)'$ of length $m = kp + 1 + 6$ and $\mathcal{B}$ is $m \times k$ (so $\Pi = \mathcal{B}'$). The NIW prior

$$\Sigma \sim \mathcal{IW}(S_0, d_0), \qquad \mathrm{vec}(\mathcal{B}) \mid \Sigma \sim \mathcal{N}\big(\mathrm{vec}(\mathcal{B}_0),\ \Sigma \otimes \Omega_0\big), \tag{4}$$

is encoded as artificial data (Y_d, X_d) appended to the sample. The replication's dummy blocks, written for equation scales s_i (the residual standard deviations of univariate AR(1) regressions, a standard pre-sample device), are:

- *Minnesota lag shrinkage.* For each lag $l = 1, \ldots, p$, the block $Y_d^{(l)} = \mathbf{1}\{l{=}1\}\ \mathrm{diag}(\delta_i s_i)/\lambda$, $X_d^{(l)} = [\,0 \cdots \mathrm{diag}(s_i)\, l^2/\lambda \cdots 0\,]$ implies prior mean $\delta_i = 1$ on each variable's own first lag (a random walk, appropriate for the log-levels specification) and zero elsewhere, with prior standard deviation $\lambda\, s_i/(l^2 s_j)$ for the coefficient on variable j at lag l in equation i: shrinkage tightens quadratically in the lag and is scale-adjusted across equations.
- *Inverse-Wishart scale.* A block $Y_d = \mathrm{diag}(s_i)$, $X_d = 0$ centres Σ on $\mathrm{diag}(s_i^2)$.
- *Sum-of-coefficients (no-cointegration) prior* (Doan et al., 1984): $Y_d = \mathrm{diag}(\delta_i \bar{y}_i)/\mu$ with the same block repeated across all p lag positions of X_d, where $\bar{y}$ is the mean of the p initial observations. As $\mu \to 0$ this forces unit sums of own-lag coefficients (one unit root per variable); $\mu = 1$ applies it loosely.
- *Single-unit-root (dummy-initial-observation) prior* (Sims and Zha, 1998): one row $Y_d = \bar{y}'/\theta$, $X_d = [\bar{y}' \cdots \bar{y}'\ \ 1/\theta\ \ 0']$, active at $\theta = 1$, pulling the system towards a common stochastic trend passing through the initial conditions.

- *Deterministic block.* The constant and the six Covid dummies receive a near-flat prior (dummy weight 10^{-3}), so the pandemic quarters are absorbed by essentially unrestricted dummy coefficients — the exogenous-dummy simplification of the Cascaldi-Garcia (2022) pandemic prior.

With augmented data $(Y_*, X_*) = ([Y; Y_d], [X; X_d])$, the posterior is again NIW:

$$\Sigma \mid Y \sim \mathcal{IW}\big(S_*,\ T_* - m\big), \qquad \operatorname{vec}(\mathcal{B}) \mid \Sigma, Y \sim \mathcal{N}\big(\operatorname{vec}(\hat{\mathcal{B}}_*),\ \Sigma \otimes (X_*'X_*)^{-1}\big), \tag{5}$$

with $\hat{\mathcal{B}}_* = (X_*'X_*)^{-1}X_*'Y_*$, $S_* = (Y_* - X_*\hat{\mathcal{B}}_*)'(Y_* - X_*\hat{\mathcal{B}}_*)$ and T_* the augmented sample size. Sampling is exact and direct (no MCMC): Σ from the inverse Wishart, then $\mathcal{B} = \hat{\mathcal{B}}_* + C^{-\prime}Z\ \operatorname{chol}(\Sigma)'$ with C the Cholesky factor of $X_*'X_*$ and Z an $m \times k$ standard normal matrix, so posterior draws are i.i.d. and no convergence diagnostics beyond the sign-restriction acceptance statistics of Appendix C are required.

The replication's defaults are $\lambda = 0.2$, $\mu = 1$, $\theta = 1$ (the source paper does not report its hyperparameters). Because the prior is conjugate, the marginal likelihood is available in closed form as a ratio of NIW normalising constants (Giannone et al., 2015, eq. 5),

$$\log p(Y) = -\tfrac{Tk}{2}\log\pi + \log\frac{\Gamma_k(\frac{d_1}{2})}{\Gamma_k(\frac{d_0}{2})} + \tfrac{k}{2}\big(\log|X_d'X_d| - \log|X_*'X_*|\big) + \tfrac{d_0}{2}\log|S_0| - \tfrac{d_1}{2}\log|S_1|, \tag{6}$$

with $d_0 = T_d - m$ and $d_1 = T + T_d - m$ the prior and posterior degrees of freedom and S_0, S_1 the corresponding scale matrices. The codebase exposes an optional empirical-Bayes step that maximises (6) over (λ, μ, θ) by grid search refined with Nelder–Mead in log space, the Giannone et al. (2015) procedure; Appendix B reports the resulting values next to the defaults.

3.3 Identification: zero and sign restrictions

Six of the eight shocks are identified — world demand (WD), world energy (WE), world supply (WS), UK demand (UD), UK supply (US) and UK monetary policy (MP) — while two are left unidentified, one global (U1) and one UK-specific (U2), to absorb residual volatility. All restrictions apply to the impact matrix B only. Table 2 reproduces the restriction matrix (Table 2 of the source paper) exactly as encoded in the replication.

The economic content is conventional: a positive world demand shock raises world and UK output and prices; an expansionary (disinflationary) world energy shock raises real activity while lowering oil prices, world CPI and UK CPI including its energy component; a positive world supply shock raises output and lowers prices with a zero energy-price impact; a UK demand shock raises UK GDP, CPI and Bank Rate; a UK supply shock raises GDP and lowers CPI; and a monetary tightening raises Bank Rate while lowering CPI and GDP.

3.3.1 The sampling algorithm

For each posterior draw (Π, Σ), Q is drawn with the algorithm of Arias et al. (2018) so that the *zero* restrictions hold exactly by construction. Let $L = \operatorname{chol}(\Sigma)$ and let Z_j denote the set of variables restricted to zero on impact for shock j; the restriction $B_{vj} = 0$ reads $L_{v\cdot}\, q_j = 0$, a linear constraint on the jth column of Q. Columns are drawn recursively, most-restricted first, each as a standard normal vector projected onto the intersection of the null space of its zero constraints and the orthogonal complement of the previously drawn columns. The *sign*

Table 2: Identifying restrictions on impact (source paper Table 2)

Variable	WD	WE	WS	U1	UD	US	MP	U2
World GDP	+	+	+		0	0	0	0
World CPI	+	−	−		0	0	0	0
Oil price		−			0	0	0	0
Bank Rate					+		+	
Exchange rate								
UK CPISA	+	−	−		+	−	−	
CPI Energy	+	−	0					
Real UK GDP	+	+	+		+	+	−	

Notes: Blank cells are unrestricted. The zero block on the three world variables encodes the small-open-economy assumption: UK shocks (UD, US, MP) and the UK unidentified shock (U2) have no contemporaneous impact on world variables. The global unidentified shock U1 is exempt from that block — that exemption is what makes it "global". The zero on CPI Energy under WS exogenises energy prices with respect to the world supply shock, separating it from the world energy shock.

Algorithm 1 Zero-and-sign-restriction sampler (one candidate per posterior draw)

1: Draw (Π, Σ) from the NIW posterior (5); set $L \leftarrow \text{chol}(\Sigma)$.
2: Order shocks by number of zero restrictions, most first: $j_1, \ldots, j_k$.
3: **for** $n = 1, \ldots, k$ **do**
4: Stack $M \leftarrow \left[\{L_{v\cdot} : v \in Z_{j_n}\};\ q'_{j_1}, \ldots, q'_{j_{n-1}}\right]$.
5: Draw $z \sim \mathcal{N}(0, I_k)$; project $z \leftarrow NN'z$ with N an orthonormal basis of null(M).
6: Normalise $q_{j_n} \leftarrow z/\|z\|$.
7: **end for**
8: $B \leftarrow LQ$.
9: **Permutation search** (Chan et al., 2025): within each block of shocks sharing an identical zero-restriction set — {WD, WE, U1}, {WS}, {UD, US, MP, U2} — enumerate column permutations and ± 1 sign flips; keep a perfect matching in which every sign-restricted shock is assigned a column satisfying its restrictions (chosen uniformly at random if several exist).
10: **if** no valid matching exists in some block **then reject** the entire draw (Π, Σ, Q).
11: **else accept**; compute the importance weight w of (7).
12: **end if**

restrictions are then checked on $B = LQ$ by accept/reject as in Uhlig (2005) and Rubio-Ramírez et al. (2010). Algorithm 1 states the procedure as implemented.

Three implementation details matter for correctness. First, exactly *one* rotation is attempted per posterior draw and rejection discards the pair (Σ, Q) jointly: retrying rotations until one is accepted for a given Σ would oversample reduced-form draws for which acceptance is hard, targeting the wrong posterior. Second, the permutation search exploits the invariance results of Chan et al. (2025): the Haar measure is invariant under column permutations and sign flips, but with zero restrictions a column may only be relabelled to a shock position with the *identical* zero-restriction set, so the replication conservatively permutes only within the three blocks listed in Algorithm 1; a post-condition asserts that every zero restriction survives the relabelling. Third, sign flips of the two unidentified columns are always available, so the accept/reject step binds only through the six identified shocks' sign sets.

3.3.2 Importance weights

The recursive construction of Algorithm 1 does not draw Q uniformly (with respect to the Haar measure) from the zero-restricted set: conditioning each column on the previously drawn ones distorts the density. Arias et al. (2018) show that the draw can be corrected by importance weights equal to the ratio of two volume elements. Writing structural parameters as $x = (\mathrm{vec}(A_0), \mathrm{vec}(A_+))$ with $A_0 = (L')^{-1}Q$ and $A_+ = \mathcal{B}A_0$, the weight of an accepted draw is

$$w \;\propto\; \frac{v_f}{v_{g\circ f}}, \qquad \log v_f = -(2k+m+1)\,\log|\det A_0|, \qquad \log v_{g\circ f} = \tfrac{1}{2}\log\det\left(D_N'D_N\right), \tag{7}$$

where $D_N = D_{g\circ f}\,N_z$ is the Jacobian of the map from structural parameters to the reduced-form parameters and the sphere coordinates of the recursive construction, restricted (via an orthonormal basis N_z of the null space of the zero-restriction Jacobian D_z) to the zero-restricted manifold. The replication computes $D_{g\circ f}$ and D_z by central finite differences, parameterising Σ through its $k(k+1)/2$ free elements $\mathrm{vech}(\Sigma)$ so that D_N has exactly full column rank, and using smooth (projection-based) null-space bases so the finite differencing is well defined — two documented refinements relative to the reference `bsvarSIGNs` implementation it was ported from. Weights are normalised to mean one across accepted draws, all posterior quantiles are weighted quantiles, and the Kish effective sample size $\mathrm{ESS} = (\sum_i w_i)^2/\sum_i w_i^2$ is reported (Appendix C). Because each weight requires numerical Jacobians of a $k(k+m)$-dimensional map, the weights are computed on the validation and production paths but not for the exploratory figure runs, which Section 6 shows shift one-year FEVD shares by only a few percentage points.

3.4 Outputs

Impulse responses. Write (1) in companion form: stacking $s_t = (y_t', \ldots, y_{t-p+1}')'$,

$$s_t = F s_{t-1} + J'(c + \delta' d_t + u_t), \qquad F = \begin{pmatrix} \Pi_1 & \cdots & \Pi_{p-1} & \Pi_p \\ I_{k(p-1)} & & & 0 \end{pmatrix}, \qquad J = \begin{pmatrix} I_k & 0 \end{pmatrix}, \tag{8}$$

so that $\Psi_h = JF^hJ'$ is the reduced-form moving-average coefficient at horizon h: iterating (8) forward and selecting the first block gives $y_{t+h} = \sum_{s=0}^{h}\Psi_s u_{t+h-s} + (\text{deterministic terms})$, the Wold representation truncated at the sample start. Substituting $u_t = B\varepsilon_t$, the structural impulse response of variable i to a one-standard-deviation shock j at horizon h is

$$\mathrm{IRF}_{ij}(h) = \left[\Psi_h B\right]_{ij}, \tag{9}$$

reported as pointwise posterior medians with 68 and 90 per cent credible bands across accepted (Π, Σ, Q) triples (importance-weighted quantiles where the weights are computed).

Forecast-error variance decomposition. The forecast error at horizon H is

$$y_{t+H} - \mathbb{E}_t y_{t+H} = \sum_{h=0}^{H}\Psi_h B\,\varepsilon_{t+H-h};$$

because the structural shocks are orthonormal, its variance for variable i splits additively across shocks, giving the FEVD at horizon H:

$$\mathrm{FEVD}_{ij}(H) \;=\; \frac{\sum_{h=0}^{H}\left[\Psi_h B\right]_{ij}^2}{\sum_{h=0}^{H}\left[\Psi_h\Sigma\Psi_h'\right]_{ii}}, \qquad \sum_{j=1}^{k}\mathrm{FEVD}_{ij}(H) = 1. \tag{10}$$

where the equality of the denominator with $\sum_j$ of the numerators follows from $\Sigma = BB'$, so the shares sum to one mechanically — a property the test suite nevertheless verifies numerically (Section 6).

Historical decomposition. Estimated structural shocks are recovered as $\hat{\varepsilon}_t = B^{-1}\hat{u}_t$ at each accepted draw, where $\hat{u}_t = y_t - \hat{c} - \sum_l \hat{\Pi}_l y_{t-l} - \sum_j \hat{\delta}_j d_{j,t}$ are that draw's reduced-form residuals. Solving (8) backward from the first post-lag observation splits each variable's path into the cumulative contribution of each structural shock, the Covid-dummy component and the deterministic component:

$$y_{i,t} \;=\; \underbrace{\sum_{j=1}^{k}\sum_{h=0}^{t-1}\big[\Psi_h B\big]_{ij}\,\hat{\varepsilon}_{j,t-h}}_{\text{structural shocks}} \;+\; \underbrace{y_{i,t}^{\text{covid}}}_{\text{dummies}} \;+\; \underbrace{y_{i,t}^{\text{det}}}_{\text{initial conditions, constant}}, \tag{11}$$

an additive identity that must reconstruct the data exactly — a property Section 6 verifies to 10^{-8}. The Covid component $y_{i,t}^{\text{covid}}$ is obtained by feeding the estimated dummy effects $\hat{\delta}\, d_t$ through the companion recursion (8) with zero shocks, so that the dynamic propagation of the pandemic dummies is attributed to the dummies rather than to the deterministic trend.

3.5 Forecasting and the forecast-revision decomposition

Unconditional forecasts. Conditional on a draw (Π, Σ) and data through the forecast origin T, the point forecast iterates (1) forward with zero shocks and zero future dummies:

$$\hat{y}_{T+h|T} \;=\; c + \sum_{l=1}^{p} \Pi_l\, \hat{y}_{T+h-l|T}, \qquad \hat{y}_{T+s|T} = y_{T+s} \text{ for } s \le 0. \tag{12}$$

The predictive *distribution* adds both sources of uncertainty: across accepted posterior draws (parameter uncertainty) and, within each draw, stochastic paths

$$y_{T+h} = \hat{y}_{T+h|T} + \sum_{s=0}^{h-1} \Psi_s u_{T+h-s}$$

with $u_{T+h} \sim \mathcal{N}(0, \Sigma)$ (future-shock uncertainty). The replication simulates five stochastic paths per accepted draw and reports weighted pointwise medians with 68 and 90 per cent bands; growth variables are converted to year-on-year rates as $y_t - y_{t-4}$, exact in the $100 \cdot \log$ parameterisation. This is precisely the construction of the fan charts in Section 8.

The structural reading at the data edge. Because the model is linear, the reduced-form innovation at the most recent observation T maps into that quarter's structural shocks, $\hat{\varepsilon}_T = B^{-1}\big(y_T - \hat{c} - \sum_l \hat{\Pi}_l y_{T-l}\big)$, one vector per accepted draw. The posterior distribution of $\hat{\varepsilon}_T$ across draws — summarised by weighted sign probabilities $P(\varepsilon_{j,T} > 0)$ and magnitude quantiles — is the `latest_shocks` output: a probabilistic statement about what hit the economy in the latest quarter.

Forecast revision. The Bank's round-by-round use case (Brignone and Piffer, 2025, 2026) decomposes the change in the forecast between two data edges into shock contributions. For one draw, let $\hat{y}_{T+h|T}$ be the forecast from data through T and $\hat{y}_{T+h|T-1}$ the pseudo-forecast from data through $T-1$ using the *same* parameters. Linearity of (12) in the state and $u_T = B\varepsilon_T$ give the exact identity

$$\hat{y}_{T+h|T} \;=\; \hat{y}_{T+h|T-1} \;+\; \Psi_h B\,\hat{\varepsilon}_T, \qquad h = 0, 1, \ldots, H, \tag{13}$$

(with $\hat{y}_{T|T} = y_T$ at $h = 0$): the entire forecast revision is the composite impulse response $\Psi_h B\hat{\varepsilon}_T$, which splits additively across the eight shocks as $\sum_j \hat{\varepsilon}_{j,T}\,[\Psi_h B]_{\cdot j}$. Identity (13) holds *per draw* to machine precision (verified to 5.3×10^{-12}; Section 6), so the published decomposition is exact rather than approximate. In the year-on-year units of the figures, the level composite is differenced, YoY effect(h) = level(h) − level$(h-4)$ with the level effect zero before impact.

Conditional forecasts. The same moving-average algebra supports conditioning on assumed future paths for a subset of variables (a market interest-rate path, an energy-price assumption) in the sense of Waggoner and Zha (1999): stacking the conditions $\bar{y}_{i,T+h} = \hat{y}_{i,T+h|T} + [\sum_s \Psi_s B\varepsilon_{T+h-s}]_i$ as linear restrictions $R\varepsilon = r$ on the future shocks $\varepsilon = (\varepsilon'_{T+1}, \ldots, \varepsilon'_{T+H})'$, the restricted shocks are distributed $\mathcal{N}\big(R'(RR')^{-1}r,\; I - R'(RR')^{-1}R\big)$, i.e. the minimum-norm shock configuration consistent with the conditions plus orthogonal noise. The replication exposes the ingredients (Ψ_h, B, the forecast iterator) but does not yet ship a conditional-forecast tool; Section 9 lists it as the natural next extension, and within a structural model the choice of *which* shocks are allowed to implement the conditions is an economic, not mechanical, decision.

4 The source paper and replication scope

The model replicated here is documented in Brignone and Piffer (2025): Davide Brignone and Michele Piffer, *A structural VAR model for the UK economy*, Bank of England Macro Technical Paper No. 3, published 18 July 2025. It is one of the first papers in the Bank's Macro Technical Paper series, which documents the modelling toolkit supporting the Monetary Policy Committee, and it describes the SVAR used on a round-by-round basis to provide a structural narrative for revisions to the Monetary Policy Report forecast. A companion paper by the same authors — Staff Working Paper No. 1,165, January 2026 — develops the forecast-revision methodology in general form (Brignone and Piffer, 2026); Section 7.2 compares the replication against it.

Replication scope. The replication targets the paper's core pipeline and its standard outputs (the paper's Figures 2–6): Bayesian estimation, identification, impulse responses, FEVDs, estimated structural shocks and historical decompositions, plus the Section-5 forecast-revision machinery in reduced scope. The estimation sample matches the paper exactly: **1992Q1–2023Q2** (126 quarters), beginning with UK inflation targeting and ending a year before the data edge because recent quarters are revision-prone. Data through **2024Q2** are retained for the forecasting and forecast-revision exercises, matching the vintage of the Bank's August 2024 Monetary Policy Report round (Bank of England, 2024). *Update (July 2026):* the production run was re-estimated on the extended **1992Q1–2025Q1** sample (133 quarters), and the headline FEVD shares quoted in this paper (37.4 and 42.3 per cent) are from that extended run; the original-sample replication figures remain in the appendix.

Approximated internal series. Two of the eight variables — UK-trade-weighted world GDP and world CPI — are internal Bank series that are not published. The replication proxies them with trade-weighted aggregates constructed from OECD and IMF sources, and constructs the real sterling oil price from Brent crude converted to sterling and deflated by UK CPI, since the paper does not fully specify its construction. UK series (Bank Rate, sterling ERI, CPISA, CPI Energy, real GDP) come from the Bank of England and the Office for National Statistics (Office for National Statistics, 2026c, 2025a). Because the world aggregates are proxies and the Bank released no replication package, the global block can be validated only against the published figures, not against the underlying data — a caveat the quantitative comparison in Section 6 makes explicit.

Documented deviations. Beyond the proxy world series, the replication documents its deviations openly: an assumed lag length of $p = 4$ (unstated in the paper), default rather than hierarchically optimised prior hyperparameters, Covid dummies as exogenous regressors rather than the full pandemic-prior algebra, permutation search restricted to zero-pattern-equivalent groups, and current data vintages in place of the August 2024 vintages. Each deviation is a configurable parameter or a flagged approximation in the codebase, not a silent choice.

5 Method and implementation

5.1 Code architecture

The replication is a small, dependency-light Python package (`boe_var`, src layout; NumPy, SciPy, pandas, matplotlib) with a module per stage of the pipeline:

- `data.py` — `load_data()` returns the transformed quarterly panel — a `PeriodIndex` DataFrame in the ordering of Table 1, $100 \cdot \log$ except Bank Rate — assembled by a separate download script from public sources;
- `bvar.py` — the `BVAR` class: NIW posterior sampling of (Π, Σ) with Minnesota and sum-of-coefficients priors and exogenous Covid dummies, plus companion-form and residual helpers;
- `identification.py` — the restrictions of Table 2 as machine-readable constants (`ZERO_RESTRICTIONS` and `SIGN_RESTRICTIONS`), the Arias et al. (2018) `draw_Q` routine, and `identify()`, which pairs posterior draws with accepted rotations;
- `analysis.py` — IRFs, FEVDs, estimated shocks, historical decompositions, forecast bands and the composite (forecast-revision) impulse response, with plotting helpers.

An end-to-end script reproduces the paper's Figures 2–6 and writes a numerical summary; the test suite (Section 6) runs the same code on the real dataset with a fixed seed.

5.2 Estimation and identification runs

The production configuration takes **10,000** posterior draws, of which **765** survive the sign-restriction accept/reject step (a 7.6 per cent acceptance rate), with $p = 4$ lags and hyperparameters $\lambda = 0.2$, $\mu = 1$. Zero restrictions are satisfied by construction on every draw; sign

restrictions are enforced by rejection, so they too hold exactly on every accepted draw — neither fact is a test (Section 6). Because the Arias et al. (2018) zero-restriction sampler does not draw Q uniformly from the restricted set, the slow validation path additionally computes the algorithm's importance weights (each requiring a numerical Jacobian per accepted draw); the importance-weighted effective sample size is **368.0**, and weighted and unweighted results are compared in Section 6. The forecast-revision exercise re-estimates through 2024Q2 with 6,000 draws (487 accepted, weighted ESS 207.5). The hosted interactive tools default to a smaller 2,000-draw run on the extended sample; the 10,000-draw run described here is the publication artifact. Its figures, and those of the rolling-origin evaluation in Section 6, are the ones committed in the model repository by its 2026Q2 data refresh (pull request 24; data through 2026Q2, estimation sample 1992Q1–2025Q1, so the five quarters 2025Q2–2026Q2 are held out of estimation). The values before that refresh are at revision `845e2a4`.

5.3 Hosted MCP tools

Within PolicyEngine Macro the model is wrapped by an integration layer — a `pe-macro` command-line interface and a Model Context Protocol (MCP) server, deployed automatically to Modal on every merge — exposing three tools:

- `forecast_uk` — unconditional forecasts of the eight variables from the data edge, with median paths and credible bands, in year-on-year units for growth variables;
- `latest_shocks` — the posterior sign probabilities and magnitudes of the six identified structural shocks in the most recent quarter, i.e. a structural reading of "what just hit the economy";
- `model_summary` — the model card: specification, sample, restriction matrix, acceptance statistics and validation status.

The same tools back the model page at `https://policyengine-macro.vercel.app/svar/`. This makes the SVAR usable as the suite's conditioning member: agents and scripts can pull a structural baseline before invoking the reform-scoring models.

6 Calibration and validation

Validation is against ground truth — the restrictions and published figures of Brignone and Piffer (2025) — rather than the mere absence of exceptions. The checks fall into three tiers: exact invariants that must hold on every draw, sign checks on quantities the identification does *not* restrict, and quantitative comparisons with the paper's reported magnitudes inside deliberately wide bands that respect the replication's proxy data and Monte-Carlo noise. All checks run in continuous integration on the real dataset.

Exact invariants — checks that cannot fail, and are reported as such. On *every* accepted identification draw: (i) each element of B restricted to zero in Table 2 is zero to machine precision (largest violation below 10^{-9}); (ii) every sign restriction holds strictly ($s \cdot B_{vj} > 0$); (iii) $BB' = \Sigma$ within 10^{-7}, so every accepted B is a valid structural factorisation. None of these three can come out any other way, and the numbers should be read accordingly: the zeros are imposed by the Arias et al. (2018) construction, the signs *define* acceptance (a draw violating

one is discarded, so the pass rate over accepted draws is one by definition), and $BB' = \Sigma$ is the factorisation identity the construction starts from. What they test is that the implementation does what the algorithm says — a floating-point and indexing check on the code, not evidence that this replication reproduces the Bank's model. The residuals quoted (10^{-9}, 10^{-7}) are solver tolerances, not accuracy against any external quantity. The one check in this group that carries information is the last: median impact impulse responses also respect every Table 2 sign, which is not guaranteed draw-by-draw acceptance and could in principle fail.

Unrestricted sign checks. The paper's key overidentifying check is that the real oil price — whose response to a world supply shock is left *unrestricted* — rises after a positive world supply shock. The replication reproduces this: the median impact response of the sterling real oil price to a one-standard-deviation world supply shock is +5.83 ($100 \cdot \log$ units), positive as in the paper's Figure 2.

Comparison with published magnitudes. Table 3 collects the model-versus-paper comparison. The headline comparison is the share of UK forecast-error variance attributed to the identified global shocks (world demand + energy + supply) at the one-year horizon, where the paper reports roughly 40 per cent for UK GDP and roughly 50 per cent for UK CPI.

Getting that comparison like-for-like took three corrections, all made in August 2026 and all of which had been inflating the replication's agreement with the paper. First, the *horizon*: the decomposition of equation (10) accumulates squared moving-average coefficients, so its column h holds the variance of the $(h+1)$-step-ahead forecast error. "One year" is therefore column 3, not column 4; reading column 4 was reading five quarters. Second, the *statistic*: the shares were formed by summing the six identified shocks' pointwise posterior *medians* and renormalising the result to 100 per cent. The per-shock medians sum to about 0.62 (GDP) and 0.71 (CPI) here, so that renormalisation inflated every identified share by roughly a third. Third, the *target*: the paper's Figure 4 decomposes "the mean of the sum of the decomposition", and its note attributes the residual to 100 per cent to the unidentified shocks — its text states that the six identified shocks explain around 80 per cent of the domestic variables, of which the global shocks are around 40 per cent of real GDP and 50 per cent of CPI variation one year after the shocks. Its shares are therefore posterior *means* of shares of *total* forecast-error variance, renormalised over nothing.

Computed the paper's way — the group share formed on each accepted draw, averaged, four quarters ahead — the 10,000-draw production run gives **38.7** per cent for UK GDP against the paper's ∼40, a match, and **41.5** per cent for UK CPI against ∼50, about eight and a half points short. The previously published pair, 42.1 and 49.5 per cent, appeared to match on both; the CPI agreement was the renormalisation, not the model. What the shortfall does *not* support is a strong claim in either direction: the 68 per cent posterior band on these shares is $[24.0, 53.1]$ for GDP and $[25.6, 56.3]$ for CPI, roughly ±15 points, wider than the 8.5-point gap being discussed. The identified shocks in total explain about 77 per cent of UK GDP and 78 per cent of UK CPI one-year variance here, against the paper's "around 80 per cent" — so the disagreement is about the *split* between global and domestic, not about how much the six shocks explain.

The $[30, 60]$ per cent acceptance band used by the per-commit test suite should be read for what it is: a regression tripwire, not a calibration result. A band 30 points wide around targets of 40 and 50 admits values from three-quarters below to half again above the published number, so "in band" is close to a claim that cannot fail; it is set to catch a gross code regression without

Table 3: Model versus paper: validation summary

Check	Replication	Paper	Verdict
Cannot fail — implementation checks, reported for completeness only			
Zero restrictions (Table 2)	exact, every draw ($< 10^{-9}$)	imposed	holds by construction
Sign restrictions (Table 2)	every accepted draw	imposed	defines acceptance
$BB' = \Sigma$	every draw ($< 10^{-7}$)	by construction	factorisation identity
FEVD rows sum to one	$< 10^{-10}$, all horizons/draws	identity	algebraic identity
Historical decomposition reconstructs data	$< 10^{-8}$	identity	algebraic identity
Forecast-revision adding-up identity	$< 5.3 \times 10^{-12}$	identity	algebraic identity
Can fail — comparisons against the publication			
Sign restrictions on the median IRF	every restricted response	imposed per draw	pass
Oil price after world supply shock (impact)	+5.83 (unrestricted)	positive (Fig. 2)	pass
Global share, UK GDP FEVD, 1 yr	38.7% [24.0, 53.1]	∼40%	match
Global share, UK CPI FEVD, 1 yr	41.5% [25.6, 56.3]	∼50%	∼8.5 pp short
Identified shocks in total, UK GDP / CPI, 1 yr	76% / 79%	"around 80%"	match
Shock narrative (2008, 2022–23 episodes)	reproduced	Fig. 5	pass
Largest domestic CPI contributor	UK supply (13.9% vs 8.7%)	monetary policy (prose)	unresolved

Notes: The first block is separated because those six checks have no failing branch: zeros and $BB' = \Sigma$ are properties of the Arias et al. (2018) construction, sign restrictions are the accept/reject rule itself, and the last three are algebraic identities of equations (10), (11) and (13). The residuals are solver tolerances, not accuracy. FEVD rows: posterior mean of the group share formed on each accepted draw at the 4-quarter-ahead forecast-error variance, as a share of total variance — the statistic the paper's Figure 4 plots — with 68 per cent posterior bands in brackets; production run of 10,000 posterior draws, 765 accepted, importance-weighted ESS 368.0. The final row is marked unresolved rather than a known defect: the paper states the monetary-policy ranking in prose but publishes no FEVD table, and its Figure 4 does not obviously show it for CPI; percentages are this replication's median 1-year CPI shares for UK supply and UK monetary policy. Proxy world aggregates remain a candidate explanation.

going flaky on Monte-Carlo variation and the proxy world aggregates, and passing it is not evidence of agreement. The informative comparison is the point estimate against ∼40 / ∼50 with its posterior band, above.

Nor do the importance weights explain the remaining distance. The cheap unweighted CI configuration (600 draws, 51 accepted) gives 40.1 and 40.1 per cent, but it is a noisy estimator — across sampling seeds its global GDP share has a standard deviation of about 6.5 points at that draw count — and the unweighted 3,000-draw run behind Figures 1–3 gives 38.1 and 42.3 per cent, essentially the weighted production values. The gap between the CI configuration and the production run is therefore Monte-Carlo noise in the accepted sample, not the weights correcting the Arias et al. (2018) sampler towards the uniform-over-rotations posterior; an earlier version of this paper attributed it to the weights.

Internal consistency. Identities are verified to numerical precision rather than assumed: FEVD shares sum to one at every horizon on every draw within 10^{-10} (equation 10); the

historical decomposition (11) — structural shocks plus Covid component plus deterministic component — reconstructs the observed data within 10^{-8}, and the stochastic component equals the sum over shocks within 10^{-10}; and in the forecast-revision exercise the identity "forecast at T = pseudo-forecast at $T-1$ + composite impulse response" holds per draw with a maximum absolute error of 5.3×10^{-12} across all draws, horizons and variables.

Rolling-origin forecast evidence: no demonstrated skill. A separate expanding-window pseudo-out-of-sample evaluation estimates the reduced-form BVAR at 50 historical origins, at horizons one to eight, and compares its posterior-mean path with three benchmarks: no change (a driftless random walk), a random walk *with drift*, and an AR(1). No structural restriction or FEVD target enters this exercise. Read carelessly it is the most flattering evidence in this paper; read properly it is the least.

Three things have to be got right before any of its ratios mean anything, and until August 2026 this section got two of them wrong.

The benchmark. A driftless random walk on a trending log level is a weak opponent, and most of the apparent skill is against it: CPI-level relative RMSE is 0.63 at one quarter and 0.67 at eight (values below one beat the benchmark), while UK GDP is worse than no change at both (1.06 and 1.06). Against the drifting walk the CPI advantage shrinks to 0.83 at one quarter and reverses to 1.03 at eight. Figures from the no-change columns must not be quoted inside a drift-benchmark discussion — the ex-Covid UK GDP ratio of 0.77 and the UK GDP p-value range of 0.38–0.67, both of which have circulated, are no-change numbers; against drift the same quantities are 1.21 and 0.33–0.43.

The specification. The evaluation had always estimated *without* the six Covid dummies (2020Q1–2021Q2) that every published forecast in this paper carries, so it scored a model nobody publishes. Re-run on the published specification, UK GDP against drift goes 1.06 → 0.99 at one quarter and 1.12 → 0.95 at eight — level with the naive benchmark rather than worse, and still not better. CPI moves by less than 0.02, so its weakness at long horizons is not a Covid artefact. Both blocks are committed; the production-specification block is the like-for-like one.

Multiplicity. The exercise runs 8 variables × 8 horizons = 64 Diebold–Mariano tests per benchmark, so the smallest of them is not a discovery. Under a Benjamini–Hochberg false-discovery-rate adjustment across the whole grid, the minimum q-value against the drifting random walk is **0.39** in the headline block and **0.42** under the published specification: *nothing clears a 10 per cent false-discovery rate on any variable at any horizon.* Bank Rate comes closest — relative RMSE 0.78–0.87 against drift and an unadjusted $p = 0.018$ at one quarter, the smallest of the 64 — and does not survive. Only against the weakest benchmark, no change, do any q-values fall below 0.10 (two in the headline block, three under the published specification), and none below 0.05.

The summary is therefore narrower than the earlier one: this model has *not* demonstrated forecast skill against a drifting random walk on any variable at any horizon, and the seven-quarter frozen-edge result of Section 8 is a single origin, not a contradiction of that. The generated evidence is committed as `figures/rolling_evaluation.json` alongside this paper and as `results/rolling_evaluation.json` in the model repository, with its `production_spec` and `multiplicity` blocks.

6.1 Interval coverage across rolling origins (added July 2026)

A companion evaluation, added in July 2026 and committed as `coverage_evaluation.json`, asks whether the predictive *bands* are calibrated, not just whether the point forecasts beat benchmarks. Using the same expanding-window design, the reduced-form BVAR is re-estimated at 49 historical origins; at each origin, 100 posterior draws with 5 stochastic paths per draw generate predictive distributions, and percentile bands on the level of each series are confronted with the subsequent outturns at horizons one to eight.

The bands under-cover. Averaged across the eight variables and eight horizons, the nominal 68 per cent band contains the outturn 62.1 per cent of the time and the nominal 90 per cent band 77.2 per cent of the time, and coverage worsens with horizon: the 68 per cent band covers 71 per cent at one quarter but only 57 per cent at eight quarters, the 90 per cent band 87 per cent at one quarter and 72 per cent at eight. Three caveats bound the reading: the Covid quarters are included with no Covid dummies in the evaluation model; estimation uses final revised data, so the exercise is pseudo- rather than real-time; and consecutive origins overlap, so the effective number of independent observations is well below 49. None of these caveats reverses the direction of the finding — users should treat the published credible bands as too narrow, especially at longer horizons.

6.2 Replication figures

Figures 1–3 show the replication's core structural outputs, generated by `figures/make_figures.py` from a fixed-seed run of 3,000 posterior draws (241 accepted; unweighted, which as shown above moves one-year FEVD shares by well under a point relative to the importance-weighted production benchmark).

Figure 1 plots impulse responses to the three shocks that carry most of the narrative weight. The world demand row shows the classic demand comovement — world GDP, oil prices, UK CPI and UK GDP all rise — with the UK CPI response building over two to three years. The world energy row shows the expansionary (disinflationary) energy shock: activity up, oil and UK CPI down, with the oil-price response large on impact (-6 to -7 in $100 \cdot \log$ units) and persistent. The UK monetary policy row shows a tightening lowering both CPI and GDP with the familiar hump; its world-variable responses are exactly zero on impact by construction and small thereafter, the small-open-economy assumption at work. Only the impact signs in Table 2 are imposed; the persistence of every response is untouched data information.

Figure 2 stacks the FEVD shares for UK GDP and UK CPI by horizon. The pattern behind the comparison in Table 3 is visible directly: at the one-year horizon (marked) the three identified global shocks account for 38.1 per cent of UK GDP variance and 42.3 per cent of UK CPI variance in this unweighted run — within a point of the weighted production figures — with the world demand share of CPI growing steadily in the horizon and the unidentified shocks retaining about a fifth throughout. Domestically, UK supply is the largest identified contributor to GDP variance — and, in this replication's unresolved discrepancy with the paper's prose, also to CPI variance, where the paper says monetary policy contributes most.

Figure 3 gives the historical decomposition of year-on-year UK CPI inflation from 2016 — the model's structural account of the 2021–23 inflation surge and its aftermath. The surge is attributed overwhelmingly to global shocks: contractionary world energy shocks and strong world demand dominate the 2022 bars, with UK-specific shocks and the Covid-dummy component playing supporting roles, and the decomposition unwinding through 2023 as energy shocks

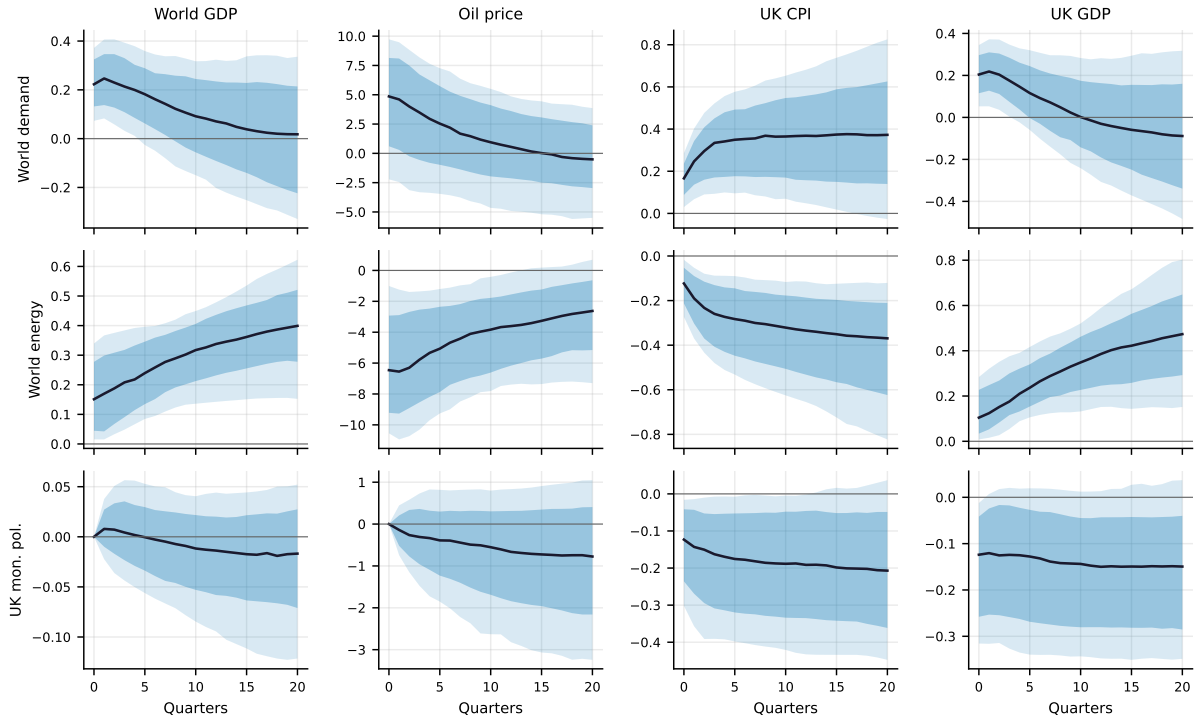


Figure 1: Impulse responses to one-standard-deviation world demand, world energy and UK monetary policy shocks. Solid line: pointwise posterior median; dark and light bands: 68 and 90 per cent credible bands across accepted draws. Units are $100 \cdot \log$ levels (percentage points for Bank Rate). Restrictions are imposed on impact only.

reverse. This reproduces the qualitative account in the source paper's Figure 6 and matches the Bank's contemporaneous narrative of imported inflation.

Narrative reproduction. The estimated shock series reproduce the paper's episodes: the 2008 collapse in world demand, the contractionary world energy shocks of 2022, UK monetary loosening readings after 2022 followed by tightening readings in 2023, and muted shocks during 2020–21 by construction of the Covid dummies. Known caveats are equally documented: median UK GDP responses to world demand and supply shocks decay through zero at long horizons where the paper's remain positive (the paper's path lies inside the replication's bands), and within the global supply side the energy shock contributes more than the non-energy supply shock to world GDP and oil-price variance, where the paper's figure suggests the reverse ordering; the paper's claims about the *combined* supply side are matched.

7 Comparison with the original publication

This section compares the replication directly with what Brignone and Piffer (2025) publish, side by side and with explicit deviations. It is distinct from the outturn comparison of Section 8: Table 6 there judges the model's forecasts against subsequent ONS data, whereas everything here judges the replication against the *paper's own published numbers and figures*.

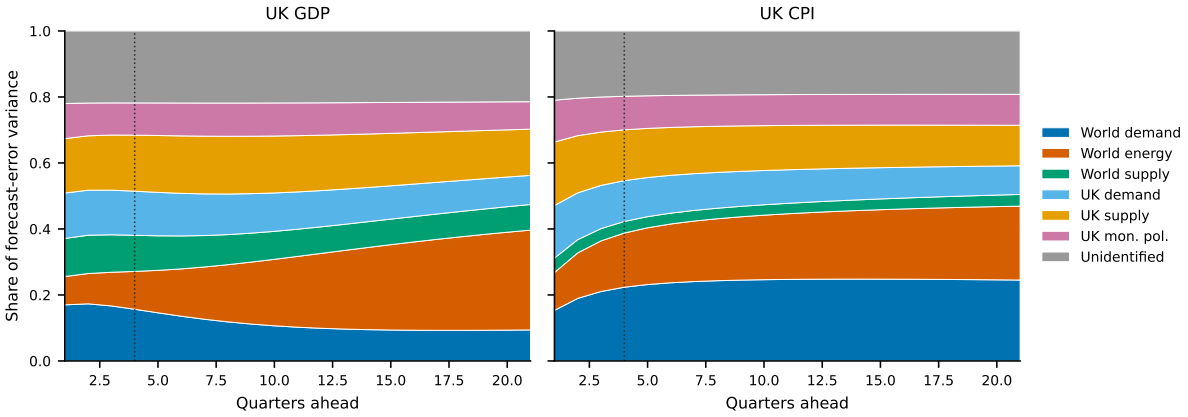


Figure 2: Forecast-error variance decompositions of UK GDP and UK CPI by horizon. Each band is the posterior *mean* across accepted draws of that shock's share of *total* forecast-error variance; shares sum to one across shocks on every draw, so nothing is renormalised and the two unidentified shocks (combined, grey) are the genuine residual rather than a plug. The horizon axis is quarters ahead — column h of equation (10) holds the $(h{+}1)$-step variance — and the dotted line marks the one-year reading at four quarters. An earlier version of this figure stacked pointwise per-shock medians renormalised to sum to one, which inflated every identified share by roughly a third.

7.1 Published anchors, side by side

The paper publishes two quantitative anchors that the replication's test suite locks against — the one-year global FEVD shares of UK GDP and UK CPI — plus a set of qualitative results (restriction patterns, the unrestricted oil-price sign check, the shock narrative, and the ranking of domestic CPI contributors). Table 4 collects all of them. Two replication configurations are shown: the CI fast configuration (600 posterior draws, fixed seeds, 51 accepted, unweighted — exactly the configuration the repository's validation tests run on every commit) and the production configuration (10,000 draws, 765 accepted, importance-weighted, ESS 368.0). Both are computed on the paper's own definition of the statistic, set out in Section 6: the posterior mean of the group share formed on each accepted draw, at the four-quarter-ahead forecast-error variance, as a share of total variance.

The headline pattern is a partial match, and less flattering than earlier versions of this table reported. On UK GDP the production run's 38.7 per cent against the paper's ∼40 is a match, and the total explained by the six identified shocks (77 and 78 per cent against "around 80") lines up on both variables. On UK CPI the replication finds 41.5 per cent against ∼50 — roughly eight and a half points short. Earlier versions of this paper reported 42.1 and 49.5 per cent and read the CPI figure as a match to within half a point; that agreement was produced by renormalising a sum of per-shock posterior medians to 100 per cent, an arithmetic step rather than an economic result (Section 6).

Two things stop the shortfall being over-read. First, the 68 per cent posterior band on the CPI share is $[25.6, 56.3]$, about ± 15 points — wider than the 8.5-point gap, so the data do not resolve it. Second, the difference between the fast and production configurations (40.1 versus 38.7 for GDP, 40.1 versus 41.5 for CPI) is Monte-Carlo noise in a 51-draw accepted sample, not a weighting effect: the global GDP share from the fast configuration has a standard deviation of

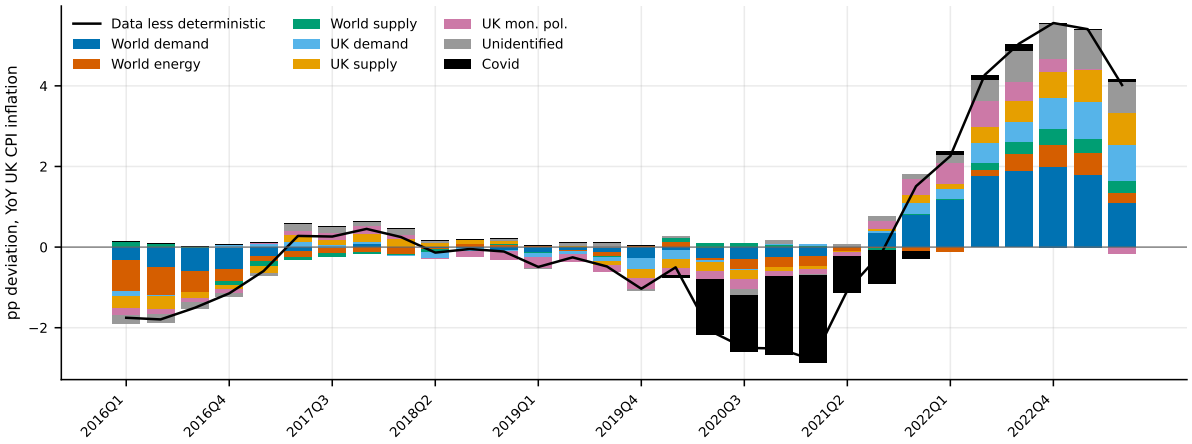


Figure 3: Historical decomposition of year-on-year UK CPI inflation (deviation from the deterministic component), 2016Q1–2023Q2. Bars: median cumulative contributions of the six identified shocks, the combined unidentified shocks, and the Covid-dummy component; line: data less the deterministic component.

about 6.5 points across sampling seeds, and the unweighted 3,000-draw run behind Figures 1–3 lands at 38.1 and 42.3 per cent, within a point of the weighted production values. An earlier version of this section claimed the importance weights moved both shares substantially towards the published values; on the corrected statistic they do not, and the weights' role is the narrower one of correcting the Arias et al. (2018) zero-restriction sampler towards the uniform-over-rotations posterior.

Figure 4 shows the same comparison graphically. No ours-versus-paper IRF overlay is drawn because Brignone and Piffer (2025) publish impulse responses only as figures; neither the paper nor the repository's validation data contains digitised numeric IRF paths, so the IRF comparison is confined to the sign and shape checks in Table 4 and Section 6.

7.2 The companion Staff Working Paper on structural forecast analysis

The natural question after Table 6 is whether the model's authors have published forecast-evaluation numbers of their own to compare against. The companion paper — Brignone and Piffer (2026), Bank of England Staff Working Paper No. 1,165, "Structural forecast analysis", January 2026 — turns out to answer a different question. It is a methods paper: it derives a five-way structural decomposition of forecast errors and forecast revisions (deterministic component, past shocks, revisions to past shocks, new shocks, and parameter/vintage revision) and illustrates it on a *four*-variable UK SVAR (Bank Rate, real GDP, CPI and real oil prices, 1992Q1–2025Q2, five lags, hierarchical Giannone et al., 2015 hyperpriors, Covid dummies for 2020), running real-time unconditional forecasts each quarter from 2021Q4 to 2025Q2 on the actual ONS vintages. It publishes no root-mean-squared errors, coverage rates or benchmark comparisons — its output is the structural attribution of the errors, shown in figures — so no numeric accuracy comparison with Table 6 is possible, and the RMSEs reported there (0.32 percentage points at horizons of one to seven quarters) stand without an official counterpart.

The comparison that *is* possible is of the attribution itself, and it is reassuring. Brignone and Piffer (2026) attribute the large upward inflation-forecast revisions of 2022 to inflationary

Table 4: Replication versus the published values of Brignone and Piffer (2025)

Quantity	Paper	Fast config	Production	Deviation
Global share, UK GDP FEVD, 1 yr	∼40%	40.1%	38.7%	+0.1 / −1.3 pp
Global share, UK CPI FEVD, 1 yr	∼50%	40.1%	41.5%	−9.9 / −8.5 pp
Identified total, UK GDP / CPI, 1 yr	"around 80%"	—	77% / 78%	−3 / −2 pp
Oil-price impact IRF, world supply shock	positive (Fig. 2)	+5.83	positive	sign matches
Zero restrictions (paper Table 2)	imposed	hold exactly ($< 10^{-9}$), every draw		cannot fail
Sign restrictions (paper Table 2)	imposed	hold on every accepted draw		cannot fail
Sign restrictions, median IRF	imposed per draw	hold		none
Shock narrative episodes (paper Fig. 5)	2008, 2022–23	reproduced		none
CPI inflation surge attribution (paper Fig. 6)	global energy/demand	reproduced		none
Largest domestic CPI contributor	monetary policy (prose)	UK supply (13.9% vs 8.7%)		unresolved
UK GDP IRF to world shocks, long horizons	persistently positive	median decays through zero		inside bands

Notes: "Fast config": 600 posterior draws, sample seed 1, identification seed 2, 51 accepted, unweighted — the repository's per-commit CI validation configuration, re-estimated fresh for this table by `figures/make_comparison.py`. "Production": 10,000 draws, 765 accepted, importance-weighted (ESS 368.0), from the repository's production artifact. FEVD rows are posterior means of the per-draw group share of *total* four-quarter-ahead forecast-error variance, the statistic the paper's Figure 4 plots; the production 68 per cent posterior bands are $[24.0, 53.1]$ for GDP and $[25.6, 56.3]$ for CPI, so the deviations shown are well inside a single band. Deviations are fast / production in percentage points against the paper's approximate published shares. The zero and sign restriction rows are marked "cannot fail": they are properties of the Arias et al. (2018) construction and of the accept/reject rule, not tests. The paper publishes IRFs only as figures — no numeric values — so IRF comparisons are of sign and shape. The domestic-contributor row is marked unresolved: the paper asserts the ranking in prose but publishes no FEVD table.

supply and energy shocks combined with expansionary demand and monetary policy shocks, plus revisions to previously estimated shocks. The replication's forecast-revision machinery (Section 3.5) is exactly the "new shocks" term of their five-way decomposition — the term their identity isolates when the parameters and past data are held fixed — and its 2024Q2 reading (a probable contractionary world energy shock alongside strong world demand, Section 8) is the same shock vocabulary applied two years later. Their broader decomposition, which also tracks vintage revisions and re-estimation, is a natural extension of the replication's fixed-parameter identity (13).

7.3 Cross-checks against other published UK shock estimates

A replication of a set-identified SVAR should also be checked *sideways*, against shock estimates published by other authors with different identification schemes. Table 5 performs the check qualitatively — sign and timing over the major overlapping episodes — against three benchmarks: the Bank's own Monetary Policy Report narrative (most recently the November 2025 Report's quantified account of the 2025 inflation overshoot), the oil-supply-news series of Känzig

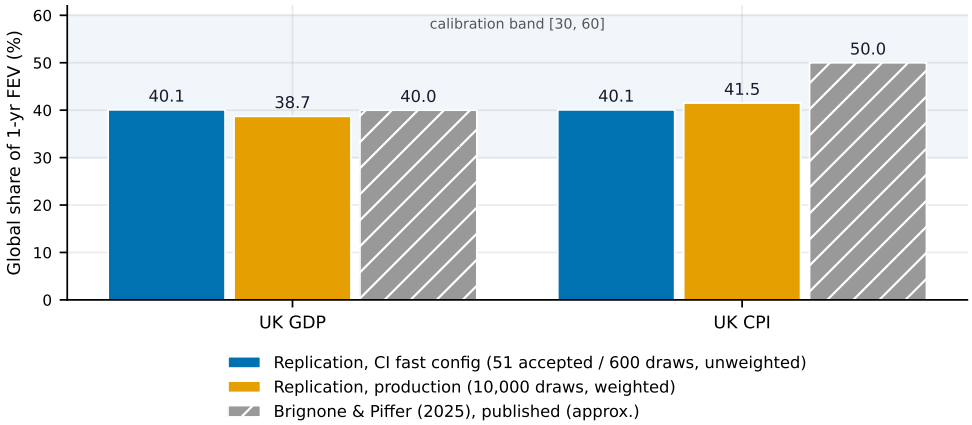


Figure 4: One-year global FEVD shares of UK GDP and UK CPI: replication versus Brignone and Piffer (2025). Shares are posterior means of the per-draw group share of total four-quarter-ahead forecast-error variance — the statistic the paper's Figure 4 plots — and are not renormalised. Blue: fresh fixed-seed run of the CI fast configuration (600 draws, 51 accepted, unweighted). Orange: production run (10,000 draws, 765 accepted, importance-weighted). Hatched grey: the paper's approximate published values. Shaded region: the [30, 60] per cent regression tripwire used by the validation suite, which is too wide to constitute a calibration result.

(2021), identified from futures-price movements in tight windows around OPEC announcements (extended vintage through 2022), and the narrative UK monetary policy shock series of Cloyne and Hürtgen (2016).

Three points deserve emphasis. First, where the schemes measure the same object — the 2008 world demand collapse, the 2022 energy shock — signs and timing agree. Second, where they measure *different* objects the disagreement is informative rather than troubling: the replication's world energy shock is defined by price–quantity comovement across all energy-price drivers, so it registers the 2022 Russian gas shock strongly, while the OPEC-announcement instrument of Känzig (2021) is deliberately silent on non-OPEC events. Third, the November 2025 MPR's arithmetic of the 2025 overshoot — roughly half administered prices and food, half labour costs — maps into the model's vocabulary as domestic cost-push (negative UK supply) layered on imported price shocks, which is precisely the shock combination the frozen 2024Q2 data edge could not yet see; it explains *which* shocks the undershoot in Table 6 missed, not just that some were missed.

7.4 The Bank's use of the model since publication

Two 2026 Bank publications confirm that the replicated model remains in operational use. The January 2026 *Forecast Evaluation Report* lists Brignone and Piffer (2025) among the "range of cross-check models … used regularly to help inform judgements and in some cases also interpret changes in the Bank's forecasts" (Bank of England, 2026a). And in July 2026, MPC member Catherine L. Mann cited the companion methodology directly: "Using a structural VAR model, Brignone and Piffer (2026) demonstrate that upward revisions in inflation forecasts following the post-pandemic inflation surge can be attributed to a combination of expansionary shocks on the

Table 5: Sign and timing of identified shocks across published UK estimates

Episode	This replication	Bank MPR narrative	Känzig (2021) oil-supply news
2008–09 crisis	Large negative world demand shocks	Global demand collapse	No large supply-news events; oil swings demand-driven — consistent
Covid, 2020	Shocks muted by construction (Covid dummies absorb 2020Q1–2021Q2)	Unprecedented demand and supply disruption	2020 oil crash read as demand, not supply news; April 2020 OPEC+ cut a positive supply-news event
2021–23 energy crisis	Contractionary world energy shocks dominate the CPI decomposition (Figure 3)	"Imported inflation": energy and tradables	Only modest shocks: the Russian gas shock fell outside OPEC announcement windows — complementary, not contradictory
2025 inflation hump	Energy/supply-side price shocks on weak demand (2024Q2 posterior reading, extended)	Nov 2025 MPR: 1.8 pp overshoot = 0.4 pp administered prices + 0.4 pp food + ∼1 pp labour costs	Outside the series' 2022 data edge

Notes: Qualitative comparison of sign and timing; the underlying series use different identification schemes (zero/sign restrictions here; narrative decomposition in the MPRs; high-frequency futures surprises around OPEC announcements in Känzig, 2021). The Cloyne and Hürtgen (2016) narrative UK monetary policy series is omitted from the episode columns because it ends in the 2000s (extended to early 2009, when Bank Rate reached its effective lower bound) and so overlaps none of the episodes; it and the high-frequency UK series of Cesa-Bianchi et al. (2020) (through 2015) remain the natural benchmarks for the replication's monetary policy shock over the pre-Covid sample.

demand side, revisions to past shocks, and contractionary supply-side shocks" (Mann, 2026) — the same attribution the replication's historical decomposition delivers for 2021–23 (Figure 3).

The Bank's own published projections also provide a late cross-check on Table 6. The February 2026 Monetary Policy Report put four-quarter CPI inflation at 3.0 per cent and four-quarter GDP growth at 0.8 per cent for 2026Q1 (Bank of England, 2026b) — forecasts made with seven more quarters of data than the replication's frozen 2024Q2 edge. The outturns were 3.1 and 0.9; the replication's eighteen-month-old medians were 3.2 and 0.9. For that quarter, at least, freezing the data edge in mid-2024 cost the model almost nothing relative to the official forecast made contemporaneously.

8 Results: a forecast from the 2024Q2 data edge

8.1 The forecast

With the model estimated through 2023Q2 and conditioned on data through 2024Q2 — the same vintage the Bank faced in its August 2024 round — the unconditional forecast puts year-on-year UK GDP growth medians between **1.3 and 2.0 per cent** and year-on-year CPI inflation medians between **2.3 and 2.8 per cent** over the forecast horizon in the importance-weighted production run (6,000 draws, 487 accepted, ESS 207.5), with credible bands that widen with

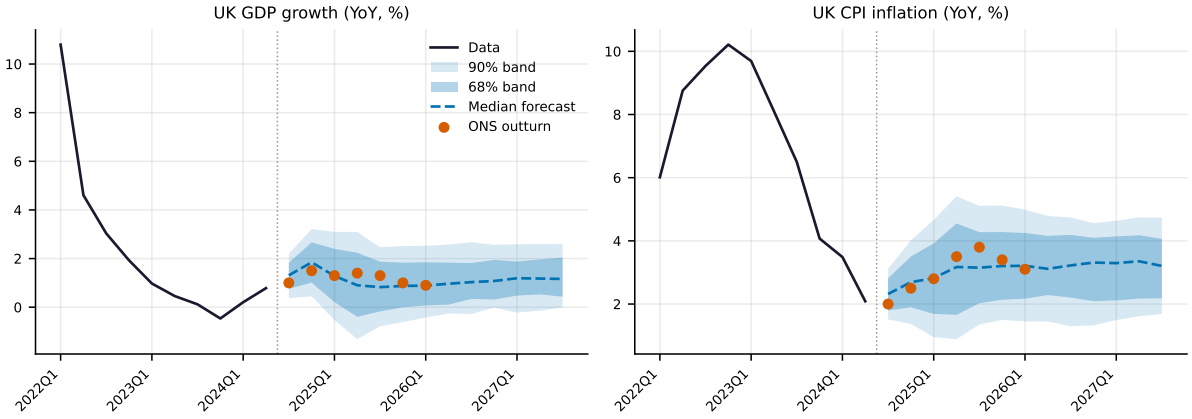


Figure 5: Fan-chart forecasts of year-on-year UK GDP growth and CPI inflation from the 2024Q2 data edge (median, 68 and 90 per cent predictive bands; parameter plus future-shock uncertainty, five stochastic paths per accepted draw), against subsequent ONS outturns (dots). The vertical dotted line marks the forecast origin.

horizon as posterior parameter and shock uncertainty compound. The structural reading at the data edge (the `latest_shocks` output) assigns a 0.81 posterior probability to a positive world demand shock and a 0.80 probability to a contractionary (price-raising) world energy shock in 2024Q2, with a 0.75 probability of a positive UK supply shock and a 0.66 probability of a negative world supply shock — a configuration in which global demand supports activity while the energy and domestic supply readings pull inflation in opposite directions. All numbers in this section use the frozen paper-vintage samples (estimation through 2023Q2, conditioning through 2024Q2); they are distinct from the extended 1992Q1–2025Q1 production run, re-estimated in July 2026 and re-tabulated in August 2026, which supplies the headline FEVD shares (37.4 and 42.3 per cent) quoted in the abstract and Section 6.

Figure 5 shows the corresponding fan charts from this paper's fixed-seed figure run (3,000 draws, 217 accepted, unweighted; per-quarter numbers in Table 6), with the ONS outturns that subsequently arrived overlaid as dots. The unweighted figure run is slightly less benign on inflation than the weighted production run — its CPI medians drift up to about 3.2 per cent at longer horizons rather than staying in the 2.3–2.8 range — a useful reminder that the importance weights and Monte-Carlo seed move medians by a few tenths even as the bands are robust.

8.2 Comparison with the outturn: an out-of-sample look, quarter by quarter

Table 6 confronts the forecast with what actually happened, quarter by quarter, through 2026Q1 — seven quarters of genuine out-of-sample data, the longest evaluation window ONS publication allows at the time of writing (the first estimate of 2026Q1 GDP appeared on 14 May 2026; June 2026 CPI has since landed, putting the 2026Q2 quarterly-average CPI inflation outturn at 2.8 per cent, inside the model's 68 per cent band — the archived forecast round's error was −0.12 percentage points, against −0.66 for the OBR March 2026 EFO and 0.30 for a naive random walk, per the public track record at `https://policyengine-macro.vercel.app/forecasts`).

The scorecard is mixed in an informative way. On *activity*, the model does well throughout:

Table 6: Quarter-by-quarter forecast versus outturn (year-on-year, per cent)

	GDP growth			CPI inflation		
Quarter	Median	68% band	Outturn	Median	68% band	Outturn
2024Q3	1.3	[0.8, 1.8]	1.0	2.3	[1.8, 2.8]	2.0
2024Q4	1.8	[1.0, 2.7]	1.5	2.7	[1.9, 3.5]	2.5
2025Q1	1.3	[0.2, 2.4]	1.3	2.8	[1.7, 3.9]	2.8
2025Q2	0.9	[−0.4, 2.2]	1.4	3.2	[1.7, 4.6]	3.5
2025Q3	0.8	[−0.2, 1.9]	1.3	3.2	[2.0, 4.3]	3.8
2025Q4	0.9	[−0.0, 1.8]	1.0	3.2	[2.1, 4.3]	3.4
2026Q1	0.9	[0.1, 1.8]	0.9	3.2	[2.2, 4.3]	3.1

Notes: Model columns are from the fixed-seed figure run of Figure 5 (3,000 draws, 217 accepted, unweighted). GDP outturns are ONS estimates of quarterly real GDP relative to the same quarter a year earlier (Office for National Statistics, 2026c,b); annual growth was 1.1 per cent in 2024 and 1.3 per cent in 2025. CPI outturns are quarterly averages of twelve-month CPI inflation (Office for National Statistics, 2025b, 2026a): 2025Q4 averages October–December 2025 rates of 3.6, 3.2 and 3.4 per cent; 2026Q1 averages January–March 2026 rates of 3.0, 3.0 and 3.3 per cent; the monthly peak was 3.8 per cent in September 2025. Later ONS vintages revise these figures — the June 2026 Quarterly National Accounts already revised quarterly growth in 2025Q3 and 2025Q4 down to 0.1 per cent each.

every quarterly GDP outturn lies inside the 68 per cent band, and the medians track the economy's actual gentle path — modest growth around 1–1.5 per cent, neither re-acceleration nor stall. The one visible miss in profile is 2025Q2–Q3, where the model's median dips towards 0.8–0.9 per cent while the economy held at 1.3–1.4 per cent; the bands absorb it. On *inflation*, the model captures the level and timing of the 2024Q3–2025Q1 path almost exactly (2.3 vs 2.0, 2.7 vs 2.5, 2.8 vs 2.8), then progressively undershoots the 2025 re-acceleration: outturns of 3.5 and 3.8 per cent against medians of 3.2, still inside the 68 per cent band but on its upper half. The two newest quarters close the loop: as the 2025 hump unwound — quarterly-average inflation fell back to 3.4 per cent in 2025Q4 and 3.1 per cent in 2026Q1 — the outturn converged onto the model's essentially flat 3.2 per cent median, and the 2026Q1 GDP outturn of 0.9 per cent landed on the model's 0.9 median exactly. Seven quarters ahead, the frozen-2024Q2-edge forecast errors are small: root-mean-squared errors of 0.32 percentage points for both GDP growth and CPI inflation over 2024Q3–2026Q1, with all fourteen outturns inside the 68 per cent bands. The widening credible bands do their job — no outturn escapes even the 68 per cent band in this run. That result should, however, be read the other way around: a correctly calibrated 68 per cent band should contain roughly 9.5 of 14 outturns, so containing all 14 suggests the bands are over-dispersed at this particular origin — and the 49-origin rolling coverage evaluation (Section 6) finds the opposite pattern, systematic *under*-coverage, so the frozen-origin result should not be read as evidence that the bands are calibrated. But the medians clearly did not anticipate how far energy, food and administered prices would push inflation in mid-2025; the model was right that the hump would mean-revert, wrong (by up to 0.6 points) about its peak.

8.3 What the model says 2025 inflation was, and what the Bank said

The model's structural machinery lets the miss be interrogated rather than merely recorded. Two readings are relevant.

First, the *data-edge reading*. At 2024Q2 the `latest_shocks` posterior already flagged a contractionary world energy shock (probability 0.80) alongside strong world demand (0.81): in

the model's language, the last observed quarter was one in which energy-driven cost pressure was hitting while global demand held up. The forecast-revision decomposition quantifies what that reading did to the outlook: relative to a pseudo-forecast made one quarter earlier, the 2024Q2 shocks revised the year-on-year GDP path up by a peak of +0.40 percentage points and the CPI path down by −0.22 percentage points on impact — the structural account of why the mid-2024 data changed the forecast, exactly the round-by-round use the Bank built the model for. What the frozen data edge could not see is that the energy- and price-level shocks of late 2024 and 2025 — Ofgem price-cap increases, food-price inflation, April 2025 administered and regulated price rises — would keep arriving.

Second, the *ex-post comparison with the Bank's own narrative*. The Bank's August 2025 Monetary Policy Report, written as inflation averaged 3.5 per cent in 2025Q2, attributed the increase from 2.8 per cent in Q1 explicitly to energy, food and administered prices, and worried chiefly that the salience of food and energy could feed inflation expectations and the wage–price process (Bank of England, 2025a); the MPC nevertheless cut Bank Rate to 4 per cent, judging the pressure to be a hump rather than a trend. By November 2025 the Report judged that CPI inflation had peaked at 3.8 per cent in the July–September period and projected a fall to around 3 per cent in coming quarters and a gradual decline to target thereafter, underpinned by subdued growth, building labour-market slack, easing pay growth and easing services inflation (Bank of England, 2025b). Mapped into the replication's shock vocabulary, the Bank's story of 2025 inflation is a sequence of adverse *energy and supply-side price* disturbances layered on weak demand — precisely the world energy and (negative) supply shocks the model's 2024Q2 posterior was already detecting, extended for three more quarters. The model's historical decomposition of the 2021–23 surge (Figure 3) tells the same story at larger scale: in both episodes the model and the Bank agree that UK inflation moved on imported cost shocks rather than domestic overheating. The residual disagreement is one of magnitude and persistence, not mechanism: the SVAR's mean-reverting shock structure pulls inflation back towards the deterministic trend faster than the administered-price pipeline of 2025 allowed, an error the Bank's August 2024 projections — which had inflation falling back towards 2 per cent over 2025 (Bank of England, 2024) — shared in full.

In annual terms: ONS outturns of 1.1 per cent growth in 2024 and 1.3 per cent in 2025 sit at or just below the bottom of the model's median range and comfortably inside the bands, while CPI inflation's climb to a 3.8 per cent September 2025 peak escapes the weighted-run median range of 2.3–2.8 per cent but not the credible bands. The miss is shared with the official forecaster and, through the shock decomposition, structurally interpretable.

9 Limitations and conclusion

Four limitations should frame any use of the model, and each is disclosed in the codebase as well as here.

The data edge is frozen at 2024Q2. The estimation sample ends in 2023Q2 and the conditioning data in 2024Q2, matching the source paper's vintage. Forecasts and shock readings therefore describe the economy as of mid-2024; the out-of-sample comparison in Section 8 is possible precisely because the edge is frozen, but operational use requires re-running the data pipeline on current vintages, and results will shift with revisions. *Update (July 2026):* the production configuration has since been re-estimated on the extended 1992Q1–2025Q1 sample

with data through 2026Q2, and the headline FEVD shares quoted in the abstract and Section 6 (38.7 and 41.5 per cent) come from that extended run; the frozen-2024Q2 out-of-sample exercise and the original-sample replication figures are retained unchanged.

Internal Bank series are approximated. The UK-trade-weighted world GDP and world CPI series are unpublished; the replication's OECD/IMF-based proxies, the assumed lag length, the default prior hyperparameters and the simplified pandemic-prior treatment mean the global block matches the paper qualitatively, not exactly. Two substantive discrepancies stand: the global share of one-year UK CPI variance is about eight points below the paper's approximate 50 per cent, and UK monetary policy is not the largest domestic contributor to CPI variance. Both are plausibly attributable to these approximations, but neither has been traced to a cause, and the posterior band on the first is wider than the gap; they are recorded as unresolved rather than explained.

Forecast skill is not established. The frozen-origin exercise below is one forecast origin. The 49-origin rolling evaluation (Section 6) finds no accuracy gain over a random walk with drift on any of the eight variables at any of the eight horizons once the 64 tests are adjusted together, and finds the predictive bands under-covering. Nothing in this paper supports a claim that the model forecasts better than a naive benchmark.

This is not an official Bank product. The replication is an independent open-source reconstruction from the published paper. It is not endorsed by, and does not speak for, the Bank of England; where the replication and the paper disagree, the paper is ground truth. Nor does the model score reforms: within PolicyEngine Macro it is the baseline and conditioning member, and reform scoring belongs to the suite's other models.

Extensions. Three extensions follow naturally from the machinery already in place. Conditional forecasting in the sense of Waggoner and Zha (1999) — conditioning the fan charts on a market-implied Bank Rate path or an energy-price assumption — requires only the linear-restriction algebra of Section 3.5 plus an economic choice of which shocks implement the conditions. Narrative sign restrictions (Antolín-Díaz and Rubio-Ramírez, 2018) on the clearly signed 2008 and 2022 episodes in the estimated shock series would sharpen the identified set. And re-running the data pipeline on current vintages would move the data edge forward, turning the out-of-sample exercise of Section 8 into a live forecast evaluation loop.

Within those limits, the exercise shows that a central bank's operational structural VAR can be rebuilt in the open: the identification implemented and verified draw by draw, the published magnitude for GDP reproduced and the shortfall on CPI quantified with its posterior band rather than renormalised away, out-of-sample performance on the record including where it is unflattering, and the whole apparatus callable through hosted tools. Replication of policy-institution models, with numbers attached, is both feasible and useful — and this paper is offered as a template for doing it.

References

Antolín-Díaz, J. and Rubio-Ramírez, J. F. (2018). Narrative sign restrictions for SVARs. *American Economic Review*, 108(10):2802–2829.

Arias, J. E., Rubio-Ramírez, J. F., and Waggoner, D. F. (2018). Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications. *Econometrica*, 86(2):685–720.

Bańbura, M., Giannone, D., and Reichlin, L. (2010). Large Bayesian vector auto regressions. *Journal of Applied Econometrics*, 25(1):71–92.

Bank of England (2024). Monetary policy report – august 2024. https://www.bankofengland.co.uk/monetary-policy-report/2024/august-2024.

Bank of England (2025a). Monetary policy report – august 2025. https://www.bankofengland.co.uk/monetary-policy-report/2025/august-2025.

Bank of England (2025b). Monetary policy report – november 2025. https://www.bankofengland.co.uk/monetary-policy-report/2025/november-2025.

Bank of England (2026a). Forecast evaluation report – january 2026. https://www.bankofengland.co.uk/paper/2026/forecast-evaluation-report-january-2026.

Bank of England (2026b). Monetary policy report – february 2026. https://www.bankofengland.co.uk/monetary-policy-report/2026/february-2026.

Baumeister, C. and Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5):1963–1999.

Brignone, D. and Piffer, M. (2025). A structural VAR model for the UK economy. Macro Technical Paper 3, Bank of England.

Brignone, D. and Piffer, M. (2026). Structural forecast analysis. Staff Working Paper 1165, Bank of England.

Burgess, S., Fernandez-Corugedo, E., Groth, C., Harrison, R., Monti, F., Theodoridis, K., and Waldron, M. (2013). The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models. Working Paper 471, Bank of England.

Canova, F. and De Nicoló, G. (2002). Monetary disturbances matter for business fluctuations in the G-7. *Journal of Monetary Economics*, 49(6):1131–1159.

Cascaldi-Garcia, D. (2022). Pandemic priors. International Finance Discussion Papers 1352, Board of Governors of the Federal Reserve System.

Cesa-Bianchi, A., Thwaites, G., and Vicondoa, A. (2020). Monetary policy transmission in the United Kingdom: A high frequency identification approach. *European Economic Review*, 123:103375.

Chan, J. C. C., Matthes, C., and Yu, X. (2025). Large structural VARs with multiple sign and ranking restrictions. *Journal of Econometrics*.

Cloyne, J. and Hürtgen, P. (2016). The macroeconomic effects of monetary policy: A new measure for the United Kingdom. *American Economic Journal: Macroeconomics*, 8(4):75–102.

Doan, T., Litterman, R. B., and Sims, C. A. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews*, 3(1):1–100.

Faust, J. (1998). The robustness of identified VAR conclusions about money. *Carnegie-Rochester Conference Series on Public Policy*, 49:207–244.

Fry, R. and Pagan, A. (2011). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4):938–960.

Ghenis, M. and PolicyEngine contributors (2025). PolicyEngine: Open-source tax–benefit microsimulation and macroeconomic models. https://policyengine.org.

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior selection for vector autoregressions. *Review of Economics and Statistics*, 97(2):436–451.

Kadiyala, K. R. and Karlsson, S. (1997). Numerical methods for estimation and inference in Bayesian VAR-models. *Journal of Applied Econometrics*, 12(2):99–132.

Känzig, D. R. (2021). The macroeconomic effects of oil supply news: Evidence from OPEC announcements. *American Economic Review*, 111(4):1092–1125.

Kilian, L. and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis*. Cambridge University Press, Cambridge.

Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions — five years of experience. *Journal of Business & Economic Statistics*, 4(1):25–38.

Mann, C. L. (2026). Mixed signals, research findings, and policy judgements. Speech at the Natixis CIB conference, Bank of England, July 2026. https://www.bankofengland.co.uk/speech/2026/july/catherine-l-mann-fireside-chat-at-the-natixis-cib-conference-on-private-debt.

Office for National Statistics (2025a). Consumer price inflation, UK: December 2024. https://www.ons.gov.uk/economy/inflationandpriceindices/bulletins/consumerpriceinflation/december2024.

Office for National Statistics (2025b). Consumer price inflation, UK: November 2025. https://www.ons.gov.uk/economy/inflationandpriceindices/bulletins/consumerpriceinflation/november2025.

Office for National Statistics (2026a). Consumer price inflation, UK: May 2026. https://www.ons.gov.uk/economy/inflationandpriceindices/bulletins/consumerpriceinflation/may2026.

Office for National Statistics (2026b). GDP first quarterly estimate, UK: January to march 2026. https://www.ons.gov.uk/economy/grossdomesticproductgdp/bulletins/gdpfirstquarterlyestimateuk/januarytomarch2026.

Office for National Statistics (2026c). GDP first quarterly estimate, UK: October to december 2025. https://www.ons.gov.uk/economy/grossdomesticproductgdp/bulletins/gdpfirstquarterlyestimateuk/octobertodecember2025.

Rubio-Ramírez, J. F., Waggoner, D. F., and Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *Review of Economic Studies*, 77(2):665–696.

Sims, C. A. (1980). Macroeconomics and reality. *Econometrica*, 48(1):1–48.

Sims, C. A. and Zha, T. (1998). Bayesian methods for dynamic multivariate models. *International Economic Review*, 39(4):949–968.

Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2):381–419.

Waggoner, D. F. and Zha, T. (1999). Conditional forecasts in dynamic multivariate models. *Review of Economics and Statistics*, 81(4):639–651.

A Data appendix

Table 7 documents every series in the dataset: source, series codes, and the construction applied. The dataset described here is the paper-vintage dataset, covering 1992Q1–2024Q2 (130 quarters, no missing values); the current data pipeline extends the same construction through 2026Q2. The binding start is the FRED Brent series `POILBREUSDQ`, which begins in 1992Q1. All variables enter the model as $100 \cdot \log$ of the level except Bank Rate, which enters as the level in per cent. The full download-and-construction script ships with the repository (`scripts/download_data.py`); raw source files are cached so the build is reproducible offline.

B Prior hyperparameters

Table 8 lists every prior hyperparameter, its default, and its role. Defaults follow the Giannone et al. (2015) conventions; the source paper does not report the Bank's settings. Maximising the closed-form marginal likelihood (6) over (λ, μ, θ) on the estimation sample — the optional empirical-Bayes step — selects substantially looser lag shrinkage than the default ($\lambda \approx 1.12$, $\mu \approx 0.48$, $\theta \approx 1.0$; log ML -1678.5 against -1795.4 at the defaults, computed on the estimation sample with the Covid dummies); the shipped configuration retains the conventional defaults for comparability with standard practice. That the $[30, 60]$ per cent tripwire of Section 6 covers both settings is not reassurance — the band is too wide to distinguish them — and the sensitivity of the headline FEVD shares to the hyperparameters has not been measured.

C Sampling and convergence diagnostics

Posterior draws are i.i.d. (direct sampling from the conjugate NIW posterior; no Markov chain), so the relevant diagnostics concern the identification step: the sign-restriction acceptance rate and the effective sample size of the Arias et al. (2018) importance weights. Table 9 reports them for every run used in this paper. Acceptance rates of 7–8 per cent are typical for a restriction set of this size (23 sign restrictions across six shocks plus 13 zero restrictions), and are materially raised by the Chan et al. (2025) permutation search relative to plain accept/reject. The weighted ESS of roughly 43–48 per cent of accepted draws (47.9 per cent in the production run, 42.6 per

cent in the forecast-revision run) indicates moderately variable weights — informative but not degenerate; no single draw dominates ($\max_i w_i / \sum_i w_i < 2$ per cent in the production run). Weighted and unweighted headline statistics are compared in Section 6. The weights are there to recover the uniform-over-rotations posterior, but on this dataset they do not visibly move the headline shares: the unweighted 3,000-draw run gives one-year global FEVD shares of 38.1 (UK GDP) and 42.3 (UK CPI) per cent against the weighted production run's 38.7 and 41.5. The much larger spread of the 600-draw CI configuration (40.1 and 40.1) is Monte-Carlo noise from 51 accepted draws, not a weighting effect.

Table 7: Series, sources and transformations

Variable	Description	Source (codes)	Construction
`uk_gdp`	UK real GDP, chained volume, SA, £m	ONS UKEA (`ABMI`)	As published.
`cpisa`	UK CPI, seasonally adjusted	ONS MM23 (`D7BT`)	Monthly NSA index averaged to quarterly, then multiplicative classical seasonal adjustment (ratio to 2×4 centred moving average, constant quarterly factors normalised to mean 1). Approximates the Bank's internal SA series; no X-13.
`cpi_energy`	UK CPI energy	ONS MM23 (`D7CH`, `D7EC`)	CPI division 04.5 (electricity, gas & other fuels) and class 07.2.2 (fuels & lubricants) combined with equal 0.5/0.5 weights via chained weighted log growth; NSA. The official aggregate uses time-varying expenditure weights.
`bank_rate`	Bank Rate, quarterly average, %	BoE IADB (`IUQABEDR`)	As published.
`eri`	Sterling broad effective exchange-rate index	BoE IADB (`XUQABK67`)	Quarterly average.
`oil_price`	Real Brent price in sterling	FRED (`POILBREUSDQ`); BoE (`XUQAUSS`)	Brent USD/bbl (quarterly average) divided by USD/GBP, deflated by `cpisa`. Index-like level; only $100 \cdot \log$ enters the VAR.
`world_gdp`	UK-trade-weighted world real GDP (proxy)	FRED: US `GDPC1`; EA19 `CLVMNACSCAB1GQEA19`; DE `CLVMNACSCAB1GQDE`; JP `JPNRGDPEXP`; CN `RGDPNACNA666NRUG`	Divisia-style chain of weight-averaged quarterly log growth with era trade weights (EA/US/JP/CN: 0.50/0.17/0.04/0.01 to 1999; 0.48/0.16/0.03/0.05 to 2009; 0.45/0.16/0.02/0.07 after), renormalised each quarter over available countries. EA backcast pre-1995 with German growth; China (annual PWT) log-linearly interpolated, drops out after 2023Q2; Japan enters from 1994Q1.
`world_cpi`	UK-trade-weighted world CPI (proxy)	FRED: US `CPALTT01USQ661S`; EA19 `CP0000EZ19M086NEST`; DE `DEUCPIALLMINMEI`; JP `CPALTT01JPQ661S` + `FPCPITOTLZGJPN`; CN `CHNCPIALLMINMEI`	Same era weights and chain method. EA backcast pre-1997 with German CPI; Japan extended after 2021Q2 with World Bank annual inflation (flat within-year profile); China enters from 1993Q2. EA/DE legs NSA.

Notes: The two world aggregates proxy unpublished internal Bank series. The era trade weights are static approximations informed by ONS Pink Book and IMF Direction of Trade shares (covering roughly 0.70–0.72 of UK trade; the remainder implicitly grows at the weighted average of the included countries), documented in the repository's `data/README.md` together with a comparison file retaining the earlier two-country (US/EA) proxies, with which the current aggregates correlate at 0.996–1.000 in levels.

Table 8: Prior hyperparameters

Parameter	Default	Role
λ	0.2	Overall Minnesota tightness: prior s.d. of the coefficient on variable j, lag l, equation i is $\lambda s_i/(l^2 s_j)$.
lag decay	l^2	Harmonic-quadratic tightening in the lag order.
δ_i	1	Prior mean on each variable's own first lag (random walk in levels).
μ	1	Sum-of-coefficients (no-cointegration) tightness, centred on pre-sample means.
θ	1	Single-unit-root (dummy-initial-observation) tightness.
s_i	AR(1) resid. s.d.	Per-equation scale from univariate AR(1) regressions.
IW scale	$\mathrm{diag}(s_i^2)$	Prior scale of Σ via dummy rows.
$\varepsilon_{\text{const}}$	10^{-3}	Dummy weight on constant and Covid dummies (near-flat prior).
p	4	Lag length (unstated in the source paper; configurable).
Covid dummies	6	One exogenous 0/1 dummy per quarter, 2020Q1–2021Q2.

Table 9: Runs, acceptance and effective sample sizes

Run	Posterior draws	Accepted	Acceptance	Weights	ESS
Production replication	10,000	765	7.6%	yes	368.0
Forecast revision (2024Q2 edge)	6,000	487	8.1%	yes	207.5
Paper figures, estimation sample	3,000	241	8.0%	no	—
Paper figures, 2024Q2 edge	3,000	217	7.2%	no	—

Notes: All runs use $p = 4$, $\lambda = 0.2$, $\mu = 1$, $\theta = 1$ and fixed seeds (figure runs: seed 20250717 and 20250718 respectively; scripts in the paper's `figures/` directory). ESS is the Kish effective sample size of the normalised importance weights. Each weight requires two numerical Jacobians of a $k(k+m) = 376$-dimensional map, which is why the exploratory figure runs omit them.