



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Maximilian Schröder Multivariate uncertainty and
distributional transmission

No 3290 / September 2026

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

Macroeconomic policy increasingly requires assessments of risk rather than point forecasts alone. This paper develops a framework to compute impulse responses for scenario probabilities and the full predictive distribution of macroeconomic outcomes. The approach combines a regime-dependent mixture VAR with endogenous state probabilities and heteroskedastic shock identification, complemented by sparse sign and narrative restrictions. Applied to the euro area and U.S., the framework addresses a central policy question following the outbreak of the 2026 Iran war: do geopolitical shocks imply a return to the risk profile observed after Russia's invasion of Ukraine, or is this time different? Results show that shocks differ not only in their effects on expected outcomes, but also in how they reshape scenario probabilities, tail risks, and transitions into high-volatility states.

Keywords: Scenario Analysis; Density Forecasting; Macroeconomic Risks; State-Dependent Transmission; Bayesian Econometrics.

JEL Classification: C11, C32, C34, C53, E37

Non-technical summary

Macroeconomic shocks affect not only the most likely path of inflation and activity, but also the probability of adverse outcomes such as recession or stagflation. This matters for central banks, which increasingly need to assess the balance of risks around their baseline forecasts. When a new shock occurs, policymakers and market participants often look to similar historical episodes for guidance. Yet such comparisons can be misleading if they focus only on the shock itself, as its consequences may also depend on the state of the economy in which it occurs. A geopolitical disturbance hitting an economy already facing elevated inflation, weak activity or financial stress may generate very different risks from an equally large shock occurring under more benign conditions.

This paper develops a model designed to separate these two forces. It allows the transmission of economic shocks to vary with prevailing economic conditions and measures how shocks change not only expected inflation and activity, but also the probability of adverse outcomes. A distinctive feature of the framework is that scenarios are generated within the model rather than imposed externally. This allows identified shocks to alter both the paths of macroeconomic variables and the probabilities attached to different future states of the economy directly. This makes it possible to assess, in a unified way, how a given shock changes the full distribution of possible outcomes and the likelihood of policy-relevant scenarios such as recession or stagflation.

The paper illustrates the approach for monthly euro area and US data by taking the geopolitical shock observed during the build-up to Russia's invasion of Ukraine in January 2022 and applying the same shock to two different starting points: January 2022 and December 2025, ahead of the outbreak of the US-Iran war. This provides a direct answer to the question of whether a 2022-sized geopolitical shock would recreate a 2022-type risk environment; a key question faced by policy makers at the time.

The results suggest a more contained risk profile under the economic conditions prevailing at the end of 2025. In particular, the same geopolitical shock is associated with a smaller increase in stagflation risk than when applied to the January 2022 environment. This indicates that historical episodes can provide useful benchmarks, but their implications depend importantly on the state of the economy in which a new shock occurs.

1. INTRODUCTION & MOTIVATION

In recent years, geopolitical tensions, trade fragmentation, and large movements in commodity prices have made the macroeconomic environment markedly more uncertain, shifting not only expected outcomes but also the balance of risks. This shift is particularly challenging when policy involves competing objectives, as recent experience illustrates. The disruptions associated with the Covid-19 pandemic and Russia's invasion of Ukraine simultaneously raised concerns about slowing growth and persistent inflation, transforming stagflation from a tail risk into a central policy concern. Because policy decisions based solely on expected outcomes may overlook a buildup of tail risks, policymaking increasingly relies on risk assessments and scenario analysis rather than point forecasts alone (Bernanke 2024, Kamps et al. 2025, Federal Open Market Committee 2025).

However, existing frameworks typically treat scenarios as either a future path or an externally specified narrative. Conditional forecasting approaches based on identified models allow particular transmission channels to be isolated and traced through the economy (Waggoner and Zha 1999), while more recent scenario synthesis frameworks evaluate alternative narratives against a statistical reference distribution and quantify their relative plausibility (Adrian et al. 2025). Although these approaches provide valuable tools for policy analysis and risk assessment, they take the path or scenario itself as the primitive object of analysis. As a result, they are less naturally suited to studying how the risks and transmission associated with a given shock vary across economic environments.

This distinction becomes particularly important when policymakers seek to assess whether the risks associated with a new shock resemble those observed in past episodes. The 2026 Iran war provides a natural example. Although the conflict revived concerns reminiscent of those following Russia's invasion of Ukraine, inflation had largely moderated, supply disruptions had eased, and monetary policy had adjusted to the post-pandemic inflation surge. The central question is therefore not simply what a repeat of the 2022 geopolitical episode would imply for inflation, growth, and financial conditions, but whether a shock of comparable magnitude would generate a similar configuration of risks under current economic conditions. Answering this question requires tracing the impact of an identified shock through the evolving state of the economy rather than conditioning on a pre-specified path or externally defined scenario.

To this end, this paper argues that answering such questions requires characterizing scenarios

jointly by their paths and their probabilities. For this reason, the framework developed in this paper links identified economic shocks to changes in the full predictive distribution of macroeconomic outcomes, making it possible to quantify not only how inflation, output, and financial conditions evolve following a shock, but also how the shock reshapes the probabilities of economically relevant tail outcomes and, with them, the broader balance of risks faced by policymakers.

The framework builds on the observation of [López-Salido and Loria \(2024\)](#) that distributional asymmetries and tail risks can be interpreted through the lens of state-dependent macroeconomic dynamics, providing an alternative perspective on standard at-risk models. This paper makes that connection explicit by linking state-dependent dynamics, predictive distributions, and identified structural shocks. Specifically, the proposed framework combines a finite regime-dependent mixture VAR, related to the regime-switching tradition of [Hamilton \(1989\)](#), with endogenous mixing probabilities that evolve as functions of observable economic conditions following [Schröder \(2026\)](#) and [Rigon and Durante \(2021\)](#). This feature provides a flexible representation of macroeconomic risk, as changes in the likelihood of adverse economic states emerge endogenously from the data. Importantly, by allowing shock variances to differ across states, the regime structure creates a natural source of identification, building on the heteroskedasticity-based approaches of ([Rigobon 2003](#), [Lanne and Lütkepohl 2008](#), [Herwartz and Lütkepohl 2014](#), [Lanne et al. 2010](#)). Combined with sparse sign and narrative restrictions, this yields economically interpretable shocks that affect both macroeconomic outcomes and the likelihood of future economic scenarios, thereby reshaping the full predictive distribution of the economy rather than only its conditional mean. As a result, the framework delivers model-consistent impulse responses not only for expected outcomes, but also for tail probabilities, scenario probabilities, and the broader balance of risks.

The paper unifies three strands of the literature that have largely evolved separately. The first studies macroeconomic outcomes “at risk” and emphasizes how shocks and economic conditions affect the tails of predictive distributions ([Adrian et al. 2019](#), [Banerjee et al. 2024](#), [Bekaert and Popov 2019](#), [Monache et al. 2024](#), [Korobilis 2017](#), [Korobilis et al. 2021](#), [Korobilis and Schröder 2025](#), [López-Salido and Loria 2024](#), [Schröder 2026](#), [Tagliabracci 2020](#)). While this literature has increasingly moved toward multivariate settings ([Forni et al. 2021](#), [Iacopini et al. 2022](#), [Chavleishvili and Manganelli 2024](#), [Schüler 2020](#), [Korobilis and Schröder 2025](#)), existing approaches typically face one of two challenges. They either characterize dependence through

comovements of marginal quantiles rather than a coherent joint distribution, or require imposing identifying restrictions across different parts of the distribution. The proposed framework provides a structural alternative by identifying shocks at the model level and allowing their effects to propagate through the full predictive distribution. The theoretical results in [Online Appendix B](#) and [Online Appendix C](#) characterize the mechanisms underlying these distributional responses. In particular, they decompose state dependent propagation into within-regime and regime-selection channels, show how structural shocks alter future regime incidence through the evolving distribution of economic conditions, establish why otherwise identical shocks can transmit differently from different initial states, and characterize the resulting changes in scenario probabilities.

The second strand develops tools for scenario analysis, reflecting the increasing role of risk assessments in policy institutions. Recent contributions provide a coherent way to translate baseline forecasts into scenario-consistent outcome distributions ([Adrian et al. 2025](#)), offering a powerful framework for incorporating expert judgment and narrative scenarios into probabilistic forecasts. However, because the scenarios themselves remain external to the underlying economic model, they are largely agnostic about the structural shocks and transmission mechanisms that generate them. Consequently, they do not naturally characterize how the transmission of a given shock varies across different economic environments. The framework developed in this paper complements this literature by embedding scenario analysis within a structural, state-dependent model. By bringing a shock-based perspective to probabilistic scenario analysis, both scenario paths and scenario probabilities arise endogenously from identified shocks and their transmission through the economy.

The third strand comprises structural VARs and regime-switching models, which provide flexible tools for modelling state dependence and identifying economically meaningful shocks in nonlinear macroeconomic systems ([Hamilton 1989](#)). In particular, heteroskedasticity-based approaches exploit differences in shock variances across regimes to recover structural shocks without relying on strong contemporaneous restrictions ([Rigobon 2003](#), [Lanne and Lütkepohl 2008](#), [Herwartz and Lütkepohl 2014](#), [Lanne et al. 2010](#)). This paper combines these ideas by embedding identified shocks within a framework of endogenous state probabilities. As a result, shocks affect not only macroeconomic outcomes but also the likelihood of future economic states, providing a natural link between structural shocks, nonlinear dynamics, and the evolution of macroeconomic risks.

The framework is illustrated using the geopolitical shocks associated with Russia’s invasion of Ukraine and the 2026 Iran war. Returning to the motivating question of whether the latter places the economy back into a 2022-type risk environment, the analysis examines how the shock reshapes not only expected macroeconomic outcomes, but also predictive distributions, scenario probabilities, and the broader balance of risks. The results show that similar geopolitical shocks can generate markedly different risk profiles depending on the underlying state of the economy. While both episodes raised concerns about inflation, growth, and stagflation, the framework reveals substantial differences in the probabilities attached to these outcomes. In particular, the 2022 episode was characterized by a substantially higher baseline risk of stagflation and a more persistent increase in adverse risks following the shock. By contrast, the 2026 Iran war generates a considerably more benign risk profile despite important similarities in the nature of the underlying disturbance. More broadly, the findings demonstrate that the effects of shocks on macroeconomic risks cannot be inferred from average responses alone, but depend critically on the state of the economy in which they occur.

Section 2 introduces the methodology, section 3 presents empirical exercises, and section 4 concludes.

2. METHODOLOGY AND ESTIMATION

The framework developed in this paper builds on the premise that economically meaningful scenarios should be characterized jointly by the outcomes they imply and the probabilities attached to them. From this perspective, scenarios are not external inputs to the analysis, but endogenous objects whose likelihood and evolution are determined by the same forces shaping the broader predictive distribution. To endogenize both objects, the following sections develop a multivariate state-dependent framework in which predictive distributions, scenario probabilities, and macroeconomic risks emerge from a common set of underlying economic states. Structural shocks affect both macroeconomic outcomes and the evolution of these states, providing a natural basis for distributional impulse response analysis.

2.1 Predictive Distributions and Economic States

Following [López-Salido and Loria \(2024\)](#), predictive distributions can be interpreted through the lens of state-dependent macroeconomic dynamics. From this angle, uncertainty about future

macroeconomic outcomes reflects uncertainty about which economic environment will prevail in the future. To bring this perspective into a multivariate setting, the economy is represented by a finite number of latent states ($s_t \in 1, \dots, H$) in the spirit of [Hamilton \(1989\)](#). Conditional on occupying state h , the n -vector of endogenous variables macroeconomic outcomes, y_t , evolves according to

$$y_t = c_h + \sum_{\ell=1}^p A_{\ell,h} y_{t-\ell} + u_t, \quad u_t \mid s_t = h \sim \mathcal{N}(0, \Sigma_h), \quad (1)$$

where each state is characterized by its own dynamic behaviour and volatility structure. Throughout, u_t is an innovation with respect to the predetermined information set $\mathcal{F}_{t-1}$, so that $\mathbb{E}(u_t \mid \mathcal{F}_{t-1}, s_t) = 0$, and innovations are serially uncorrelated conditional on the state process. Letting $\pi_h = \Pr(s_t = h)$ denote the probability of occupying state h , integrating over the latent state variable and collecting lagged information in $x_t = (1, y'_{t-1}, \dots, y'_{t-p})'$ then yields the predictive density

$$f(y_t \mid x_t) \approx \sum_{h=1}^H \pi_h \phi(y_t; \beta_h x_t, \Sigma_h), \quad (2)$$

where $\phi(\cdot; \mu, \Sigma)$ denotes the Gaussian kernel with mean μ and covariance matrix Σ , and $\beta_h = (c_h, A_{1,h}, \dots, A_{p,h})$ collects the regime-specific VAR coefficients. The predictive distribution therefore arises from uncertainty about both future macroeconomic outcomes within a given state and the future realization of the states themselves.

Up to this point, the state probabilities π_h have been left unspecified. In classical regime-switching models, state dynamics are typically governed by a latent Markov process with constant transition probabilities. For the purposes of risk analysis, however, it is desirable for the likelihood of future states to vary with prevailing economic conditions. In this paper, the state probabilities are therefore allowed to depend directly on observable economic covariates, extending the density-regression framework of [Rigon and Durante \(2021\)](#) and its application to “inflation at risk” in [Schröder \(2026\)](#) to a multivariate setting. Specifically, let the mixing weights, $\pi_h(x_t; \gamma) = \nu_h(x_t) \prod_{l=1}^{h-1} (1 - \nu_l(x_t))$ and $\pi_H(x_t; \gamma) = \prod_{l=1}^{H-1} (1 - \nu_l(x_t))$, be defined by the stick-breaking weights $\nu_h(x_t)$ which can be linked directly to economic covariates via logistic gates following [Rigon and Durante \(2021\)](#). Defining $\eta_h(x_t) = \text{logit}(\nu_h(x_t)) = \log\left(\frac{\nu_h(x_t)}{1 - \nu_h(x_t)}\right)$, the mixture probabilities $\pi_h(x_t; \gamma)$ can then be written as

$$\pi_h(x_t; \gamma) = \frac{\exp\{\eta_h(x_t)\}}{1 + \exp\{\eta_h(x_t)\}} \prod_{l=1}^{h-1} \left[\frac{1}{1 + \exp\{\eta_l(x_t)\}} \right]. \quad (3)$$

Here, $\eta_h(x_t) = \psi_h x_t$ represents the log-odds of observation t being allocated to component h , conditional on not having been assigned to any of the previous components. This formulation transforms the multinomial allocation problem into a sequence of binary logistic regressions. [Online Appendix B](#) formalizes the distributional implications of this state-dependent allocation mechanism. [Proposition 1](#) shows how the distribution of the economic state evolves by averaging the regime-specific transition laws over both current economic conditions and their associated regime probabilities, while [Proposition 2](#) establishes that aggregate regime incidence is itself an endogenous distributional object, depending on where probability mass is located in the economic state space.

The framework becomes generative once structural shocks are introduced. A shock affects macroeconomic outcomes through the regime-specific VAR dynamics, but it also alters the economic conditions entering the gating mechanism via x_t . The predictive distribution therefore evolves through two distinct channels: a direct effect on outcomes within a given state and an indirect effect operating through changes in future state probabilities. Importantly, because the logistic link creates a nonlinear mapping from economic conditions to state probabilities, the effects of shocks depend on the prevailing economic environment. For example, a financial shock may push the economy into an adverse state when macroeconomic conditions are already fragile, while having little effect when economic conditions are resilient. Conversely, an improvement in financial conditions may not be sufficient to return the economy to a more benign state even if substantial, if other state variables remain weak. The framework therefore naturally generates state-dependent and asymmetric responses of predictive distributions and macroeconomic risks without requiring explicit threshold rules or nonlinear within-regime VAR dynamics. The stochastic-process properties underlying this construction are developed in [Online Appendix A](#), which provides the Markov representation of the state-dependent process, establishes validity and communication properties of the finite LSBP transition mechanism, and discusses sufficient conditions for global stability and geometric ergodicity. [Online Appendix B](#) formalizes this mechanism: [Proposition 3](#) decomposes aggregate propagation into a within-regime component and a regime-selection component, while [Proposition 4](#) shows how structural shocks alter subsequent regime incidence through the distribution of the evolving economic state. [Online Appendix C](#)

further characterizes the local nonlinearity generated by the logistic gate.

2.2 Heteroskedasticity-Based Structural Identification

Having established the general model, this section outlines the identification of structural shocks. In many “at-risk” forecasting frameworks, identification is performed directly on individual distributional objects, such as quantiles or risk measures. For example, the quantile VARs of [Forni et al. \(2021\)](#), [Iacopini et al. \(2022\)](#), [Chavleishvili and Manganelli \(2024\)](#), [Schüler \(2020\)](#) estimate separate models across quantile levels, requiring identifying assumptions to be imposed on each quantile-specific system. More recent approaches such as the QFAVAR of [Korobilis and Schröder \(2025\)](#) recover a coherent distribution by estimating multiple quantiles jointly, but identification remains at the quantile level.

In contrast, the framework proposed in this paper identifies shocks at the level of the underlying state-dependent model. The resulting structural shocks therefore retain a coherent economic interpretation across the entire predictive distribution, including its moments, quantiles, scenario probabilities, and other measures of macroeconomic risk. Moreover, because each state is characterized by its own volatility structure, the regime structure itself provides a natural source of identification.

Following [Lanne and Lütkepohl \(2008\)](#), the structural representation factorises the regime covariance through a single, regime-invariant impact matrix B and regime-specific diagonal shock variances Λ_h :

$$\Sigma_h = B \Lambda_h B', \quad \Lambda_h = \text{diag}(\lambda_{h,1}, \dots, \lambda_{h,n}). \quad (4)$$

The structural shocks are then recovered as $\varepsilon_t = B^{-1}u_t$, with $\text{Var}(\varepsilon_{t,k} \mid s_t = h) = \lambda_{h,k}$. Up to column permutation and column sign, B is statistically fully identified whenever the regime-specific variance profiles $(\lambda_{1,k}, \dots, \lambda_{H,k})$, $k = 1, \dots, n$ are pairwise non-proportional across shocks, implying that different shocks display distinct patterns of volatility across regimes ([Lanne et al. 2010](#)).

However, the heteroskedasticity-based decomposition leaves two indeterminacies that must be resolved before the columns of B can be interpreted as structural shocks. First, B and Λ_h are jointly identified only up to a column-wise rescaling. This ambiguity is removed by selecting a reference regime h^* (by default, regime (1)) and imposing $\Lambda_{h^*} = I_n$, such that the remaining

diagonal elements of Λ_h measure shock variances relative to the reference regime.¹ Second, while heteroskedasticity may identify the structural shocks, it does not determine either their ordering or their orientation. The columns of B can therefore be permuted arbitrarily and multiplied by -1 without affecting the likelihood.

To resolve this remaining permutation-sign indeterminacy, the empirical analysis imposes additional sign and narrative restrictions in the spirit of (Uhlig 2005) and (Antolín-Díaz and Rubio-Ramírez 2018). In the ideal case of non-proportional variances for all shock pairs, these restrictions serve solely as labeling devices (Lanne et al. 2010), while full identification itself continues to be achieved through the heteroskedastic structure. For weakly identified shock pairs, the sign and narrative restrictions take on an additional identifying role. Online Appendix D provides post-estimation evidence on the posterior volatility ratios for the empirical exercises.

2.3 Density Forecasts & Impulse Response Functions

Having identified and economically labelled the structural shocks, this section outlines how the framework can be used to trace their effects on the entire predictive distribution of future macroeconomic outcomes. A natural starting point is the regime-specific impulse response familiar from the regime-switching VAR literature. Conditional on holding a given regime h^* fixed throughout the forecast horizon up to $\tilde{T}$, the response of the economy to a one-standard-deviation structural shock k is given by

$$\text{IRF}_k^{h^*}(\tau) = \Phi_{h^*}^{(\tau)} B e_k \sqrt{\lambda_{h^*,k}}, \quad \tau = 0, 1, \dots, \tilde{T}, \quad (5)$$

where $\Phi_{h^*}^{(\tau)}$ denotes the companion form matrix implied by the regime-specific VAR coefficients. While informative, the regime-conditional response remains a counterfactual object. By assuming that the economy remains in regime h throughout the forecast horizon they abstract from one of the central mechanisms of the proposed framework: Structural shocks affect not only macroeconomic outcomes within a given regime, but also the probabilities of future regime realizations through the endogenous gating mechanism.

To capture this channel, one can instead construct a paired baseline and shocked conditional path that passes through the model's gating mechanism, allowing state realizations to evolve

¹To align volatility scales across variables, the data are divided by their in-sample standard deviations prior to estimation.

endogenously:²

$$y_{t+\tau}^{\text{baseline}} = \beta_{s_{t+\tau}^b} x_{t+\tau-1}^{\text{baseline}}, \quad s_{t+\tau}^b \sim \pi_h(x_{t+\tau-1}^b; \gamma), \quad (6)$$

$$y_{t+\tau}^{\text{shocked}} = \beta_{s_{t+\tau}^s} x_{t+\tau-1}^{\text{shocked}} + \mathbf{1}\{\tau = 0\} B e_k \sqrt{\lambda_{h,k}}, \quad s_{t+\tau}^s \sim \pi_h(x_{t+\tau-1}^s; \gamma), \quad (7)$$

The regime independent IRF for a given starting state s_t is then path-wise difference

$$\text{IRF}_k^h(\tau) = [y_{t+\tau}^{\text{shocked}} - y_{t+\tau}^{\text{baseline}} \mid s_t = h]. \quad (8)$$

Unlike [equation 5](#), [equation 8](#) allows for nonlinear propagation through endogenous regime selection. Under the timing of the model, the regime governing the impact transition is selected from the predetermined state and is therefore unaffected by the contemporaneous structural perturbation. The shock nevertheless changes the subsequent state distribution and thereby the probabilities with which future regime-specific laws of motion are selected. This feedback channel is formalized in [Proposition 4](#) in [Online Appendix B](#). As a result, propagation can vary with both the sign and the size of the shock, since otherwise comparable perturbations can move the economy differently across regions of the state space in which the regime-selection mechanism is more or less responsive. For the logistic gate, this local sensitivity is itself state dependent and is strongest near the switching boundary, as formalized in [Proposition 7](#) in [Online Appendix C](#). Consequently, shocks of equal absolute magnitude but opposite sign need not induce symmetric changes in future regime probabilities, and the resulting IRFs to $+\sigma$ and $-\sigma$ need not be mirror images.

While [equation 8](#) resolves the regime-conditioning problem, it continues to summarize the response of the economy through a single trajectory. In nonlinear models, however, this is generally not sufficient to characterize changes in the predictive distribution. In addition, a given structural shock may leave the trajectory at the conditional mean largely unchanged while substantially altering the dispersion, skewness, tail behaviour, or scenario probabilities of future outcomes.

These limitations point to the integration framework of [Koop et al. \(1996\)](#), who define the

²The main analysis considers temporary structural interventions. For permanent policy changes that modify the transition kernel, [Proposition 6](#) in [Online Appendix B](#) shows that the invariant distribution, and hence long-run regime incidence, may change even when the regime-selection mechanism itself is unchanged.

response of a nonlinear system as

$$\text{GIRF}^{\text{KPP}}(\tau, \delta, \omega_{t-1}) = \mathbb{E}[y_{t+\tau} | \varepsilon_{k,t} = \delta, \omega_{t-1}] - \mathbb{E}[y_{t+\tau} | \omega_{t-1}], \quad (9)$$

where both expectations integrate over all future stochastic innovations $\{\varepsilon_{t+\tau}\}_{\tau \geq 1}$ and, in this state-dependent mixture setting, also over the implied stochastic regime trajectories $\{s_{t+\tau}\}_{\tau \geq 1}$. [Koop et al. \(1996\)](#) themselves consider the distribution of (9) across histories ω_{t-1} and shock sizes δ , but the response object remains a difference of conditional means.

This paper extends this construction to the entire conditional predictive distribution. For notational convenience, let $z_t := \Lambda_{s_t}^{-1/2} \varepsilon_t$, with $z_t \sim N(0, I)$. By comparing the conditional density of $y_{t+\tau}$ under the shocked and baseline paths for each horizon τ ,

$$\text{DIRF}_k^h(\tau) = f^\sigma(y_{t+\tau}^{\text{shocked}} | s_t = h) - f(y_{t+\tau}^{\text{baseline}} | s_t = h), \quad (10)$$

where f^σ denotes the predictive density generated by perturbing structural shock k by an amount σ at impact. The key implementation detail is that both densities are obtained by Monte Carlo integration over the same future innovation sequence $\{z_{t+\tau}^{(m)}\}_{\tau \geq 1}$, $m = 1, \dots, M$, applied identically to the shocked and baseline paths, for a large number of simulated alternative shock sequences. That is, for each Monte Carlo path m :

$$y_{t+\tau}^{\text{baseline},(m)} = \beta_{s_{t+\tau}^{\text{b},(m)}} x_{t+\tau-1}^{\text{baseline},(m)} + B \Lambda_{s_{t+\tau}^{\text{b},(m)}}^{1/2} z_{t+\tau}^{(m)}, \quad (11)$$

$$y_{t+\tau}^{\text{shocked},(m)} = \beta_{s_{t+\tau}^{\text{s},(m)}} x_{t+\tau-1}^{\text{shocked},(m)} + B \Lambda_{s_{t+\tau}^{\text{s},(m)}}^{1/2} z_{t+\tau}^{(m)} + \mathbf{1}\{\tau = 0\} B e_k \sigma, \quad (12)$$

with $\{z_{t+\tau}^{(m)}\}_{\tau \geq 1}$ shared across the two paths.

The two predictive distributions differ only through the perturbation of structural shock k at $\tau = 0$ and the subsequent endogenous propagation through the state-dependent system. Because the shock has been structurally identified as described in Section 2.2, the difference between the shocked and baseline predictive distributions has the interpretation of the model-implied causal effect of that structural perturbation, conditional on the maintained identification assumptions and the initial information set, in the spirit of [Koop et al. \(1996\)](#).³

³Unlike [Koop et al. \(1996\)](#), the present construction treats the shock in deviation form. Each shocked path receives the same innovation that would otherwise have occurred, plus an additional perturbation σ in shock k . This preserves underlying shock uncertainty while shifting the predictive distribution, which is natural for distributional impulse-response analysis.

The object in [equation 10](#) inherits the nonlinear-propagation channels described above, including endogenous changes in future regime incidence and sign- and size-dependent transmission, while preserving the full shape of the predictive distribution at every horizon. Importantly, these distributional responses are themselves state dependent. Even when the structural parameters, regime-selection mechanism, and structural intervention are held fixed, the same shock can generate different future state distributions and hence different regime-incidence and distributional responses when applied from different initial economic conditions. This implication is formalized in [Corollary 1](#) in [Online Appendix B](#).

2.4 Scenario Analysis

The availability of the full predictive distribution permits the evaluation of a broad range of distributional risk measures and economically meaningful scenarios. Beyond shifts in conditional moments such as the mean and standard deviation, and movements in conditional quantiles (including the tail quantiles that drive growth-at-risk and inflation-at-risk measures) it directly delivers the probabilities of economically meaningful scenarios, such as recessions, stagflationary outcomes, or breaches of inflation tolerance bands. By construction, each of these distributional objects can be derived from either a single or multitude of structurally identified shocks and therefore retains a coherent economic interpretation. To make this precise, let a scenario be formally defined in terms of three primitive components:

Definition. A *scenario* is characterised by the triple

$$\mathcal{S} = (t_0, \Delta, \mathcal{A}),$$

where

- (i) t_0 is the *starting period* or initial state from which the scenario is evaluated, together with the implied conditioning information ω_{t_0} (in particular the initial regime s_{t_0});
- (ii) $\Delta \in \mathbb{R}^{n \times (H_{\max} + 1)}$ is the structural-shock perturbation profile, whose column $\Delta_{\cdot, \tau} = \delta_{t_0 + \tau} \in \mathbb{R}^n$ specifies the additive structural perturbation applied at horizon τ , for $\tau = 0, 1, \dots, H_{\max}$. For each simulated path m , $\varepsilon_{t_0 + \tau}^{\Delta, (m)} = \varepsilon_{t_0 + \tau}^{0, (m)} + \Delta_{\cdot, \tau}$, so that Δ is interpreted in deviation form relative to the innovations that would otherwise occur. The

profile accommodates single shocks, sustained shock sequences, and combinations of several structural shocks occurring at the same or different horizons;

- (iii) $\mathcal{A} \subseteq \mathbb{R}^{p \times (H_{\max}+1)}$ is the *scenario design profile*, an event of interest defined on the predictive paths; for example, a recession (output growth below a threshold for two consecutive quarters), a stagflationary outcome (output growth below and inflation above given thresholds simultaneously), or a breach of the inflation target over a specified horizon.

This definition is conceptually closer to classical conditional forecasting and contrasts with the *scenario-synthesis* approach of Adrian et al. (2025). In their framework, scenarios are exogenous primitives, typically obtained from judgmental sources such as central-bank staff forecasts (e.g. the Federal Reserve Tealbook). The methodology then assigns relative weights to them based on their concordance with a reference predictive density built from Growth-at-Risk and quantile regressions (Adrian et al. 2019). Both the scenarios and the reference distribution, however, are reduced-form objects: the reference density characterizes the distribution of outcomes conditional on current macroeconomic and financial conditions, but does not specify the structural shocks or transmission mechanisms through which alternative outcomes arise.

An important implication is that similar paths for observable variables need not correspond to similar economic scenarios. Comparable trajectories for inflation and output may emerge from fundamentally different structural disturbances and initial conditions, implying different transmission channels, risk profiles, and probabilities of future outcomes. In addition, an exogenously specified scenario might not be reachable under the dynamics of a given reference model or might implicitly require shock combinations that the model would deem implausibly extreme.

In the framework proposed here, by contrast, The *scenario paths* are the conditional predictive paths under the shocked and baseline trajectories defined in (10), evaluated from t_0 over the scenario horizon, with the shock profile Δ replacing the single impact perturbation. The event $\mathcal{A}$ is used to summarise the resulting predictive distribution into a path and a corresponding probability. Scenarios outcomes are therefore objects characterised jointly by their economic impact (the distribution of paths) and their probabilities (the mass the predictive distribution places on $\mathcal{A}$), with shock identification and transmission both fully internal to the model. They describe what the economy looks like along the scenario, including its central tendency, its dispersion, and its higher-order features.

Specifically, the scenario probabilities are the conditional probabilities that the event $\mathcal{A}$ is

realised under the shocked and baseline predictive distributions:

$$\Pr_{\mathcal{S}}^{\text{shocked}} = \Pr(\{y_{t_0+\tau}^{\Delta}\}_{\tau=0}^{H_{\max}} \in \mathcal{A} \mid s_{t_0}), \quad (13)$$

$$\Pr_{\mathcal{S}}^{\text{baseline}} = \Pr(\{y_{t_0+\tau}^0\}_{\tau=0}^{H_{\max}} \in \mathcal{A} \mid s_{t_0}), \quad (14)$$

both computed by Monte Carlo integration over the same future-innovation sequences, analogous to [equation 10](#). By computing the difference between the baseline and shocked probabilities, this additionally gives rise to scenario probability IRF:⁴

$$\text{IRF}^{\mathcal{S}}(\tau \mid h) = \Pr_{\mathcal{S}}^{\text{shocked}}(\tau \mid h) - \Pr_{\mathcal{S}}^{\text{baseline}}(\tau \mid h), \quad (15)$$

Tracking $\text{IRF}^{\mathcal{S}}$ across horizons answers a central question: by how much, and for how long, does an identified structural-shock profile raise or lower the probability of a given macroeconomic scenario, conditional on the prevailing regime? Beyond its interpretation as a probability response, the scenario-probability IRF admits an exact accounting decomposition. [Proposition 5](#) in [Online Appendix B](#) shows that the change in the probability of a scenario can be separated into a component associated with changes in regime incidence and a component associated with changes in the probability of the scenario within regimes. Hence, a structural shock may raise the probability of an event such as recession or stagflation because it makes adverse regimes more likely, because it worsens outcomes conditional on a given regime, or through both channels simultaneously. This decomposition makes explicit that scenario probabilities are functionals of the shock-induced predictive distribution rather than separate forecasting objects.

2.5 Estimation & Priors

The proposed framework involves joint inference over regime-specific dynamics, latent state allocations, and structural parameters. For estimation, Bayesian methods are particularly well suited as they allow uncertainty about all model components to be propagated coherently into the distributional objects of interest, including impulse responses, predictive densities, and scenario probabilities. Accordingly, this section outlines the Bayesian estimation strategy and prior specification employed throughout the empirical analysis.

⁴When scenario probabilities are reported horizon by horizon, τ indexes the corresponding truncation of the predictive path and the event $\mathcal{A}$ is understood relative to that horizon. For scenarios defined by contemporaneous conditions, this reduces to evaluating the same criterion at $t_0 + \tau$.

To regularise estimation of the regime-specific VAR coefficients and ensure stable inference in a potentially high-dimensional setting, the model combines a Minnesota prior (Litterman 1986) with sparse Bayesian learning (SBL) shrinkage (Tipping 2001). The Minnesota prior encodes the conventional belief that macroeconomic variables are persistent and that own lags are generally more informative than cross-variable lags, while the SBL component allows the degree of shrinkage to adapt flexibly across coefficients. Specifically, an inverse-gamma hyperprior is placed on the variance of each coefficient, allowing weakly identified coefficients to be shrunk more aggressively toward zero. This additional layer of shrinkage is particularly useful in the present setting, where some regimes may be supported by relatively few observations relative to the number of parameters to be estimated, thereby improving both numerical stability and estimation efficiency.⁵

For the regime-specific variances, each $\lambda_{h,k}$ in non-reference regimes is assigned an inverse-gamma prior, $\lambda_{h,k} \sim \mathcal{IG}(a_0, b_0)$, with $(a_0, b_0) = (4, 9)$. The prior is centered at a regime variance ratio of 3, corresponding to a regime standard-deviation ratio of $\sqrt{3} \approx 1.73$ relative to the reference regime, accommodating order of magnitude differences between regimes while still penalizing degenerate solutions in which a high-volatility regime captures only a handful of extreme observations.⁶ As discussed previously, the variance in a chosen reference regime h^* is held fixed at $\Lambda_{h^*} = I_n$ throughout the chain, which jointly resolves the column-rescaling indeterminacy of the factorisation $\Sigma_h = B\Lambda_h B'$ and provides the absolute scale that the heteroskedasticity-identifying contrast in $\Lambda_h - \Lambda_{h^*}$ requires.⁷

Sampling the elements of B under sign restrictions in the present setup is challenging and requires a modified algorithm. In particular, for sign-restricted SVARs, efficient algorithms such as those of Arias et al. (2018, 2025) rely on the QR rotation $B = FQ$ with orthonormal Q . In the present heteroskedastic setup, however, applying this parametrization is more involved because candidate rotations must jointly fit the covariance structure across regimes while satisfying the identifying restrictions. For this reason, this paper parametrizes B directly column-by-column, where each column $b_j \in \mathbb{R}^n$ is sampled by an elliptical slice sampler (ESS) following (Murray et al. 2010) against a Gaussian *base measure* that is used purely as the proposal-generating

⁵Cross-equation Minnesota weights are additionally rescaled by the ratio of pooled-OLS residual standard deviations, $\sigma_{OLS,i}^2/\sigma_{OLS,j}^2$, rendering the prior invariant to the units of measurement of individual variables.

⁶To further ensure numerical stability, the lower tail of λ is truncated at 10^{-6} , preventing near-zero variance realisations.

⁷Should a sampled regime allocation leave the reference regime empty at a given iteration, the most-populous regime is temporarily promoted to reference duty for that iteration.

mechanism: Let $\sigma_{\text{ols},k} = \sqrt{[\hat{\Sigma}_{\text{pool}}]_{kk}}$ denote the pooled-OLS innovation standard deviation of equation k and let $s_{kj} \in \{-1, 0, +1, \text{NaN}\}$ be the (k, j) entry of the sign restriction matrix, with NaN marking unrestricted entries. The base measure is centred at zero at every non-zero entry,

$$\mu_{0,k,j} = 0, \quad \text{for all } (k, j) \text{ with } s_{kj} \neq 0, \quad (16)$$

with standard deviation

$$\sigma_{0,k,j} = c_{\text{base}} \times \begin{cases} \kappa_{\text{sign}} \sigma_{\text{ols},k} & s_{kj} \in \{-1, +1\}, \\ \kappa_{\text{free}} \sigma_{\text{ols},k} & s_{kj} = \text{NaN}, \\ 0 & s_{kj} = 0, \end{cases} \quad (17)$$

and hyperparameters $(\kappa_{\text{sign}}, \kappa_{\text{free}}, c_{\text{base}})$ defaulting to $(0.8, 1.0, 1.1)$. The final Bayesian prior on the free entries of b_j in the column-wise update is effectively flat (improper).⁸

Because ESS parameterizes proposals through an angle on an ellipse and progressively shrinks the admissible angle bracket whenever the slice criterion is violated, every accepted draw automatically satisfies the imposed sign and rank restrictions without requiring an explicit accept-reject step. During sampling, sign orientation itself is enforced by a hard indicator $\mathbf{1}\{b_j \in \mathcal{F}\}$ that is folded into the slice target as a $\log p = -\infty$ penalty whenever any entry of b_j violates $s_{\cdot,j}$ or renders B numerically singular.⁹ Conditional on data and other model parameters, for each column j the conditional log-likelihood as a function of b_j (holding the remaining columns of B fixed) is then given by

$$\log L(b_j) = - \sum_{h=1}^H n_h \log |\det B| - \frac{1}{2} \sum_{h=1}^H \sum_{t:s_t=h} \sum_{k=1}^n \frac{\varepsilon_{t,k}^2}{\lambda_{h,k}}, \quad \varepsilon_t = B^{-1} u_t, \quad (18)$$

where $n_h = |\{t : s_t = h\}|$.

Because the columns of B are coupled in the conditional likelihood through the Jacobian

⁸the base-measure log-density is subtracted from the slice target so that the actual distribution sampled is proportional to the likelihood (times the sign and rank indicators) only. Concretely, the ESS routine targets $\mathcal{N}(b_j^{\text{free}}; \mu_0, \Sigma_0) \cdot \exp(\log f(b_j))$, with $\log f(b_j) = \log L(b_j) - \log \mathcal{N}(b_j^{\text{free}}; \mu_0, \Sigma_0)$. The Gaussian factor cancels, leaving $L(b_j)$ as the actual target. This is a deliberate design choice: the impact matrix is identified by data, sign restrictions, and heteroskedasticity, without shrinkage from a prior on B ; regularisation in the model lives in the Minnesota-SBL priors on the regime-specific VAR coefficients β_h .

⁹Equivalently, one could multiply the indicator $\mathbf{1}\{b_j \in \mathcal{F}\}$ into the prior; the two formulations yield the same target distribution. Folding the indicator into the slice target instead keeps the prior strictly Gaussian, which is what the ESS ellipse construction requires.

term $-T \log |\det B|$, any deterministic column order would update some columns under a systematically more constrained B than others during a Gibbs sweep. To mitigate systematic within-sweep ordering effects on mixing, at each iteration the n columns of B are visited in a fresh uniform random order. This diffuses the within-sweep asymmetry uniformly across columns while preserving the target distribution (Geman and Geman 1984, Roberts and Sahu 1997).

Narrative restrictions of the form

$$\text{sign}\left(\mathcal{N}(\varepsilon_{t,k_r}; t \in \mathcal{W}_r)\right) = \delta_r, \quad r = 1, \dots, R, \quad (19)$$

can be imposed following Antolín-Díaz and Rubio-Ramírez (2018), where $\mathcal{W}_r$ is a window of dates, k_r the shock index, and $\delta_r \in \{-1, +1\}$ the required sign. Because $\varepsilon_t = B^{-1}u_t$ depends only on B and the reduced-form residuals u_t , evaluating each restriction is computationally trivial and draws are accepted only when all R restrictions hold jointly; otherwise the entire parameter draw at that iteration is discarded and the chain proceeds to the next iteration.

Finally, for the mixing coefficients, sampling can be simplified by exploiting the Pólya–Gamma data augmentation scheme of Polson et al. (2013) in the spirit of (Rigon and Durante 2021). Each stick-breaking equation $h = 1, \dots, H - 1$ is a Bernoulli regression with linear predictor $\psi_{h,t} = x_t' \gamma_h$ and outcome $z_{h,t} = \mathbf{1}\{s_t = h \mid s_t \geq h\}$. Conditioning on latent variables $\omega_{h,t} \sim \text{PG}(1, \psi_{h,t})$ yields a conditionally Gaussian representation in which the response $(z_{h,t} - 1/2)/\omega_{h,t}$ regressed on x_t produces an exact Gaussian conditional for γ_h . This approach is tuning-free, computationally efficient, and typically more stable than Metropolis-Hastings alternatives. Conditional on $A_{\ell,h}, B, \Lambda$, the regime indicator s_t is then updated independently across t from its full conditional. For each t the conditional probability $\Pr(s_t = h \mid \cdot)$ combines the gate’s logistic stick-breaking prior weight $\pi_{h,t}$ with the regime-specific Gaussian likelihood of u_t under $\Omega_h = B\Lambda_h B'$ and $A_{\ell,h}$, and the new s_t is sampled from the resulting categorical distribution.¹⁰

The overall algorithm hence cycles through the following sequence:

1. **Mixing equation updates.** For each $h = 1, \dots, H - 1$, sample γ_h from its Gaussian full conditional under Pólya–Gamma augmentation.
2. **VAR coefficient updates.** For each regime h , draw $A_{\ell,h}$ equation-by-equation from its conditional Normal posterior using the sparse-precision algorithm of Rue (2001) on the

¹⁰Because the gate makes regime assignments conditionally independent across t given z_t and γ ; no forward-filter-backward-sampler step is required.

Minnesota-tightened, SBL-modulated precision.¹¹

3. **Structural impact matrix.** Recompute reduced-form residuals $u_t = y_t - A_{\ell,h}x_t$, then update B column-by-column in fresh uniform random order via the elliptical-slice scheme described above, with sign restrictions enforced through the $\log p = -\infty$ rejection in the slice target.
4. **Regime-specific structural variances.** For each non-reference regime $h \neq h^*$ and shock k , draw $\lambda_{h,k}$ from its conditional posterior.
5. **Regime assignment.** Sample s_t independently across t from the categorical full conditional that combines the LSBP gate weights $\pi_{h,t}$ with the regime-specific Gaussian likelihood of u_t .¹²
6. **Narrative restrictions.** Post burnin, evaluate the narrative restrictions and discard the entire set of draws if any restriction is violated.

Taken together, the resulting algorithm is tuning-free and efficient in practice for systems of practically relevant dimensions.

3. EMPIRICAL ILLUSTRATION

Having introduced the model and estimation algorithm, this section illustrates the proposed framework using euro area and U.S. macroeconomic data. The analysis focuses on tracing the effects of identified structural shocks through predictive densities and scenario probabilities, with particular attention devoted to geopolitical risk shocks motivated by recent events such as Russia's invasion of Ukraine and the escalation of tensions in the Middle East.

The empirical analysis proceeds in three stages. First, the estimated regime structure and the identification of structural shocks are examined. Second, the transmission of identified shocks is studied through impulse responses. Finally, the framework is used to address a policy question that emerged following the escalation of tensions between the United States and Iran, when data

¹¹Draws for which any regime-specific companion matrix violates $\rho(F_h) < 1$ are excluded during sampling. As discussed in [Online Appendix A](#), this restriction ensures stability of each fixed-regime VAR but is not, by itself, sufficient for global stability of the state-dependent process. [Online Appendix A](#) therefore also characterizes the aggregate transition kernel and provides sufficient drift conditions for geometric ergodicity.

¹²As an additional fallback against label switching, regime assignments are reordered post-estimation based on the Hungarian algorithm against a reference draw. However, for small H with strongly asymmetric regime occupancy rates and given the structure of the LSBP, label switching is largely a non-issue. For larger H and many small regimes, the role of post processing is more substantive.

was not yet available: should policymakers expect a repeat of the build-up of inflationary and stagflationary risks associated with Russia’s invasion of Ukraine or is this time different?

3.1 Data & Labelling Restrictions

The full data overview for the euro area and U.S. is given in [Table 1](#). For both economies, the model is estimated on monthly data at a uniform frequency, with $p = 12$ lags. The U.S. sample runs from June 1976 to December 2025 and the EA sample from January 1999 to December 2025. The two samples are deliberately not aligned. Instead, each is chosen to maximize the available time span while respecting the relevant institutional break: the post-Bretton-Woods floating exchange-rate era in the United States and the formation of the EMU in the euro area.

Table 1: Data Overview

MACROECONOMIC INDICATORS	SOURCE	MNEMONICS	TCODE ¹
Euro area:			
Harmonised Index of Consumer Prices excl. Energy & Food	Eurostat	HICPx	4
EA industrial production	OECD	EAIP	4
2-Year Overnight Index Swap ²	ECB	OIS2Y	2
New Composite Indicator of Systemic Stress (New CISS)	ECB	CISS	1
Geopolitical risk	Caldara and Iacoviello (2022)	GPR	3
WTI oil price	FRED	WTI	4
Broad EUR nominal effective exchange rate	ECB	EANEER	4
United States:			
Personal Consumption Expenditure excl. Energy & Food	FRED	PCEEx	4
U.S. industrial production	OECD	USIP	4
U.S. treasury securities 2-year	FRED	US2Y	2
Excess bond premium	Gilchrist and Zakrajšek (2012)	EBP	1
Geopolitical risk	Caldara and Iacoviello (2022)	GPR	3
WTI oil price	FRED	WTI	4
Broad USD nominal effective exchange rate ²	FRED	USNEER	4

¹ Tcode identifies the time series transformation: level (1), differences (2), log-level (3), log-differences (4).

² Due to data availability, the Advanced Foreign Economics USD index (DTWEXAFEGS) is extended backwards using the discontinued Major Currencies USD index (DTWEXM). Similarly, where headline OIS quotes are unavailable in the early part of the EA sample, the OIS rate is extended backwards using the German Bund 2-year rate.

For both economies, the model employs an identical seven-variable specification. The resulting setup is deliberately parsimonious, comprising only the variables necessary to identify the structural shocks of interest and trace their macroeconomic transmission. In each economy, the data can broadly be split into a domestic and a global block. The domestic block comprises the preferred policy-relevant measure of core inflation in each economy - for Personal Consumption Expenditures (PCE) in the United States and the Harmonised Index of Consumer Prices (HICP) in the euro area, both excluding food and energy - industrial production as a measure of economic activity, a 2 year interest rate as a proxy for the monetary policy stance, and a measure of financial stress. For the United States, financial stress is measured by the excess bond premium (EBP) of [Gilchrist and Zakrajšek \(2012\)](#). For the euro area, the corresponding

measure is the new CISS, the revised composite indicator of [Holló et al. \(2012\)](#). The inclusion of a financial-stress indicator is motivated by a growing literature showing that financial conditions contain important information about future macroeconomic risks and uncertainty ([Banerjee et al. 2024](#), [Korobilis et al. 2021](#), [López-Salido and Loria 2024](#), [Tagliabracci 2020](#)). In particular, [Gilchrist and Zakrajšek \(2012\)](#) show that the EBP is highly informative about future economic activity.¹³ In addition, the inclusion of a financial-stress measure helps disentangle the effects of other macroeconomic shocks from those of purely financial disturbances.

The global block contains three variables designed to capture the external risk environment in which both economies operate. Central to the subsequent empirical analysis is the text-based geopolitical risk (GPR) index of [Caldara and Iacoviello \(2022\)](#), which serves as the primary measure of geopolitical shocks.¹⁴

The global price of oil is included because commodity markets constitute one of the principal channels through which such geopolitical shocks propagate to the broader economy. Military conflict, and political instability in major producing regions can generate large and persistent movements in energy prices, which subsequently feed through to inflation, production costs, and economic activity. In this context, including oil prices helps distinguish geopolitical shocks from other commodity-market disturbances and prevents unrelated oil-price fluctuations from being absorbed by the GPR shock.

Similarly, the exchange rate is included not as a variable of primary interest, but as a high-frequency market-based indicator of global risk sentiment, cross-border capital flows, and safe-haven demand. It provides additional information on the financial transmission of geopolitical events and helps distinguish GPR-induced asset repricing from other sources of exchange-rate variation.

These broader considerations are mirrored by the sign and narrative restrictions, which are collected in [Table 2](#). As discussed previously, the model is primarily identified via the different volatility scales, with the set of relatively sparse sign and narrative restrictions serving two complementary roles. For well separated shock pairs, the restrictions resolve sign and permutation ambiguity and assign economic interpretations to each shock. For shock pairs whose volatility profiles are only weakly distinct, they provide additional identifying information.

¹³Compared with broader financial conditions indices such as the NFCI, the EBP is constructed entirely from financial-market data and captures information not already contained in standard macroeconomic variables ([Reichlin et al. 2020](#)).

¹⁴In addition a growing literature documents the substantial effects of geopolitical risk shocks on macroeconomic activity, inflation, financial markets, and asset prices.

Table 2: Data Overview

VARIABLE	SIGN RESTRICTIONS						
	DOM. SUPPLY	DOM. DEMAND	MONETARY POLICY	FINANCIAL	GEOPOLITICAL RISK	OIL PRICE	EXCHANGE RATE
HICPx/PCE _x	+	+	-			+	-/. ¹
EAIP/USIP	-	+	-	-	-		-/. ¹
OIS2Y/US2Y			+				
CISS/EBP		-	+	+	+		
GPR					+		
WTI					+	+	
EANEER/USNEER		+	+	-/+			+
SHOCK	NARRATIVE RESTRICTIONS						
	RANGE	SIGN	VARIABLE	NARRATIVE			
Financial	Sep-Oct 2008	+	CISS/EBP	Large financial stress episode in both samples.			
Geopolitical risk	Sep-Oct 2001	+	GPR	9/11. Clean GPR shock.			
Geopolitical risk	Feb-Mar 2022	+	GPR	Russian invasion of Ukraine.			

¹ (-) denotes that the restriction is left empty for the US.

Specifically, the domestic demand and domestic supply shocks are labelled in the standard fashion, with an adverse domestic supply shock raising prices and lowering output, while a positive demand shock raises both. A contractionary monetary-policy shock is oriented such that it raises the policy rate, lowers inflation and output, and tightens financial conditions. The financial shock, in turn, is associated with a tightening in the stress measure and a contraction in output. No restriction is imposed on prices, allowing the data to determine whether disruptions in financial conditions are stagflationary or disinflationary. Turning to more global shocks, the geopolitical-risk shock is oriented to increase GPR, tighten financial conditions, reduce output, and raise oil prices, broadly capturing the joint signature documented by [Caldara and Iacoviello \(2022\)](#). At the same time, the inflation response is left unrestricted, allowing the data to determine the extent to which geopolitical disturbances transmit through price or uncertainty channels. Similarly, the oil shock is associated with an increase in WTI prices and core inflation on impact but is otherwise left unrestricted.

Finally, the exchange-rate shock is only weakly signed. In the US specification, its orientation is determined solely by the exchange-rate response itself, with all remaining responses left unrestricted. The euro area specification introduces two additional impact restrictions, requiring an appreciation to lower both inflation and output. This reflects the greater openness of the euro-area economy and the possibly stronger role of exchange-rate pass-through, whereby an appreciation lowers import prices and dampens net exports. A further distinction between the two specifications concerns the response of the exchange rate to a financial shock, reflecting the

safe-haven status of the US dollar and the associated flight-to-safety flows.

The three narrative restrictions in [Table 2](#) are used to discipline two shocks for which the impact sign restrictions leave residual labelling ambiguity, namely the financial shock and the geopolitical-risk shock. The collapse of Lehman Brothers in September-October 2008 pins the sign of the financial shock by anchoring the single largest episode of financial stress in either sample, with both the excess bond premium and the CISS reaching unprecedented levels and with little scope for alternative interpretations. The September 2001 terrorist attacks and the February-March 2022 Russian invasion of Ukraine serve an analogous role for the geopolitical-risk shock. Narrative discipline is therefore applied only where the historical attribution is essentially uncontroversial. By contrast, no corresponding narrative restrictions are imposed on the oil or exchange-rate shocks. While a substantial literature exists on identifying the underlying drivers of oil-price and exchange-rate movements, doing so typically requires a more detailed modelling of the respective markets. A clear narrative attribution to a single shock is therefore more ambiguous in the current setup.


Together, the three identification layers therefore play complementary, non-redundant roles. Shocks are identified by heteroskedasticity where regime volatility profiles are distinct, sign restrictions label the identified columns and sharpen weakly identified pairs, and the narrative restriction sharpens the empirically most contested columns.

3.2 Estimated Regimes & Identification Diagnostics

Throughout the exercises, attention is restricted to a 2-regime ($H = 2$) specification, motivated by economic and statistical considerations. Economically, the endogenous logistic gate makes the regime-allocation probability a flexible nonlinear function of the lagged state. Episodes with very different observable triggers - a financial-stress shock, a sudden widening of sovereign spreads, or a pandemic-driven volatility spike - can each push the economy into the volatile state through their own configuration of state variables, without requiring a separate regime for each stress type. Heterogeneity in the trigger for switching is therefore absorbed by the gate, while the high-volatility regime itself captures the common high-volatility, high-co-movement dynamics that follow. From this perspective, additional regimes might primarily subdivide a broader stress state rather than introduce economically distinct forms of propagation that are central to the policy questions considered later in the paper.

Statistically, two regimes keep the number of state-dependent parameters and the dimension

of the gating regression small relative to the available sample, which aids the stability of the posterior regime probabilities and the identification of the state-dependent impulse responses examined later in the paper.¹⁵




Ultimately, however, the adequacy of this parsimonious two-regime specification remains an empirical question. In particular, the estimated regimes should display sufficiently distinct covariance structures to support heteroskedasticity-based identification. To assess this, [Figure 1](#) reports the posterior distribution of the regime-specific log-variance ratios. For all shocks, most posterior mass lies above zero, implying systematically larger variances in the high volatility regime. At the same time, the volatility shifts differ across shocks, indicating that the estimated heteroskedasticity is not equivalent to a single common scale factor applied uniformly to the covariance matrix. Taken together, these results point to a clear separation in the covariance structures.

However, separation is not uniformly sharp. Considerable overlap remains across several posterior distributions, implying that the shocks should not be interpreted as being cleanly ranked by the magnitude of their volatility shifts. Accordingly, heteroskedasticity provides substantially sharper statistical identification for some shock pairs than for others. While the impact matrix is generically identified up to column sign and permutation when the regime-specific volatility profiles are distinct, limited separation between some profiles implies weak statistical identification in finite samples. The sign and narrative restrictions therefore retain an active role for these less sharply identified pairs. They label the statistically identified columns

¹⁵In addition, in pre-runs with additional regimes, the third regime is either degenerate - a “phantom” state with near-zero posterior allocation probability and parameters poorly identified from a handful of observations - or it singles out the COVID-19 episode as an outlier regime driven almost entirely by the March-April 2020 observations. Neither outcome alters the dating of the high-volatility regime shown here.

where the rank condition holds, and sharpen identification for shock pairs whose volatility profiles are only weakly distinct. Additional diagnostics on the volatility structure and on the sign and narrative restrictions are reported in [Online Appendix D](#). Consistent with the results shown here, they suggest that the volatility states are separated, though not uniformly across shocks, and that the labelling restrictions are compatible with the data and informative for shock interpretation.

Figure 2: Posterior mean regime path with NBER recessions, financial & GPR events



Notes: the solid black line plots the posterior probability of the high-volatility regime, computed as the share of retained Gibbs draws in which the date- t regime indicator equals the crisis state. Vertical dotted lines mark salient events which are color-coded by category (blue for macro-financial events, red for GPR events, and green for monetary events). Shaded bands denote NBER recessions (US) and CEPR-EABCN recessions (EA); the additional shaded band in the EA panel highlights the sovereign debt crisis.

Having established the heteroskedasticity structure underlying the two-regime specification, the next question is when these regimes are active. [Figure 2](#) reports the estimated probability of the high-volatility regime for both economies. Several features stand out. First, regime switching is highly episodic rather than persistent. In both economies, the probability remains close to

zero for extended periods before rising sharply during episodes of macroeconomic, financial, or geopolitical stress.

Second, movements in the regime probabilities align closely with major historical events. In the euro area, the high-volatility regime is activated during episodes such as the 9/11 attacks, the Global Financial Crisis, the Greek bailout and euro area sovereign debt crisis, the COVID-19 pandemic, and the inflationary and energy-market disruptions following Russia's invasion of Ukraine. The U.S. exhibits a similar pattern, with elevated regime probabilities during the Volcker disinflation, Black Monday, the Gulf War recession, the Global Financial Crisis, the COVID-19 selloff, and the geopolitical and monetary-policy turbulence associated with 2022 and beyond. Additional visible bumps appear around the 2003 Iraq war, Hurricane Katrina, and an early-2024 episode in the post Israel-Hamas window.

Outside these episodes, the probability of the high-volatility regime remains close to zero, indicating that the low-volatility regime provides an adequate description of economic dynamics during normal times or after the initial high volatility spell. More importantly, the fact that financial, economical, and geopolitical events all activate the same regime provides a natural empirical counterpart to the discussion above, suggesting that a simple two-regime specification can accommodate a wide range of disturbances, with the regimes capturing the resulting state of the economy rather than the underlying source of the shock.

3.3 Shock Transmission


Before turning to the density impulse responses and scenario analysis, it is useful to briefly examine conventional regime-dependent impulse responses. Although the main contribution of the framework lies in tracing the effects of shocks through predictive densities and scenario probabilities, regime-dependent impulse responses provide a useful first illustration of the nonlinear transmission mechanisms embedded in the model. The full set of impulse response functions to all shocks is provided in [Online Appendix E](#).

[Figure 3](#) reports the response of the system to a one-standard-deviation geopolitical risk shock in the low- and high-volatility regimes for the euro area in the top panel and the United States in the bottom panel. A few features stand out.

First, the transmission of geopolitical shocks is clearly state dependent. In both economies, the responses of economic activity, financial conditions, and commodity prices are amplified when the shock occurs in the high-volatility regime. The same identified disturbance therefore

generates markedly different macroeconomic outcomes depending on the prevailing state of the economy. One notable difference between the two economies concerns the response of industrial production. While activity declines persistently in the euro area, the corresponding US response is considerably more muted and eventually turns positive in the high-volatility regime.

Figure 3: Regime-specific impulse response functions, GPR shock



Notes: The blue (purple) line plot the impulse response function of the full set of variables to a one standard deviation GPR shock in the calm regime for the euro area (United States). The yellow (red) line correspond to the shock response in the high-volatility regime.

This pattern might suggest the presence of offsetting channels that are less relevant in the euro area. In particular, many of the geopolitical episodes associated with elevated regime probabilities in the United States coincide with periods of increased defense spending, which may

partially offset the contractionary effects of heightened uncertainty. In addition, the expansion of domestic energy production over the latter part of the sample implies that increases in oil prices need not operate solely as an adverse cost shock, but may also support activity through the energy sector.

Second, the amplification is particularly pronounced for variables associated with financial stress and commodity markets. Measures of financial conditions deteriorate more strongly in the high-volatility regime, while oil prices display larger and more persistent increases. These effects are accompanied by a sharper decline in industrial production on impact particularly for the euro area, suggesting that geopolitical shocks propagate more forcefully through both financial and commodity-market channels during periods of elevated uncertainty.

In contrast, the response of inflation and interest rates is insignificant but comparatively similar across regimes. The primary effect of entering the high-volatility regime is therefore not simply a stronger inflationary response, but a broader amplification of macroeconomic and financial vulnerabilities.


3.4 Distributional transmission & scenario responses

As tensions between the United States and Iran escalated, policymakers were confronted with a familiar question: should the build-up of inflationary and stagflationary risks associated with Russia's invasion of Ukraine be expected to re-emerge? Yet the challenge in answering such questions is that they arise precisely when the relevant macroeconomic outcomes have not yet materialized. Historical episodes therefore become a natural point of reference.

Among the available benchmarks, Russia's invasion of Ukraine and the preceding build-up in risk is particularly informative. It represents the most recent major geopolitical disruption in the sample and generated exactly the combination of inflationary pressures, growth concerns, and financial-market stress that policymakers sought to assess in the new environment. At the same time, inflation, monetary policy, energy markets, and broader macroeconomic conditions differed substantially between early 2022 and the current sample edge. The relevant question is therefore not whether the two episodes are identical, but whether a shock comparable in magnitude to that observed around the invasion of Ukraine would generate the same distribution of outcomes today.

To isolate this issue, the scenario exercise holds the geopolitical shock fixed at its realized magnitude in January 2022. While this date predates the invasion in late February, [Figure 4](#),

Figure 4: Russia-Ukraine geopolitical events



Notes: The panels show the GPR shocks identified by the model for the Euro area and the US, as well as markers for specific event dates surrounding the Russian invasion of Ukraine.


shows that the model already identifies a substantial geopolitical shock during the period of escalating military tensions and increasingly explicit warnings of an imminent invasion.¹⁶ Using this period as the calibration point therefore captures the anticipation phase of the crisis and more cleanly isolates the geopolitical risk component from the actual geopolitical event itself.¹⁷ Once recovered from the historical decomposition, the shock is applied to two different starting conditions: the pre-invasion state of January 2022 and the sample edge in December 2025, immediately preceding the renewed escalation of tensions between the United States and Iran and the large-scale U.S. military build-up in the Middle East in January 2026. Any differences in the resulting predictive distributions and scenario probabilities can therefore be attributed to differences in initial conditions and endogenous regime dynamics rather than to differences in the shock itself.

Figure 5 reports the corresponding density impulse response functions as outlined in subsection 2.3 for inflation and industrial production after 6, 12, and 24 months. Rather than summarizing the effects of a shock through conditional mean responses, the density impulse responses characterize how probability mass is redistributed across the joint predictive distribution following a geopolitical risk shock, conditional on January 2022 and December 2025 macro data. Positive regions indicate combinations of inflation and activity that become more likely

¹⁶Figure F3 shows the realized GPR, CISS, EBP, and VIX in addition to the identified GRP shocks. In line with Figure 4, the charts suggest that the immanent invasion was in part priced in already in January, while financial market stress peaked particularly in March.

¹⁷Calibrating the exercise to either February or the average between January and March instead, leaves the results largely unaffected.

Figure 5: Bivariate density impulse response functions



Notes: The panels show the density impulse response functions for the euro area (top) and the United States (bottom) after 6, 12, and 24 months. The charts on the left show the response to the GPR shock associated with the Russian invasion of Ukraine in January 2022, the charts on the right show the response to the same shock applied to December 2025. In each panel, blue (purple) denotes areas in which probability mass decreases after the shock and yellow (red) areas indicate regions in which probability mass increases.

after the shock, while negative regions indicate outcomes whose probability declines.

Several observations emerge. First, the redistribution of probability mass differs markedly across starting conditions, despite the shock itself being held fixed. The resulting changes are therefore attributable to differences in initial conditions and endogenous regime dynamics rather than to differences in the shock itself.

Second, the density responses reveal dimensions of shock transmission that are obscured by conventional impulse responses. Rather than focusing on expected outcomes alone, they show how geopolitical shocks redistribute probability mass across the inflation-output space. Importantly, these shifts extend beyond the vicinity of the modal forecast and reach regions associated with adverse tail outcomes, including combinations of elevated inflation and weak or negative growth.

The nature of these redistributions also differs across economies. In the euro area, on impact probability mass shifts more visibly towards regions associated with weaker activity and elevated inflation, consistent with the stagflationary patterns already suggested by the impulse responses. In the United States, the redistribution is less concentrated in the adverse growth region, reflecting the more resilient activity response discussed above. At the same time, the shock response in the U.S. appears more persistent in the direction of output growth.


A second notable feature is the difference between the January 2022 and December 2025 environments. In both economies, the density responses at the later anchor appear economically more benign. The same geopolitical shock therefore generates a different balance of risks when applied to the current macroeconomic environment than when applied to the pre-invasion economy of early 2022. This finding already suggests that a repetition of the 2022 shock need not imply a repetition of the 2022 risk profile. This difference is consistent with the endogenous-selection mechanism formalized in [Proposition 4](#) in [Online Appendix B](#).

While the density impulse responses provide a comprehensive picture of how geopolitical shocks reshape the joint inflation-output distribution, interpreting changes in probability mass directly can be difficult from a policy perspective. The next step is therefore to summarize these distributional shifts in terms of economically meaningful scenarios in the spirit of [subsection 2.4](#). [Figure 6](#) reports the paths of inflation and industrial production growth together with the corresponding stagflation-risk profiles under both starting conditions. Stagflation is defined as the joint occurrence of inflation exceeding 2% and negative output growth. All scenario definitions are expressed in year-on-year terms, reflecting the medium-term horizon over which inflationary and recessionary pressures are typically assessed and reducing the influence of transitory month-to-month fluctuations.

Several findings stand out. First, the same geopolitical shock generates markedly different outcomes across starting conditions. In both economies, the responses associated with the December 2025 anchor are considerably more benign than those obtained from the January 2022 environment.

Second, these differences translate directly into the evolution of stagflation risk. In both economies, the January 2022 shock produces a large and persistent increase in the probability of stagflation, whereas the corresponding response at the December 2025 anchor is substantially smaller throughout the forecast horizon. A key difference is the substantially lower baseline risk in December 2025. Consistent with the density impulse responses, the same shock therefore

Figure 6: Scenario paths and stagflation risk responses



Notes: The panels show the paths (left) of yoy inflation and industrial production growth and stagflation risk responses (right) for the euro area (top) and United States (bottom). Blue (purple) lines denote responses conditional on January 2022 and yellow (red) lines responses conditional on December 2025. Dashed lines show the response of the baseline in absence of the shock, i.e. the baseline probability conditional on the starting state, and solid lines the probability after the shock. All probability paths are smoothed with 3m moving averages. Stagflation risk is defined as the joint scenario of inflation exceeding 2% yoy and yoy industrial production growth falling below 0%.

generates a materially different balance of risks when applied to a different macroeconomic environment. Second, these differences translate directly into the evolution of stagflation risk. This result provides an empirical counterpart to [Corollary 1](#) in [Online Appendix B](#). Even when the structural intervention and model parameters are held fixed, differences in the initial economic state can induce different future state distributions, regime-incidence responses, and ultimately different distributions of macroeconomic risk.

Taken together, the results highlight the value of expressing forecasts in terms of economically meaningful scenarios. While conventional conditional forecasts describe the expected path of the economy following a disturbance, the present framework quantifies the associated probabilities of inflationary, recessionary, and stagflationary outcomes. This distinction proves economically

important in the present application. Although the geopolitical shock is held fixed at its January 2022 magnitude, the implied stagflation risks differ substantially across starting conditions, indicating that historical episodes cannot be evaluated solely through the lens of their observed shocks but must also be interpreted in light of the economic environment in which they occur.

4. CONCLUSION

This paper develops a structural mixture VAR framework for analyzing the distributional consequences of macroeconomic shocks. The framework combines endogenous regime selection, heteroskedasticity-based structural identification, and probabilistic scenario analysis within a common multivariate model. The formal results characterize how aggregate transmission operates through both within-regime dynamics and endogenous changes in regime incidence, how structural shocks affect future regime probabilities through the evolving distribution of economic conditions, and how changes in policy-relevant scenario probabilities can be decomposed into regime-incidence and within-regime components. The framework thereby extends conventional impulse-response analysis from expected paths to the full predictive distribution and allows identified shocks to be evaluated in terms of their effects on the balance of macroeconomic risks.

The scenario exercise illustrates the practical value of this perspective. Using the geopolitical shock associated with the period immediately preceding Russia's invasion of Ukraine as a benchmark, the analysis shows that a shock of comparable magnitude generates substantially different risk profiles when applied to the economic environment prevailing at the end of 2025. The results suggest that historical episodes cannot be evaluated solely through the lens of the shocks they contain. Changes in the underlying state of the economy can materially alter the distribution of future outcomes and the probabilities attached to policy-relevant scenarios.

More broadly, the framework provides a tractable way of integrating structural identification, density forecasting, and scenario analysis within a unified empirical model. While the application has focused on geopolitical risk, the approach is sufficiently general to accommodate a wide range of macroeconomic, financial, and policy shocks. This opens several avenues for future research, including applications to monetary policy, financial stability, climate-related risks, and other settings in which policymakers are concerned not only with expected outcomes but also with the distribution of risks surrounding them.

REFERENCES

- ADRIAN, M. T., D. GIANNONE, M. LUCIANI, AND M. WEST (2025): “Scenario synthesis and macroeconomic risk,” Working Paper 105 (2025), International Monetary Fund.
- ADRIAN, T., N. BOYARCHENKO, AND D. GIANNONE (2019): “Vulnerable Growth,” *American Economic Review*, 109, 1263–1289.
- ANTOLÍN-DÍAZ, J. AND J. F. RUBIO-RAMÍREZ (2018): “Narrative sign restrictions for SVARs,” *American Economic Review*, 108, 2802–2829.
- ARIAS, J. E., J. F. RUBIO-RAMÍREZ, AND M. SHIN (2025): “A Gibbs Sampler for Efficient Bayesian Inference in Sign-Identified SVARs,” *arXiv preprint arXiv:2505.23542*.
- ARIAS, J. E., J. F. RUBIO-RAMÍREZ, AND D. F. WAGGONER (2018): “Inference based on structural vector autoregressions identified with sign and zero restrictions: Theory and applications,” *Econometrica*, 86, 685–720.
- BANERJEE, R., J. CONTRERAS, A. MEHROTRA, AND F. ZAMPOLLI (2024): “Inflation at risk in advanced and emerging market economies,” *Journal of International Money and Finance*, 103025.
- BEKAERT, G. AND A. POPOV (2019): “On the Link between the Volatility and Skewness of Growth,” *IMF Economic Review*, 67, 746–790.
- BERNANKE, B. (2024): “Forecasting for monetary policy making and communication at the Bank of England: a review,” *Bank of England*, April.
- CALDARA, D. AND M. IACOVIELLO (2022): “Measuring geopolitical risk,” *American economic review*, 112, 1194–1225.
- CHAVLEISHVILI, S. AND S. MANGANELLI (2024): “Forecasting and stress testing with quantile vector autoregression,” *Journal of Applied Econometrics*, 39, 66–85, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.3009>.
- FEDERAL OPEN MARKET COMMITTEE (2025): “Statement on longer-run goals and monetary policy strategy,” vol. 13, 2023.
- FORNI, M., L. GAMBETTI, AND L. SALA (2021): “Downside and Upside Uncertainty Shocks,” Discussion Paper 15881, Centre for Economic Policy Research.
- GEMAN, S. AND D. GEMAN (1984): “Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images,” *IEEE Transactions on pattern analysis and machine intelligence*, 721–741.

- GILCHRIST, S. AND E. ZAKRAJŠEK (2012): “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102, 1692–1720.
- HAMILTON, J. D. (1989): “A new approach to the economic analysis of nonstationary time series and the business cycle,” *Econometrica: Journal of the econometric society*, 357–384.
- HERWARTZ, H. AND H. LÜTKEPOHL (2014): “Structural vector autoregressions with Markov switching: Combining conventional with statistical identification of shocks,” *Journal of Econometrics*, 183, 104–116.
- HOLLÓ, D., M. KREMER, AND M. LO DUCA (2012): “CISS-a composite indicator of systemic stress in the financial system,” Working Paper 1426, European Central Bank.
- IACOPINI, M., F. RAVAZZOLO, AND L. ROSSINI (2022): “Bayesian Multivariate Quantile Regression with Alternative Time-varying Volatility Specifications,” *arXiv:2211.16121*.
- KAMPS, C., M. BUSSIÈRE, B. NIESSNER, O. TRISTANI, K. P. CHRISTOFFEL, S. KAPADIA, G. FERRERO, T. VLASSOPOULOS, R. MOTTO, G. NUÑO, ET AL. (2025): “Report on monetary policy tools, strategy and communication,” Occasional Paper 372, European Central Bank.
- KOOP, G., M. H. PESARAN, AND S. M. POTTER (1996): “Impulse response analysis in non-linear multivariate models,” *Journal of Econometrics*, 74, 119–147.
- KOROBILIS, D. (2017): “Quantile Regression Forecasts of Inflation under Model Uncertainty,” *International Journal of Forecasting*, 33, 11–20.
- KOROBILIS, D., B. LANDAU, A. MUSSO, AND A. PHELLA (2021): “The Time-Varying Evolution of Inflation Risks,” Working Paper Series 2600, European Central Bank.
- KOROBILIS, D. AND M. SCHRÖDER (2025): “Monitoring multi-country macroeconomic risk: A quantile factor-augmented vector autoregressive (QFAVAR) approach,” *Journal of Econometrics*, 249, 105730.
- LANNE, M. AND H. LÜTKEPOHL (2008): “Identifying Monetary Policy Shocks via Changes in Volatility,” *Journal of Money, Credit and Banking*, 40, 1131–1149.
- LANNE, M., H. LÜTKEPOHL, AND K. MACIEJOWSKA (2010): “Structural vector autoregressions with Markov switching,” *Journal of Economic Dynamics and Control*, 34, 121–131.
- LITTERMAN, R. B. (1986): “Forecasting with Bayesian vector autoregressions—five years of experience,” *Journal of Business & Economic Statistics*, 4, 25–38.
- LÓPEZ-SALIDO, D. AND F. LORIA (2024): “Inflation at risk,” *Journal of Monetary Economics*, 145, 103570.

- MONACHE, D. D., A. D. POLIS, AND I. PETRELLA (2024): “Modeling and Forecasting Macroeconomic Downside Risk,” *Journal of Business & Economic Statistics*, 42, 1010–1025.
- MURRAY, I., R. ADAMS, AND D. MACKAY (2010): “Elliptical slice sampling,” in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, JMLR Workshop and Conference Proceedings, 541–548.
- POLSON, N. G., J. G. SCOTT, AND J. WINDLE (2013): “Bayesian Inference for Logistic Models Using Polya-Gamma Latent Variables,” *Journal of the American Statistical Association*, 108, 1339–1349.
- REICHLIN, L., G. RICCO, AND T. HASENZAGL (2020): “Financial Variables as Predictors of Real Growth Vulnerability,” Discussion Paper 05/2020, Deutsche Bundesbank.
- RIGOBON, R. (2003): “Identification through heteroskedasticity,” *Review of Economics and Statistics*, 85, 777–792.
- RIGON, T. AND D. DURANTE (2021): “Tractable Bayesian density regression via logit stick-breaking priors,” *Journal of Statistical Planning and Inference*, 211, 131–142.
- ROBERTS, G. O. AND S. K. SAHU (1997): “Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler,” *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 59, 291–317.
- RUE, H. (2001): “Fast sampling of Gaussian Markov random fields,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63, 325–338.
- SCHRÖDER, M. (2026): “Mixing It Up: Inflation at Risk,” *Journal of Money, Credit and Banking*, preprint.
- SCHÜLER, Y. S. (2020): “The Impact of Uncertainty and Certainty Shocks,” Discussion Papers 14/2020, Deutsche Bundesbank.
- TAGLIABRACCI, A. (2020): “Asymmetry in the Conditional Distribution of Euro-area Inflation,” Working Paper 1270, Bank of Italy.
- TIPPING, M. E. (2001): *Journal of Machine Learning Research*, 1, 211.
- UHLIG, H. (2005): “What are the effects of monetary policy on output? Results from an agnostic identification procedure,” *Journal of Monetary Economics*, 52, 381–419.
- WAGGONER, D. F. AND T. ZHA (1999): “Conditional forecasts in dynamic multivariate models,” *Review of Economics and Statistics*, 81, 639–651.

ONLINE APPENDIX

Maximilian Schröder

APPENDIX A: MARKOV REPRESENTATION AND STABILITY CONDITIONS

Let $X_t \in \mathcal{X} \subseteq \mathbb{R}^d$ denote the predetermined state of the economy and let $S_{t+1} \in \{1, \dots, H\}$ denote the regime governing the transition from t to $t + 1$. For any measurable set $B \subseteq \mathcal{X}$, define

$$p_h(x) = \Pr(S_{t+1} = h \mid X_t = x), \quad Q_h(x, B) = \Pr(X_{t+1} \in B \mid X_t = x, S_{t+1} = h). \quad (\text{A1})$$

By the law of total probability, the transition kernel of the state process is therefore

$$\begin{aligned} P(x, B) &= \Pr(X_{t+1} \in B \mid X_t = x) \\ &= \sum_{h=1}^H \Pr(X_{t+1} \in B, S_{t+1} = h \mid X_t = x) = \sum_{h=1}^H p_h(x) Q_h(x, B). \end{aligned} \quad (\text{A2})$$

Condition 1: Markov state and valid transition mechanism. The first requirement is that the chosen state contain all information needed to determine the one-step transition law and that the regime probabilities define valid probability weights. In the baseline specification,

$$X_t = (Y'_t, Y'_{t-1}, \dots, Y'_{t-p+1})'. \quad (\text{A3})$$

Because the gating variables are measurable functions of X_t , no lagged regime information enters the gate, and the regime-specific transition law depends on past observations only through X_t , the conditional distribution of X_{t+1} given the full history depends on that history only through X_t . In particular, the one-step joint transition law satisfies

$$\Pr(X_{t+1} \in B, S_{t+1} = h \mid X_t = x) = p_h(x) Q_h(x, B).$$

Hence X_t constitutes a Markov state. A valid transition mechanism further requires

$$p_h(x) \geq 0, \quad \sum_{h=1}^H p_h(x) = 1.$$

These conditions hold by construction under the finite logistic stick-breaking specification. Let

$$v_h(x) = \Lambda(\eta_h(x)), \quad \Lambda(u) = \frac{1}{1 + \exp(-u)},$$

for $h = 1, \dots, H - 1$. Then

$$p_h(x) = \begin{cases} v_h(x) \prod_{\ell < h} [1 - v_\ell(x)], & h < H, \\ \prod_{\ell < H} [1 - v_\ell(x)], & h = H. \end{cases} \quad (\text{A4})$$

Hence the LSBP weights are nonnegative and sum to one for every x . With finite logits, they are in fact strictly positive.

Condition 2: Irreducibility and aperiodicity. The second requirement concerns communication across the economically relevant state space. ψ -irreducibility requires that every measurable set with positive ψ -measure in the recurrent state space can be reached with positive probability, while aperiodicity rules out a nontrivial cyclic decomposition of that state space (Meyn and Tweedie 2012).

In the present companion VAR, the innovation entering a single transition generally has lower dimension than the companion state when $p > 1$. Irreducibility therefore follows from a finite-step rather than a one-step support argument. Conditional on any fixed sequence of regimes over the next p transitions, the state X_{t+p} is an affine function of X_t and the innovation vector

$$(u'_{t+1}, \dots, u'_{t+p})'.$$

Under the companion structure, successive reduced-form innovations enter the contemporaneous block through the I_n block of G and are subsequently shifted through the lag stack by the companion dynamics. Hence the mapping from the p successive innovations to X_{t+p} has full rank. Since $\Sigma_h \succ 0$ in every regime, the Gaussian innovations have everywhere-positive densities. Together with the strictly positive LSBP weights, this implies that the aggregate p -step transition kernel has an everywhere-positive density on the state space.

Aperiodicity concerns return times of the continuous state process rather than return times of the regime indicator. From the finite-step support argument above, for any measurable set C with positive Lebesgue measure,

$$P^p(x, C) > 0$$

for every x in the state space. Applying one additional transition and using the everywhere-positive innovation density and strictly positive regime probabilities also gives

$$P^{p+1}(x, C) > 0.$$

Thus a set of positive measure can be revisited at two consecutive return horizons, p and $p + 1$. Since

$$\gcd(p, p+1) = 1,$$

the possible return times cannot share a nontrivial common divisor. Hence the ψ -irreducible aggregate chain is aperiodic (Meyn and Tweedie 2012).

Condition 3: Stability and geometric ergodicity. Let $b_h = (c'_h, 0, \dots, 0)'$ denote the companion-form intercept, let F_h denote the regime-specific companion matrix, and define $G = [I_n, 0_{n(p-1) \times n}]'$, which maps the reduced-form innovation into the contemporaneous block of the companion state. Conditional on $S_{t+1} = h$, in the present model, the companion dynamics are

$$X_{t+1} = b_h + F_h X_t + G u_{t+1}. \quad (\text{A5})$$

In application, the regime-specific stability restriction

$$\rho(F_h) < 1, \quad h = 1, \dots, H, \quad (\text{A6})$$

is imposed, where $\rho(F_h)$ denotes the spectral radius of F_h . This rules out explosive dynamics when a particular regime is held fixed. It does not, however, establish stability of the switching process as a whole (Francq and Zakoian 2001). Even when every regime is individually stable, switching among different transition matrices can generate dynamics that are not globally stable. Given the irreducibility and aperiodicity conditions established above, global stability can instead be assessed at the level of the aggregate process, for example through drift conditions of the type developed in Meyn and Tweedie (2012).

Let $V : \mathcal{X} \rightarrow [1, \infty)$ be a measurable function that is unbounded off petite sets. For the transition kernel P , define

$$PV(x) := \int_{\mathcal{X}} V(y) P(x, dy) = \mathbb{E}[V(X_{t+1}) \mid X_t = x], \quad (\text{A7})$$

where $PV(x)$ is the expected next-period value of V when the current state is x , following Meyn and Tweedie (2012). For a function V that is unbounded off petite sets, the geometric drift condition can then be written as

$$PV(x) \leq \lambda V(x) + L, \quad 0 < \lambda < 1, \quad L < \infty. \quad (\text{A8})$$

Intuitively, when $V(x)$ becomes large, the coefficient $\lambda < 1$ prevents the expected next-period value of V from increasing proportionally without bound. Together with the conditions above, this drift condition yields geometric ergodicity of the aggregate state process.

For the present model, recall that the aggregate transition kernel is

$$P(x, B) = \sum_{h=1}^H p_h(x) Q_h(x, B). \quad (\text{A9})$$

Substituting this mixture representation into (A7) gives, for fixed x ,

$$\begin{aligned} PV(x) &= \int_{\mathcal{X}} V(y) \left[\sum_{h=1}^H p_h(x) Q_h(x, dy) \right] = \sum_{h=1}^H p_h(x) \int_{\mathcal{X}} V(y) Q_h(x, dy) \\ &= \sum_{h=1}^H p_h(x) \mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h]. \end{aligned} \quad (\text{A10})$$

Hence, stability is a property of the probability-weighted aggregate transition mechanism rather than of any single regime in isolation. A simple sufficient route to the aggregate drift condition (A8) is to establish, for the same $\lambda < 1$ and $L < \infty$,

$$\mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h] \leq \lambda V(x) + L, \quad h = 1, \dots, H. \quad (\text{A11})$$

Indeed, if (A11) holds uniformly across regimes, then (A10) and $\sum_{h=1}^H p_h(x) = 1$ imply

$$PV(x) \leq \sum_{h=1}^H p_h(x) [\lambda V(x) + L] = \lambda V(x) + L. \quad (\text{A12})$$

Thus, a uniform regime-wise drift bound is sufficient for the required aggregate drift condition. It is not necessary, however, because (A8) may hold through the probability-weighted combination in (A10) even when some regimes do not satisfy the same bound individually.

For the companion dynamics in (A5), a convenient way to verify (A11) is to use a quadratic Lyapunov function. Following [Meyn and Tweedie \(2012\)](#), let

$$V(x) = 1 + x' M x, \quad M \succ 0. \quad (\text{A13})$$

A sufficient condition is that the same matrix M provides a contraction in every regime. Specifically, suppose there exist $M \succ 0$ and $0 < \bar{\rho} < 1$ such that

$$F_h' M F_h \preceq \bar{\rho} M, \quad h = 1, \dots, H. \quad (\text{A14})$$

The important feature of (A14) is that the same matrix M must work for every regime. This condition is therefore stronger than the separate restrictions $\rho(F_h) < 1$. The common Lyapunov condition is one sufficient way of establishing the required aggregate contraction.

Assume throughout this argument that

$$\mathbb{E}[u_{t+1} \mid X_t, S_{t+1} = h] = 0, \quad \mathbb{E}[u_{t+1}u'_{t+1} \mid X_t, S_{t+1} = h] = \Sigma_h,$$

with Σ_h finite for every h .

Given the regime-specific companion dynamics

$$\begin{aligned} \mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h] &= 1 + \mathbb{E}[X'_{t+1}MX_{t+1} \mid X_t = x, S_{t+1} = h] \\ &= 1 + (b_h + F_hx)'M(b_h + F_hx) \\ &\quad + \mathbb{E}[u'_{t+1}G'MGu_{t+1} \mid X_t = x, S_{t+1} = h], \end{aligned} \quad (\text{A15})$$

where the cross term involving u_{t+1} vanishes under the conditional mean-zero assumption.

By Young's inequality, for any $\delta > 0$,

$$(b_h + F_hx)'M(b_h + F_hx) \leq (1 + \delta)x'F'_hMF_hx + (1 + \delta^{-1})b'_hMb_h. \quad (\text{A16})$$

And subsequently

$$\begin{aligned} \mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h] &\leq 1 + (1 + \delta)x'F'_hMF_hx + (1 + \delta^{-1})b'_hMb_h \\ &\quad + \mathbb{E}[u'_{t+1}G'MGu_{t+1} \mid X_t = x, S_{t+1} = h], \end{aligned} \quad (\text{A17})$$

Using (A14),

$$x'F'_hMF_hx \leq \bar{\rho}x'Mx.$$

For the term involving the intercepts we have

$$(1 + \delta^{-1})b'_hMb_h \leq (1 + \delta^{-1}) \max_{1 \leq h \leq H} b'_hMb_h < \infty, \quad (\text{A18})$$

given $H < \infty$ and finite b_h . Under the maintained assumption that $u_{t+1} \mid S_{t+1} \sim N(0, \Sigma_h)$, for the terms involving expectations, we have

$$\mathbb{E}[u'_{t+1}G'MGu_{t+1} \mid X_t = x, S_{t+1} = h] = \text{tr}(G'MG\Sigma_h) + \mu'_uG'MG\mu_u. \quad (\text{A19})$$

Because $\mu_u = 0$, finite H , and finite G , Σ_h , and M

$$\mathbb{E}[u'_{t+1}G'MGu_{t+1} \mid X_t = x, S_{t+1} = h] = \text{tr}(G'MG\Sigma_h) < \infty. \quad (\text{A20})$$

Hence,

$$\mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h] \leq 1 + (1 + \delta)\bar{\rho}x'Mx + K_\delta, \quad (\text{A21})$$

where $K_\delta < \infty$ uniformly bounds the contributions of the regime-specific intercept and the conditional innovation variance.

Since $\bar{\rho} < 1$, $\delta > 0$ can be chosen sufficiently small that

$$\lambda := (1 + \delta)\bar{\rho} < 1. \quad (\text{A22})$$

Because $V(x) = 1 + x'Mx$, and with $L = 1 + K_\delta - \lambda$ then

$$\mathbb{E}[V(X_{t+1}) \mid X_t = x, S_{t+1} = h] \leq \lambda V(x) + L, \quad h = 1, \dots, H. \quad (\text{A23})$$

Hence (A11) holds, and the argument above implies the aggregate geometric drift condition (A8).

The common quadratic Lyapunov condition is therefore a convenient sufficient condition for global stability, but it is stronger than necessary. The aggregate drift in (A10) is probability weighted, so (A8) may hold even when no single quadratic Lyapunov function contracts uniformly under every regime, provided the respective regime probabilities are small enough and the other regimes have sufficiently strong contraction.

Implication: Long-run distribution. The preceding stability and communication conditions imply that the process admits a unique invariant probability distribution F^* and is geometrically ergodic in the corresponding V -norm (Meyn and Tweedie 2012). The invariant distribution satisfies

$$F^*(B) = \int \sum_{h=1}^H p_h(x) Q_h(x, B) F^*(dx). \quad (\text{A24})$$

Thus, if $X_t \sim F^*$, endogenous regime selection followed by regime-specific propagation reproduces the same state distribution at $t + 1$.

Condition 4: Moment requirements. The final requirements are claim-specific rather than necessary for the existence of the transition process itself. Expectations involving a measurable function g require $\mathbb{E}|g(X)| < \infty$, second-moment results require $\mathbb{E}|X|^2 < \infty$, and conditional-mean results involving the VAR innovations additionally impose

$$\mathbb{E}[u_{t+1} \mid X_t, S_{t+1}] = 0.$$

APPENDIX B: ENDOGENOUS REGIME SELECTION AND DISTRIBUTIONAL EFFECTS

Proposition 1 (Evolution of the state distribution). *Let*

$$F_t(B) := \Pr(X_t \in B), \quad B \in \mathcal{B}(\mathcal{X}), \quad (\text{B25})$$

denote the distribution of the state at time t . Since

$$P(x, B) = \Pr(X_{t+1} \in B \mid X_t = x),$$

the law of total probability implies that the distribution evolves according to

$$\begin{aligned} F_{t+1}(B) &= \Pr(X_{t+1} \in B) \\ &= \int_{\mathcal{X}} \Pr(X_{t+1} \in B \mid X_t = x) F_t(dx) = \int_{\mathcal{X}} P(x, B) F_t(dx). \end{aligned}$$

Using the aggregate mixture kernel

$$P(x, B) = \sum_{h=1}^H p_h(x) Q_h(x, B),$$

this becomes

$$F_{t+1}(B) = \int_{\mathcal{X}} \left[\sum_{h=1}^H p_h(x) Q_h(x, B) \right] F_t(dx) = \sum_{h=1}^H \int_{\mathcal{X}} p_h(x) Q_h(x, B) F_t(dx). \quad (\text{B26})$$

Thus, the next-period distribution is obtained by averaging the regime-specific transition laws Q_h over both the state-dependent regime probabilities $p_h(x)$ and the current distribution of economic conditions F_t .

Proof: The first equality follows from the law of total probability by conditioning on the current state $X_t = x$. Substituting the mixture representation of the aggregate transition kernel and using linearity of the integral yields (B26). ■

Proposition 1 shows how the distribution of economic conditions evolves over time in the state-dependent mixture. At any date, the current distribution F_t places probability mass over possible states $X_t = x$. For each such state, the gate $p_h(x)$ determines how likely each regime is to govern the next transition, while the regime-specific kernel $Q_h(x, \cdot)$ determines where the economy can move conditional on that regime. The next-period distribution F_{t+1} therefore results from averaging these state- and regime-specific transition laws over all currently possible economic conditions. In this sense, the evolution of the economy depends jointly on where probability mass is located today, how those states affect regime selection, and how each selected regime propagates the state forward.

Proposition 2 (Distribution-dependent regime incidence). *Let F_t denote the distribution of X_t . Then the unconditional probability that regime h governs the transition from t to $t + 1$ is*

$$\pi_{h,t+1} \equiv \Pr(S_{t+1} = h) = \int_{\mathcal{X}} p_h(x) F_t(dx) = \mathbb{E}_{F_t} [p_h(X_t)]. \quad (\text{B27})$$

Thus, even for fixed regime-selection functions $p_h(\cdot)$, regime incidence depends on the distribution

of prevailing economic conditions.

If $\pi_{h,t+1} > 0$, the distribution of the current state conditional on the subsequent selection of regime h is

$$F_{h,t}^-(B) \equiv \Pr(X_t \in B \mid S_{t+1} = h) = \frac{\int_B p_h(x) F_t(dx)}{\pi_{h,t+1}}, \quad B \in \mathcal{B}(\mathcal{X}). \quad (\text{B28})$$

Equivalently,

$$F_{h,t}^-(dx) = \frac{p_h(x)}{\pi_{h,t+1}} F_t(dx). \quad (\text{B29})$$

Consequently, for any F_t -integrable function g ,

$$\mathbb{E}[g(X_t) \mid S_{t+1} = h] - \mathbb{E}_{F_t}[g(X_t)] = \frac{\text{Cov}_{F_t}(g(X_t), p_h(X_t))}{\pi_{h,t+1}}. \quad (\text{B30})$$

Proof: By the law of total probability,

$$\Pr(S_{t+1} = h) = \int_{\mathcal{X}} \Pr(S_{t+1} = h \mid X_t = x) F_t(dx) = \int_{\mathcal{X}} p_h(x) F_t(dx).$$

Hence,

$$\pi_{h,t+1} = \Pr(S_{t+1} = h) = \int_{\mathcal{X}} p_h(x) F_t(dx) = \mathbb{E}_{F_t}[p_h(X_t)],$$

which establishes (B27).

Now suppose $\pi_{h,t+1} > 0$. For any measurable set $B \in \mathcal{B}(\mathcal{X})$, Bayes' rule gives

$$F_{h,t}^-(B) = \Pr(X_t \in B \mid S_{t+1} = h) = \frac{\Pr(X_t \in B, S_{t+1} = h)}{\Pr(S_{t+1} = h)}.$$

The joint probability in the numerator can be written as

$$\Pr(X_t \in B, S_{t+1} = h) = \int_B \Pr(S_{t+1} = h \mid X_t = x) F_t(dx) = \int_B p_h(x) F_t(dx).$$

Therefore,

$$F_{h,t}^-(B) = \frac{\int_B p_h(x) F_t(dx)}{\pi_{h,t+1}}, \quad (\text{B31})$$

which establishes (B28). Since $F_{h,t}^-$ is absolutely continuous with respect to F_t , its Radon–Nikodym derivative is

$$\frac{dF_{h,t}^-}{dF_t}(x) = \frac{p_h(x)}{\pi_{h,t+1}}. \quad (\text{B32})$$

Hence, for any F_t -integrable function g ,

$$\int_{\mathcal{X}} g(x) F_{h,t}^-(dx) = \int_{\mathcal{X}} g(x) \frac{p_h(x)}{\pi_{h,t+1}} F_t(dx).$$

with

$$F_{h,t}^-(dx) = \frac{p_h(x)}{\pi_{h,t+1}} F_t(dx).$$

Using this selected-state distribution, for any F_t -integrable function g ,

$$\mathbb{E}[g(X_t) \mid S_{t+1} = h] = \int_{\mathcal{X}} g(x) F_{h,t}^-(dx) = \frac{1}{\pi_{h,t+1}} \int_{\mathcal{X}} g(x) p_h(x) F_t(dx) = \frac{\mathbb{E}_{F_t}[g(X_t) p_h(X_t)]}{\pi_{h,t+1}}.$$

Subtracting $\mathbb{E}_{F_t}[g(X_t)]$ from both sides and using

$$\pi_{h,t+1} = \mathbb{E}_{F_t}[p_h(X_t)]$$

gives

$$\begin{aligned} \mathbb{E}[g(X_t) \mid S_{t+1} = h] - \mathbb{E}_{F_t}[g(X_t)] &= \frac{\mathbb{E}_{F_t}[g(X_t) p_h(X_t)] - \mathbb{E}_{F_t}[g(X_t)] \mathbb{E}_{F_t}[p_h(X_t)]}{\pi_{h,t+1}} \\ &= \frac{\text{Cov}_{F_t}(g(X_t), p_h(X_t))}{\pi_{h,t+1}}, \end{aligned} \tag{B33}$$

which proves (B30). ■

Proposition 2 shows that regime incidence is itself distribution dependent. Even when the gate $p_h(x)$ is fixed, the unconditional probability of observing regime h depends on how probability mass is distributed across economic states through F_t . Moreover, conditioning on the event that regime h is selected re-weights the current state distribution toward states for which $p_h(x)$ is relatively large. Thus, regimes are not selected independently of prevailing economic conditions. Rather, the states observed before entering a regime are systematically tilted toward regions of the state space in which that regime is more likely.

The covariance identity makes this selection effect explicit. For any feature $g(X_t)$, the difference between its average value conditional on selecting regime h and its unconditional average is proportional to the covariance between $g(X_t)$ and the gate probability $p_h(X_t)$. In particular, setting $g(X_t) = X_t$ shows that the average state preceding regime h differs from the unconditional average precisely to the extent that the components of the state covary with the probability of selecting that regime.

An immediate implication is that identical gate parameters need not imply identical regime frequencies. If F and $\tilde{F}$ are two different distributions of the state under otherwise identical model parameters, then

$$\pi_h(F) - \pi_h(\tilde{F}) = \int p_h(x) \{F - \tilde{F}\}(dx). \tag{B34}$$

Economically, this implies that two economies with identical model parameters but different current or initial state distributions can face different unconditional regime probabilities at a given horizon. Under a common ergodic transition kernel, such differences need not persist in

the long run because the state distributions converge toward the same invariant distribution.

Proposition 3 (State-dependent propagation and regime selection). *Let*

$$m_h(x) := \mathbb{E}[X_{t+1} \mid X_t = x, S_{t+1} = h] = b_h + F_h x, \quad h = 1, \dots, H,$$

denote the regime-specific conditional mean of the next-period state. Then the aggregate conditional mean is

$$m(x) := \mathbb{E}[X_{t+1} \mid X_t = x] = \sum_{h=1}^H p_h(x) m_h(x). \quad (\text{B35})$$

If the functions $p_h(x)$ are differentiable, then

$$D_x m(x) = \sum_{h=1}^H p_h(x) F_h + \sum_{h=1}^H m_h(x) [\nabla_x p_h(x)]'. \quad (\text{B36})$$

The first term captures propagation through the regime-specific state-transition laws, whereas the second captures the effect of changes in the current state on the probabilities with which those laws are selected.

Proof: By the law of iterated expectations,

$$\begin{aligned} m(x) &= \mathbb{E}[X_{t+1} \mid X_t = x] = \sum_{h=1}^H \Pr(S_{t+1} = h \mid X_t = x) \mathbb{E}[X_{t+1} \mid X_t = x, S_{t+1} = h] \\ &= \sum_{h=1}^H p_h(x) m_h(x), \end{aligned}$$

which establishes (B35).

Differentiating with respect to x and applying the product rule gives

$$\begin{aligned} D_x m(x) &= \sum_{h=1}^H D_x [p_h(x) m_h(x)] \\ &= \sum_{h=1}^H [p_h(x) D_x m_h(x) + m_h(x) [\nabla_x p_h(x)]']. \end{aligned}$$

Since $m_h(x) = b_h + F_h x$, we have

$$D_x m_h(x) = F_h.$$

Substituting this into the preceding expression yields

$$D_x m(x) = \sum_{h=1}^H p_h(x) F_h + \sum_{h=1}^H m_h(x) [\nabla_x p_h(x)]',$$

which proves (B36).

■

Equation (B36) shows that a change in the economic state affects future outcomes through two channels. First, it propagates through the regime-specific VAR coefficients. Second, because the same state enters the regime gate, it changes the probabilities with which the different laws of motion are selected.

For the binary logistic case,

$$p_1(x) = v(x) = \Lambda(\alpha + \beta'x), \quad p_2(x) = 1 - v(x), \quad (\text{B37})$$

with

$$\nabla_x v(x) = v(x)\{1 - v(x)\}\beta. \quad (\text{B38})$$

Hence

$$D_x m(x) = v(x)F_1 + \{1 - v(x)\}F_2 + v(x)\{1 - v(x)\}[m_1(x) - m_2(x)]\beta'. \quad (\text{B39})$$

The final term is absent in a fixed-weight mixture. Hence, in the LSBP-VAR, the effect of a change in economic conditions differs as it depends not only on how each regime propagates that change, but also on how the change shifts the likelihood of entering the different regimes.

Proposition 4 (Structural shocks and future regime incidence). *Consider a one-time structural shock δ applied to the innovation governing the transition from t to $t + 1$. Let F_{t+k}^δ denote the distribution of X_{t+k} under the shocked path and F_{t+k}^0 the corresponding distribution under the baseline path, with the same initial distribution F_t and the same state-dependent regime-selection functions $p_h(\cdot)$.*

Under the timing

$$X_t \xrightarrow{p_h(X_t)} S_{t+1} \xrightarrow[X_t, \varepsilon_{t+1} + \delta_{t+1}]{\text{regime-specific propagation}} Y_{t+1} \longrightarrow X_{t+1} \xrightarrow{p_h(X_{t+1})} S_{t+2}$$

the shock does not affect the regime governing the contemporaneous transition, since

$$\Pr_{\delta}(S_{t+1} = h) = \Pr_0(S_{t+1} = h) = \int_{\mathcal{X}} p_h(x) F_t(dx). \quad (\text{B40})$$

The shock can, however, affect subsequent regime incidence through its effect on the distribution of the state. For any $k \geq 1$,

$$\pi_{h,t+k+1}^\delta := \Pr_{\delta}(S_{t+k+1} = h) = \int_{\mathcal{X}} p_h(x) F_{t+k}^\delta(dx) \quad (\text{B41})$$

$$\pi_{h,t+k+1}^0 := \Pr_0(S_{t+k+1} = h) = \int_{\mathcal{X}} p_h(x) F_{t+k}^0(dx). \quad (\text{B42})$$

Hence, the effect of the shock on future regime incidence is

$$\Delta\pi_{h,t+k+1} := \pi_{h,t+k+1}^\delta - \pi_{h,t+k+1}^0 = \int_{\mathcal{X}} p_h(x) [F_{t+k}^\delta(dx) - F_{t+k}^0(dx)]. \quad (\text{B43})$$

Thus, even though the regime-selection function itself is unchanged, a structural shock can alter future regime probabilities by changing the distribution of the economic states on which those probabilities depend.

Proof: Because S_{t+1} is selected conditional on X_t , and the shock δ enters only after this regime has been selected, the shocked and baseline paths have the same regime probabilities for the transition from t to $t + 1$. Using Proposition 2,

$$\Pr_{\delta}(S_{t+1} = h) = \Pr_0(S_{t+1} = h) = \int_{\mathcal{X}} p_h(x) F_t(dx),$$

which establishes (B40).

The shock affects the realization and distribution of X_{t+1} through the regime-specific transition law. Consequently, in general,

$$F_{t+1}^{\delta} \neq F_{t+1}^0.$$

Since subsequent regime probabilities are evaluated at the updated state, Proposition 2 implies, for every $k \geq 1$,

$$\pi_{h,t+k+1}^a = \int_{\mathcal{X}} p_h(x) F_{t+k}^a(dx), \quad a \in \{0, \delta\}.$$

Subtracting the baseline expression from the shocked expression yields

$$\Delta\pi_{h,t+k+1} = \int_{\mathcal{X}} p_h(x) [F_{t+k}^{\delta}(dx) - F_{t+k}^0(dx)],$$

which proves (B43). ■

Proposition 4 shows how structural shocks feed back into future regime selection. In the baseline state-dependent specification, the regime governing the transition from t to $t + 1$ is selected using the predetermined state X_t . A structural shock applied during that transition therefore does not alter the contemporaneous regime probability. Instead, the shock changes Y_{t+1} and hence the updated state X_{t+1} . Because the gate is re-evaluated using this new state, the shock can affect the probabilities of the regimes governing subsequent transitions.

The result

$$\Delta\pi_{h,t+k+1} = \int_{\mathcal{X}} p_h(x) [F_{t+k}^{\delta}(dx) - F_{t+k}^0(dx)]$$

makes this mechanism explicit. The regime-selection function $p_h(\cdot)$ itself is unchanged; what changes is the distribution of states at which that function is evaluated. Structural shocks therefore affect future regime incidence by moving probability mass across regions of the state space with different regime-selection probabilities. In this sense, shocks can generate endogenous changes in regime paths even though they do not directly alter the gate.

Corollary 1 (State dependence of shock-induced regime incidence). *Consider the same one-time policy shock δ applied to two economies, or to the same economy under two different initial state distributions, F and $\tilde{F}$. Let $F_k^{\delta;F}$ and $F_k^{0;F}$ denote the shocked and baseline distributions of X_{t+k} when the initial distribution is F , and define $F_k^{\delta;\tilde{F}}$ and $F_k^{0;\tilde{F}}$ analogously.*

The effect of the policy shock on the probability of selecting regime h at horizon $k + 1$ is

$$\Delta\pi_{h,k+1}(F) = \int_{\mathcal{X}} p_h(x) \left[F_k^{\delta;F}(dx) - F_k^{0;F}(dx) \right]. \quad (\text{B44})$$

Hence, for two different initial distributions,

$$\begin{aligned} \Delta\pi_{h,k+1}(F) - \Delta\pi_{h,k+1}(\tilde{F}) &= \int_{\mathcal{X}} p_h(x) \left[F_k^{\delta;F}(dx) - F_k^{0;F}(dx) \right] \\ &\quad - \int_{\mathcal{X}} p_h(x) \left[F_k^{\delta;\tilde{F}}(dx) - F_k^{0;\tilde{F}}(dx) \right]. \end{aligned} \quad (\text{B45})$$

Therefore, even with identical structural parameters, an identical regime-selection function, and the same policy shock, policy transmission can differ across initial economic conditions. The difference does not arise because the shock itself changes, but because the same intervention operates on different distributions of the state. Thus, even with identical structural parameters, an identical regime-selection function, and the same policy shock, the effect of the shock on future regime incidence can differ across initial economic conditions because the intervention generates different future state distributions.

Proof: Equation (B44) follows directly from the result that, under a fixed gate $p_h(\cdot)$, a shock affects future regime incidence only through the distribution of the future state. Applying this result separately under initial distributions F and $\tilde{F}$ and taking the difference yields (B45). ■

Corollary 1 formalizes the idea that the same structural shock can have different aggregate effects depending on the economic conditions from which it is applied. Even when the structural parameters, regime-selection function, and shock are identical, different initial state distributions can generate different subsequent distributions of the state. Because future regime probabilities are evaluated at those evolving states, the resulting regime-incidence responses can also differ.

The source of the difference is therefore not a change in the shock itself, but a change in the distribution on which the shock operates. In a state-dependent system, identical interventions can propagate differently when they are applied in different economic environments, since both the regime-specific dynamics and the subsequent regime probabilities depend on the state.

Proposition 5 (Decomposition of shock-induced scenario probabilities). *Let $A \in \mathcal{B}(\mathcal{X})$ denote an economically meaningful region of the state space at horizon $t + k$, such as a recession, stagflation, or policy-target-breach region. For $a \in 0, \delta$, where $a = 0$ denotes the baseline path*

and $a = \delta$ the shocked path, define

$$\pi_{h,t+k}^a := \Pr_a(S_{t+k} = h) \quad (\text{B46})$$

Under the finite-logit LSBP specification, $p_h(x) > 0$ for every x and h , so $\pi_{h,t+k}^a > 0$ and the conditional probabilities below are well defined. Let

$$s_{h,t+k}^a(A) := \Pr_a(X_{t+k} \in A \mid S_{t+k} = h). \quad (\text{B47})$$

Then the probability of the scenario at horizon $t + k$ satisfies

$$\Pr_a(X_{t+k} \in A) = \sum_{h=1}^H \pi_{h,t+k}^a s_{h,t+k}^a(A). \quad (\text{B48})$$

Consequently, the effect of the shock on the scenario probability,

$$\Delta P_{t+k}(A) := \Pr_{\delta}(X_{t+k} \in A) - \Pr_0(X_{t+k} \in A), \quad (\text{B49})$$

admits the exact decomposition

$$\Delta P_{t+k}(A) = \sum_{h=1}^H \left(\pi_{h,t+k}^{\delta} - \pi_{h,t+k}^0 \right) s_{h,t+k}^0(A) + \sum_{h=1}^H \pi_{h,t+k}^{\delta} \left[s_{h,t+k}^{\delta}(A) - s_{h,t+k}^0(A) \right]. \quad (\text{B50})$$

This is an exact accounting decomposition using the baseline regime-specific scenario probabilities as the reference. The first term measures the component associated with changes in regime incidence while holding the within-regime scenario probabilities at their baseline values. The second term measures changes in the within-regime scenario probabilities, weighted by the shocked regime incidences. The decomposition is reference dependent and an equally valid decomposition can be obtained using the shocked within-regime probabilities as the reference.

Proof: For any $a \in 0, \delta$, the events

$$\{S_{t+k} = 1\}, \dots, \{S_{t+k} = H\}$$

form a partition of the sample space. Hence, by the law of total probability,

$$\begin{aligned} \Pr_a(X_{t+k} \in A) &= \sum_{h=1}^H \Pr_a(X_{t+k} \in A, S_{t+k} = h) \\ &= \sum_{h=1}^H \Pr_a(S_{t+k} = h) \Pr_a(X_{t+k} \in A \mid S_{t+k} = h) \\ &= \sum_{h=1}^H \pi_{h,t+k}^a s_{h,t+k}^a(A), \end{aligned} \quad (\text{B51})$$

which establishes (B48).

Applying this identity to the shocked and baseline paths gives

$$\Delta P_{t+k}(A) = \sum_{h=1}^H \pi_{h,t+k}^{\delta} s_{h,t+k}^{\delta}(A) - \sum_{h=1}^H \pi_{h,t+k}^0 s_{h,t+k}^0(A). \quad (\text{B52})$$

Add and subtract

$$\sum_{h=1}^H \pi_{h,t+k}^{\delta} s_{h,t+k}^0(A).$$

Then

$$\begin{aligned} \Delta P_{t+k}(A) &= \sum_{h=1}^H \pi_{h,t+k}^{\delta} s_{h,t+k}^{\delta}(A) - \sum_{h=1}^H \pi_{h,t+k}^{\delta} s_{h,t+k}^0(A) \\ &\quad + \sum_{h=1}^H \pi_{h,t+k}^{\delta} s_{h,t+k}^0(A) - \sum_{h=1}^H \pi_{h,t+k}^0 s_{h,t+k}^0(A). \end{aligned} \quad (\text{B53})$$

Collecting terms yields

$$\Delta P_{t+k}(A) = \sum_{h=1}^H \pi_{h,t+k}^{\delta} [s_{h,t+k}^{\delta}(A) - s_{h,t+k}^0(A)] + \sum_{h=1}^H (\pi_{h,t+k}^{\delta} - \pi_{h,t+k}^0) s_{h,t+k}^0(A), \quad (\text{B54})$$

which is (B50). ■

The preceding results provide the probabilistic foundation for the scenario analysis used in the main text. A scenario is defined as

$$\mathcal{S} = (t_0, \Delta, \mathcal{A}),$$

where t_0 determines the initial conditioning information, Δ specifies a structural shock path, and $\mathcal{A}$ denotes an economically meaningful event defined on the subsequent predictive path. The corresponding scenario probability is

$$\Pr_{\Delta}(\mathcal{A} \mid \omega_{t_0}), \quad (\text{B55})$$

where ω_{t_0} denotes the information set used to initialize the predictive simulation.

In the state-dependent mixture, this probability is generated by repeated application of the same transition mechanism characterized above. At each horizon, the current state determines the regime probabilities, the selected regime determines the relevant law of motion, and the resulting observation updates the state entering the next-period gate. Structural shocks affect scenario probabilities by altering the predictive distribution of future states and, through those states, subsequent regime probabilities and regime paths.

The decomposition

$$\Delta P_{t+k}(A) = \sum_{h=1}^H \left(\pi_{h,t+k}^{\delta} - \pi_{h,t+k}^0 \right) s_{h,t+k}^0(A) + \sum_{h=1}^H \pi_{h,t+k}^{\delta} \left[s_{h,t+k}^{\delta}(A) - s_{h,t+k}^0(A) \right]$$

makes the two channels explicit. The first term captures the contribution of changes in regime incidence, while the second captures changes in the probability of the scenario within regimes. Hence, a shock can raise or lower the probability of an event such as recession, stagflation, or a target breach either because it changes which regimes are more likely to occur, because it changes outcomes within those regimes, or through both channels simultaneously. Thus, the scenario analysis is a functional of the shock-induced predictive path distribution rather than a separate forecasting device.

Proposition 6 (Permanent policy changes and the invariant distribution). *Suppose the baseline economy evolves according to the aggregate transition kernel*

$$P^0(x, B) = \sum_{h=1}^H p_h(x) Q_h^0(x, B), \quad (\text{B56})$$

and admits a unique invariant distribution $F^{\star,0}$ satisfying

$$F^{\star,0}(B) = \int_{\mathcal{X}} P^0(x, B) F^{\star,0}(dx). \quad (\text{B57})$$

Consider a permanent change in the policy rule that changes the regime-specific transition laws from Q_h^0 to Q_h^1 . If the regime-selection functions $p_h(\cdot)$ remain unchanged, the post-policy aggregate transition kernel is

$$P^1(x, B) = \sum_{h=1}^H p_h(x) Q_h^1(x, B). \quad (\text{B58})$$

If P^1 also admits a unique invariant distribution $F^{\star,1}$, then

$$F^{\star,1}(B) = \int_{\mathcal{X}} P^1(x, B) F^{\star,1}(dx). \quad (\text{B59})$$

Because the policy change alters the transition kernel, the fixed-point equation determining the invariant distribution changes. The post-policy invariant distribution therefore need not coincide with the baseline invariant distribution:

$$F^{\star,1} \text{ need not equal } F^{\star,0}.$$

Consequently, even when the gate $p_h(\cdot)$ remains unchanged, long-run regime incidence can change through changes in the invariant distribution of the state according to

$$\pi_h^{\star,1} - \pi_h^{\star,0} = \int_{\mathcal{X}} p_h(x) [F^{\star,1}(dx) - F^{\star,0}(dx)], \quad (\text{B60})$$

where

$$\pi_h^{\star,j} := \int_{\mathcal{X}} p_h(x) F^{\star,j}(dx), \quad j \in \{0, 1\}. \quad (\text{B61})$$

Thus, a permanent policy-rule change can alter long-run regime frequencies indirectly by changing the invariant distribution of the state.

Proof: Under policy regime $j \in \{0, 1\}$, the invariant distribution is defined by

$$F^{\star,j}(B) = \int_{\mathcal{X}} P^j(x, B) F^{\star,j}(dx).$$

A permanent change from Q_h^0 to Q_h^1 changes the aggregate transition kernel from P^0 to P^1 . Since the fixed-point equation determining the invariant distribution therefore changes, the associated invariant distributions need not coincide.

For each policy regime j , Proposition 2 applied to the invariant distribution gives

$$\pi_h^{\star,j} = \int_{\mathcal{X}} p_h(x) F^{\star,j}(dx).$$

Subtracting the baseline expression from the post-policy expression yields

$$\pi_h^{\star,1} - \pi_h^{\star,0} = \int_{\mathcal{X}} p_h(x) [F^{\star,1}(dx) - F^{\star,0}(dx)],$$

which establishes (B60). ■

This result highlights the distinction between temporary shocks and permanent changes in the policy rule. A one-time structural shock moves the economy through a fixed transition system. Under the maintained ergodicity conditions, the shocked and baseline predictive distributions converge back toward the same invariant distribution. By contrast, a permanent policy change modifies the transition kernel itself and therefore changes the fixed-point problem defining the invariant distribution. The resulting invariant distribution may differ from the baseline invariant distribution, in which case long-run regime incidence may also differ even though the regime-selection functions $p_h(\cdot)$ remain unchanged.

APPENDIX C: LOCAL NONLINEARITIES IN STATE-DEPENDENT REGIME SELECTION

The state-dependent gate introduces nonlinearities even when the index entering each logistic function is linear in the state. These nonlinearities determine how strongly changes in economic conditions alter regime probabilities and, through the selection channel described above, how strongly they affect aggregate propagation.

To make this mechanism transparent, first consider the binary logistic case

$$p(x) = \Lambda(\eta(x)), \quad \eta(x) = \alpha + \beta'x, \quad \Lambda(u) = \frac{1}{1 + \exp(-u)}. \quad (\text{C62})$$

Proposition 7 (Local sensitivity of the logistic gate). *Let $p(x)$ be defined by (C62). Then, for each component x_j of the state,*

$$\frac{\partial p(x)}{\partial x_j} = \beta_j p(x) \{1 - p(x)\}. \quad (\text{C63})$$

Hence,

$$\left| \frac{\partial p(x)}{\partial x_j} \right| \leq \frac{|\beta_j|}{4}, \quad (\text{C64})$$

with equality when

$$p(x) = \frac{1}{2}, \quad \text{equivalently} \quad \eta(x) = 0. \quad (\text{C65})$$

Moreover, for any two components x_j and x_k ,

$$\frac{\partial^2 p(x)}{\partial x_j \partial x_k} = \beta_j \beta_k p(x) \{1 - p(x)\} \{1 - 2p(x)\}. \quad (\text{C66})$$

Finally, for any finite change Δx in the state,

$$p(x + \Delta x) - p(x) = \Lambda(\eta(x) + \beta' \Delta x) - \Lambda(\eta(x)). \quad (\text{C67})$$

Thus, the effect of a given state change depends both on the size and direction of Δx and on the initial value of the full state through $\eta(x)$.

Proof: Because

$$\Lambda'(u) = \Lambda(u) \{1 - \Lambda(u)\},$$

the chain rule gives

$$\nabla_x p(x) = p(x) \{1 - p(x)\} \beta,$$

which establishes (C63). Since

$$0 < p(x) < 1$$

and

$$p(x) \{1 - p(x)\} \leq \frac{1}{4},$$

with equality at $p(x) = 1/2$, (C64) follows.

Differentiating (C63) with respect to x_k gives

$$\frac{\partial^2 p(x)}{\partial x_j \partial x_k} = \beta_j \frac{\partial p(x)}{\partial x_k} \{1 - 2p(x)\},$$

and substitution of

$$\frac{\partial p(x)}{\partial x_k} = \beta_k p(x) \{1 - p(x)\}$$

yields (C66). Equation (C67) follows directly from

$$\eta(x + \Delta x) = \eta(x) + \beta' \Delta x. \quad \blacksquare$$

Proposition 7 gives a direct interpretation of the nonlinear state dependence generated by the gate. The sign of the marginal effect of x_j is determined by the sign of β_j , but its magnitude depends on the location of the economy in the state space. In particular,

$$p(x) \{1 - p(x)\}$$

is largest when $p(x) = 1/2$ and approaches zero as $p(x)$ approaches either zero or one. Changes in economic conditions therefore have their largest local effect on regime probabilities when the economy is close to the switching boundary

$$\alpha + \beta' x = 0.$$

By contrast, when the economy is deep inside a region for which one regime is already overwhelmingly likely, further movements in the same region produce relatively small changes in regime probabilities.

This saturation property implies that an identical change in a particular state variable can have very different effects depending on prevailing economic conditions. Although the coefficient β_j is constant, the marginal effect

$$\beta_j p(x) \{1 - p(x)\}$$

depends on all variables entering the gate because they jointly determine $p(x)$. Thus, for example, the effect of a change in one economic variable on regime incidence can be large under one configuration of the remaining state variables and small under another. This interaction arises on the probability scale even without explicit interaction terms in the linear index.

The finite-change expression in (C67) provides the corresponding global interpretation. A movement Δx affects the gate through the index change

$$\beta' \Delta x,$$

but the resulting change in probability also depends on the starting index $\eta(x)$. Consequently, two otherwise identical changes in economic conditions need not generate identical changes in regime probabilities when they occur from different initial states.

The same mechanism also sharpens the propagation decomposition in Proposition 3. In the binary case, the selection component of the derivative of the aggregate conditional mean is

$$p(x)\{1 - p(x)\} [m_1(x) - m_2(x)] \beta'. \quad (\text{C68})$$

Its magnitude therefore depends on two distinct features of the model. First, the two regimes must imply sufficiently different conditional outcomes through

$$m_1(x) - m_2(x).$$

Second, the current state must lie in a region in which regime probabilities are sensitive to changes in economic conditions through

$$p(x)\{1 - p(x)\}.$$

The endogenous-selection component of propagation is therefore weak either when the regimes imply similar conditional dynamics or when the economy is deep inside one regime, and it can be substantially stronger when the regimes differ and the economy lies close to the switching boundary.

Extension to the logistic stick-breaking gate. The same logic extends to the multi-regime LSBP specification. Suppose

$$v_h(x) = \Lambda(\alpha_h + \beta'_h x), \quad h = 1, \dots, H-1, \quad (\text{C69})$$

and recall that

$$p_h(x) = v_h(x) \prod_{\ell < h} \{1 - v_\ell(x)\}, \quad h < H, \quad (\text{C70})$$

while

$$p_H(x) = \prod_{\ell < H} \{1 - v_\ell(x)\}. \quad (\text{C71})$$

For $h < H$, differentiation gives

$$\nabla_x p_h(x) = p_h(x) \left[\{1 - v_h(x)\} \beta_h - \sum_{\ell < h} v_\ell(x) \beta_\ell \right], \quad (\text{C72})$$

whereas for the terminal regime,

$$\nabla_x p_H(x) = -p_H(x) \sum_{\ell < H} v_\ell(x) \beta_\ell. \quad (\text{C73})$$

To see (C72), note that for $h < H$,

$$\log p_h(x) = \log v_h(x) + \sum_{\ell < h} \log \{1 - v_\ell(x)\}.$$

Since

$$\nabla_x \log v_h(x) = \{1 - v_h(x)\} \beta_h$$

and

$$\nabla_x \log \{1 - v_\ell(x)\} = -v_\ell(x) \beta_\ell,$$

multiplying the resulting derivative of $\log p_h(x)$ by $p_h(x)$ yields (C72). Equation (C73) follows analogously.

Equations (C72)–(C73) show that state dependence in the multi-regime model operates through both the regime's own break and the preceding breaks in the stick-breaking sequence. A change in the state can therefore alter the probability assigned to regime h directly through $v_h(x)$, but also indirectly by changing how much probability remains after the earlier breaks. In particular, later regimes inherit the state dependence of the preceding gates.

Each individual logistic break retains the same saturation property as in the binary case:

$$\nabla_x v_h(x) = v_h(x) 1 - v_h(x) \beta_h. \quad (\text{C74})$$

Its local sensitivity is therefore strongest near its own switching boundary,

$$\alpha_h + \beta'_h x = 0,$$

and becomes small when $v_h(x)$ is close to either zero or one. The LSBP consequently generates a sequence of state-dependent transition regions rather than a set of fixed regime probabilities.

These properties give the gate a direct economic interpretation. Regime selection is particularly responsive to changes in economic conditions in regions of the state space where the economy is close to moving between competing laws of motion. The same change can have a much smaller effect when it occurs in a state in which regime assignment is already nearly certain. Structural shocks can therefore have strongly state-dependent effects on subsequent regime incidence: shocks that move the economy through a transition region can substantially alter the distribution of future regime paths, whereas comparable shocks occurring deep within

a regime may produce only limited changes in regime selection.

APPENDIX D: SHOCK TRANSMISSION & HETEROSKEDASTICITY STRUCTURE

This section provides additional evidence on the regimes' volatility structure as well as diagnostics on the sign and narrative restrictions.

The top panel of [Table D1](#) reports a shock-by-shock measure of heteroskedasticity strength: for each structural shock i , the posterior probability that its regime-2-to-regime-1 volatility ratio ω_i differs from one by more than a factor of $e^{0.5} \approx 1.61$. High values indicate that the data identify a material change in shock-specific volatility between the two regimes, which is a necessary condition for heteroskedasticity to deliver identifying content for that shock. In the EA every shock crosses this bar with posterior probability of at least 0.96; in the US the lowest separation probability is 0.69, with the bulk of shocks above 0.74.

Table D1: Shock label diagnostics

HETEROSKEDASTICITY DIAGNOSTICS							
ECONOMY	DOM. SUPPLY	DOM. DEMAND	MONETARY POLICY	FINANCIAL	GEOPOLITICAL RISK	OIL PRICE	EXCHANGE RATE
EA	0.99	0.99	0.99	1.00	0.99	0.96	0.99
US	0.73	0.69	0.90	0.80	0.80	0.88	0.74
EURO AREA				UNITED STATES			
SIGN DIAGNOSTICS	NARRATIVE DIAGNOSTICS			SIGN DIAGNOSTICS	NARRATIVE DIAGNOSTICS		
median margin	posterior	reference	odds ratio	median margin	posterior	reference	odds ratio
1.34	0.72	0.30	5.97	1.20	0.50	0.24	3.09

Separation is defined shock-by-shock and measures the posterior probability that the volatility ratio of both states differs materially from one, i.e. by at least a factor of $e^{0.5} \approx 1.61$. The sign-restriction diagnostic measures the distance of each restricted impact coefficient from its imposed sign boundary. It is defined coefficient-by-coefficient as the posterior mean divided by the posterior standard deviation after orienting the coefficient in the required direction. Large positive values indicate that the restricted posterior lies comfortably away from the sign boundary, while values near zero indicate greater posterior mass close to the boundary and hence less sharply determined sign orientation. For the narrative restrictions, the posterior compatibility share q_{post} and the compatibility share q_{ref} under the sign-restricted reference distribution used for the diagnostic are reported. Additionally, the posterior-to-reference compatibility odds ratio $OR_{\text{comp}} = \frac{q_{\text{post}}/(1-q_{\text{post}})}{q_{\text{ref}}/(1-q_{\text{ref}})}$ is reported. Values above one indicate that the narrative restrictions are more compatible with draws informed by the data than with draws from the reference distribution. This quantity is used as a descriptive compatibility diagnostic rather than as a formal Bayes factor.

Shock-by-shock separation is, however, only a necessary condition. Full statistical identification of B (up to column sign and permutation) requires the stronger *pairwise* rank condition: the volatility profiles of any two shocks must be non-proportional across regimes. With $H = 2$ and the reference normalisation $\Lambda_1 = I$, this reduces to requiring $\omega_j \neq \omega_k$ for every pair (j, k) . [Table D4](#) and [Table D2](#) report a full pairwise diagnostic that quantifies, for each shock pair, the posterior distance from $\{\omega_j = \omega_k\}$. The two panels are complementary. The shock-by-shock test asks whether heteroskedasticity has any traction *at the level of each shock*, while the pairwise test asks whether heteroskedasticity *distinguishes each pair of shocks from one another*.

Specifically the tables evaluate

$$\delta_{jk}^{(d)} = |\log \omega_j^{(d)} - \log \omega_k^{(d)}|, \quad (\text{D75})$$

for every shock pair $\{j, k\}$ and posterior draw $d = 1, \dots, D$ and report the posterior mean $\bar{\delta}_{jk} = D^{-1} \sum_d \delta_{jk}^{(d)}$ as the headline measure of pairwise separation, the lower posterior per-

centile $\delta_{q05,jk}$ as an indication of separation in the lower tail of the posterior, the fold change $\exp(\bar{\delta}_{jk})$ for economic interpretation, and the posterior probabilities $\Pr(\delta_{jk} < 0.10 \mid \text{data})$ and $\Pr(\delta_{jk} < \ln 2 \mid \text{data})$. These cutoffs are descriptive benchmarks rather than formal identification thresholds. The first captures very small differences in the across-regime volatility ratios (less than approximately $1.10\times$), while the second records the probability that the ratios differ by less than a factor of two. [Table D5](#) and [Table D3](#) additionally report the number of pairs falling below several such descriptive separation benchmarks.

Table D2: Pairwise heteroskedasticity-ID strength via mean log-distance $\bar{\delta}_{jk} - \text{US}$.

Pair	$\bar{\delta}$	δ_{q05}	fold	$\Pr(\delta < 0.10)$	$\Pr(\delta < \ln 2)$
USDS/USDD	0.799	0.113	$2.22\times$	0.043	0.455
USDS/USMP	0.599	0.052	$1.82\times$	0.092	0.636
USDS/Financial	0.562	0.066	$1.75\times$	0.079	0.661
USDS/GPR	0.512	0.067	$1.67\times$	0.080	0.728
USDS/WTI	0.749	0.104	$2.11\times$	0.047	0.504
USDS/FX	0.712	0.088	$2.04\times$	0.059	0.509
USDD/USMP	0.631	0.043	$1.88\times$	0.110	0.625
USDD/Financial	0.807	0.142	$2.24\times$	0.034	0.423
USDD/GPR	0.779	0.131	$2.18\times$	0.040	0.414
USDD/WTI	0.869	0.117	$2.38\times$	0.042	0.443
USDD/FX	0.827	0.120	$2.29\times$	0.042	0.422
USMP/Financial	0.652	0.089	$1.92\times$	0.056	0.572
USMP/GPR	0.659	0.122	$1.93\times$	0.040	0.546
USMP/WTI	0.560	0.049	$1.75\times$	0.104	0.677
USMP/FX	0.746	0.103	$2.11\times$	0.048	0.491
Financial/GPR	0.428	0.044	$1.53\times$	0.126	0.824
Financial/WTI	0.593	0.065	$1.81\times$	0.080	0.633
Financial/FX	0.556	0.060	$1.74\times$	0.083	0.683
GPR/WTI	0.666	0.096	$1.95\times$	0.053	0.525
GPR/FX	0.525	0.051	$1.69\times$	0.101	0.705
WTI/FX	0.618	0.084	$1.86\times$	0.062	0.633

Table D3: Pair counts below identification threshold – US.

Threshold name	τ (in δ)	fold change	# pairs below
very small	0.100	$1.11\times$	0 / 21
moderate	0.500	$1.65\times$	1 / 21
two fold	0.693	$2.00\times$	13 / 21
stringent	1.000	$2.72\times$	21 / 21

Across both economies, no shock pair has a posterior mean log-distance below the smallest descriptive benchmark of 0.10: 0/21 US pairs and 0/21 EA pairs fall below this cutoff, indicating non-negligible separation of the volatility profiles for every pair. At the 0.50 benchmark, corre-

Table D4: Pairwise heteroskedasticity-ID strength via mean log-distance $\bar{\delta}_{jk}$ – EA.

Pair	$\bar{\delta}$	δ_{q05}	fold	$\Pr(\delta < 0.10)$	$\Pr(\delta < \ln 2)$
EADS/EADD	0.522	0.040	1.69×	0.119	0.714
EADS/EAMP	0.628	0.049	1.87×	0.098	0.614
EADS/Financial	0.505	0.041	1.66×	0.125	0.729
EADS/GPR	0.487	0.037	1.63×	0.130	0.746
EADS/WTI	0.555	0.043	1.74×	0.113	0.678
EADS/FX	0.545	0.043	1.73×	0.120	0.689
EADD/EAMP	0.626	0.048	1.87×	0.102	0.606
EADD/Financial	0.518	0.039	1.68×	0.126	0.715
EADD/GPR	0.512	0.040	1.67×	0.120	0.722
EADD/WTI	0.628	0.054	1.87×	0.093	0.610
EADD/FX	0.586	0.046	1.80×	0.104	0.641
EAMP/Financial	0.561	0.043	1.75×	0.108	0.672
EAMP/GPR	0.595	0.052	1.81×	0.099	0.634
EAMP/WTI	0.761	0.067	2.14×	0.075	0.497
EAMP/FX	0.591	0.045	1.81×	0.109	0.648
Financial/GPR	0.474	0.033	1.61×	0.136	0.757
Financial/WTI	0.575	0.047	1.78×	0.104	0.657
Financial/FX	0.496	0.040	1.64×	0.129	0.739
GPR/WTI	0.576	0.046	1.78×	0.112	0.654
GPR/FX	0.531	0.042	1.70×	0.117	0.697
WTI/FX	0.664	0.058	1.94×	0.087	0.571

sponding to a fold difference of approximately 1.65×, only one US pair and three EA pairs fall below the cutoff. At the more demanding two-fold benchmark, 13/21 US pairs and 20/21 EA pairs fall below the cutoff, reflecting the empirical fact that many across-regime variance-ratio differences in these samples lie between roughly 1.5× and 2×. The larger EA count also reflects the comparatively sparse occupancy of the high-volatility regime, which increases posterior uncertainty about the volatility ratios.

None of these findings amount to identification failure, but the heteroskedasticity-based identification is not uniformly sharp either, and the supplementary restrictions therefore carry meaningful weight. The bottom panel of [Table D1](#) shows that the restricted impact coefficients generally lie away from their imposed sign boundaries: the median posterior sign margin is 1.34 for the euro area and 1.20 for the United States. The narrative restrictions yield posterior-to-reference compatibility odds ratios of 5.97 (EA) and 3.09 (US), indicating that the imposed historical restrictions are substantially more compatible with posterior draws informed by the data than with draws from the reference distribution.

In terms of identification, the heteroskedastic structure generically identifies B up to column sign and permutation when the regime-specific volatility profiles are distinct. The diagnostics nevertheless indicate that the strength of this identification is not uniform across shock pairs.


Table D5: Pair counts below identification threshold – EA.

Threshold name	τ (in δ)	fold change	# pairs below
very small	0.100	$1.11\times$	0 / 21
moderate	0.500	$1.65\times$	3 / 21
two fold	0.693	$2.00\times$	20 / 21
stringent	1.000	$2.72\times$	21 / 21

While some volatility profiles are clearly separated, others are sufficiently similar that the corresponding columns of B remain only weakly distinguished in finite samples. The sign restrictions therefore play an important role in labeling the identified columns and in reducing posterior ambiguity for these less sharply separated pairs, while the narrative restriction provides additional discipline at a specific historical episode. The resulting posterior uncertainty is thus materially smaller than under sign restrictions alone, but the degree of statistical identification continues to vary across shock pairs. The pairs for which the supplementary restrictions carry greater weight are economically intuitive. In the United States, the financial and geopolitical-risk shocks exhibit relatively similar volatility profiles, while in the euro area the sparser high-volatility regime broadens posterior uncertainty more generally.


APPENDIX E: IMPULSE RESPONSES

Figure E1: Regime-specific impulse response functions, euro area



Notes: The chart shows regime dependent impulse response functions. The blue IRFs correspond to the normal regime, the yellow IRFs to the volatility regime. Shocks are scaled to regime specific standard deviations.

Figure E2: Regime-specific impulse response functions, United States




Notes: The chart shows regime dependent impulse response functions. The purple IRFs correspond to the normal regime, the red IRFs to the volatility regime. Shocks are scaled to regime specific standard deviations.


APPENDIX F: EVENT WINDOWS

Figure F3: Russia-Ukraine geopolitical events


(a) Identified GPR shocks



(b) GPR



(c) Financial indicators



Notes: The panels show the GPR shocks identified by the model for the Euro area and the US, as well as the evolution of the GPR index and financial market indicators around specific GPR event dates surrounding the Russian invasion of Ukraine.

REFERENCES

- FRANCQ, C. AND J.-M. ZAKOÏAN (2001): “Stationarity of multivariate Markov-switching ARMA models,” *Journal of Econometrics*, 102, 339–364.
- MEYN, S. P. AND R. L. TWEEDIE (2012): *Markov chains and stochastic stability*, Springer Science & Business Media.

Acknowledgements

I would like to thank Mike West and seminar participants at the European Central Bank and participants at ESOBE 2026 for helpful comments and discussions. This paper should not be reported as representing the views of the European Central Bank (ECB). During the preparation of this manuscript, the authors used ChatGPT to assist with the clarity of the exposition and to support the review and presentation of selected mathematical arguments and derivations. All mathematical results and proofs were independently reviewed and verified by the author, who takes full responsibility for the content of the manuscript.

The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Maximilian Schröder

European Central Bank, Frankfurt am Main, Germany; email: maximilian.schroder@ecb.europa.eu

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-8023-4

ISSN 1725-2806

doi:10.2866/3135211

QB-01-26-227-EN-N