GEOPOLITICAL OIL-SUPPLY SCENARIOS AND MACROECONOMIC RISKS: A BAYESIAN TOOLKIT FOR POLICY ANALYSIS

2026

BANCO DE ESPAÑA
Eurosistema

Documentos Ocasionales
N.º 2616

Irma Alonso-Alvarez, Iván Kataryniuk, David López-Salido and Javier J. Pérez

GEOPOLITICAL OIL-SUPPLY SCENARIOS AND MACROECONOMIC RISKS:

A BAYESIAN TOOLKIT FOR POLICY ANALYSIS

GEOPOLITICAL OIL-SUPPLY SCENARIOS AND MACROECONOMIC RISKS: A BAYESIAN TOOLKIT FOR POLICY ANALYSIS (*)

Irma Alonso-Alvarez

BANCO DE ESPAÑA

Iván Kataryniuk

BANCO DE ESPAÑA

David López-Salido

BANCO DE ESPAÑA

Javier J. Pérez

BANCO DE ESPAÑA

(*) We thank Enrique Alberola, Pablo Burriel, José Luis Escrivá, Pablo Garabal, Yannick Kalantzis, Jolan Mohimont, Florens Odendahl, Andrea de Polis, Fabio Rumler and MPC Eurosystem colleagues for helpful comments and suggestions. We also thank Pablo Garabal for superb research assistance. An interactive dashboard implementing the framework is available to authorised users of the Perplexity Computer workspace. The analysis in this paper is self-contained and does not require access to the dashboard. Correspondence to: Javier J. Pérez (javierperez@bde.es).

Documentos Ocasionales. N.º 2616

September 2026

https://doi.org/10.53479/44325

The Occasional Paper Series seeks to disseminate work conducted at the Banco de España, in the performance of its functions, that may be of general interest.

The opinions and analyses in the Occasional Paper Series are the responsibility of the authors and, therefore, do not necessarily coincide with those of the Banco de España or the Eurosystem.

The Banco de España disseminates its main reports and most of its publications via the Internet on its website at: http://www.bde.es.

ISSN: 1696-2230 (online edition)

Abstract

This paper develops a tractable framework that maps geopolitical risk in the Middle East to oil prices and macroeconomic conditions through three linked blocks. Geopolitical uncertainty is modelled as a probability distribution over a small set of mutually exclusive scenarios, whose weights are updated with prediction-market prices, structured expert assessments and high-frequency indicators via Bayesian methods. Each scenario is then translated into a net oil supply shortfall after inventory and reserve buffers, which in turn feeds a reduced-form but economically interpretable Brent price equation grounded in recent oil market research. The same architecture can be combined with standard statistical forecasts to check the internal consistency of scenario mixtures and, in reverse, to infer the scenario configuration implicit in observed oil prices. The framework is designed as an operational tool for central banks and other policy institutions to structure geopolitical scenario analysis, quantify resulting oil price and inflation risks and update these assessments in real time.

Keywords: geopolitical risk, oil supply disruptions, energy security, scenario analysis, Bayesian updating, oil prices, inflation, macroeconomic projections.

JEL classification: C11, C53, E31, E37, F51, Q31, Q43.

Resumen

Este trabajo desarrolla un marco analítico operativo que traduce el riesgo geopolítico en Oriente Medio en trayectorias de precios del petróleo y efectos macroeconómicos mediante tres bloques interrelacionados. En primer lugar, la incertidumbre geopolítica se representa como una distribución de probabilidades sobre un conjunto reducido de escenarios mutuamente excluyentes, cuyas probabilidades se actualizan mediante métodos bayesianos utilizando información procedente de mercados de predicción, evaluaciones estructuradas de expertos e indicadores de alta frecuencia. En segundo lugar, cada escenario se transforma en una trayectoria de perturbación neta de la oferta de petróleo, teniendo en cuenta el papel amortiguador de las existencias comerciales y las reservas estratégicas. Finalmente, estas perturbaciones se incorporan a una ecuación reducida e interpretable económicamente para obtener trayectorias del precio del brent fundamentadas en la literatura reciente sobre mercados petroleros. El marco puede combinarse con previsiones estadísticas convencionales para evaluar la coherencia interna de los distintos escenarios y, en sentido inverso, inferir la configuración geopolítica implícita en los precios observados del petróleo. La herramienta está diseñada para su uso operativo por bancos centrales y otras instituciones de política económica, facilitando el análisis de escenarios geopolíticos, la cuantificación de riesgos para el precio del petróleo y la inflación, y la actualización de dichas evaluaciones en tiempo real.

Palabras clave: riesgo geopolítico, perturbaciones de la oferta de petróleo, seguridad energética, análisis de escenarios, actualización bayesiana, precios del petróleo, inflación, proyecciones macroeconómicas.

Códigos JEL: C11, C53, E31, E37, F51, Q31, Q43.

Contents

1 Introduction

Geopolitical risks are now treated by policymakers and markets as first-order drivers of macro-financial conditions (Ioannou, Pérez, Geeroms, Vansteenkiste, and Weber, 2023; Aiyar, Chen, Ebeke, Garcia-Saltos, Gudmundsson, Ilyina, Kangur, Kunaratskul, Rodriguez, Ruta, Schulze, Soderberg, and Trevino, 2023). A growing literature shows that geopolitical tensions and geoeconomic fragmentation reshape trade, supply chains, and systemic risk, with heterogeneous effects across countries and shock origins (Caldara and Iacoviello, 2022; Bondarenko, Lewis, Rottner, and Schüler, 2024; Alonso-Alvarez, Diakonova, and Pérez, 2025). Within this broader field, work on the geopolitics of energy and energy security highlights how resource endowments, infrastructure, and transit chokepoints shape external vulnerabilities and strategic behaviour (Austvik, 2018; Kuzemko, Blondeel, Bradshaw, Bridge, Faigen, and Fletcher, 2025), while recent analyses of energy-supply security quantify how conditions at transit points map into risks for national supply portfolios (Meza, Ari, Al Sada, and Koç, 2026; Liu, Guo, Shi, Ding, and Chen, 2025).

The recent conflict in the Middle East and the associated threats to the Strait of Hormuz have renewed attention to oil markets as a central transmission channel from geopolitics to macro outcomes (Cordesman and Toukan, 2012). Historical experience shows that oil supply disruptions propagate through physical shortfalls, inventory dynamics, and risk premia, with potentially large effects on activity and inflation (Hamilton, 2009; Kilian and Murphy, 2014; Känzig, 2021; Kilian, Plante, Richter, and Zhou, 2026). The macro impact of a given episode depends not only on the size of the initial shortfall but also on its persistence, the speed and scale of inventory and reserve releases, and how expectations reprice the probability of more extreme scenarios (Caballero, Farhi, and Gourinchas, 2008; Kilian and Murphy, 2014; Känzig, 2021).

Against this background, there is a need for tools that translate conflict narratives into probabilistic scenario distributions and physically interpretable oil-price paths that can be embedded in macro projections.

This paper develops such a framework. The paper is organised in three sequential blocks, as illustrated in Figure 1. The first block assigns probabilities to a small set of mutually exclusive geopolitical scenarios using a structured combination of prior beliefs and incoming information from prediction markets, expert assessments, and real-time indicators, in the spirit of Bayesian learning models (Polymarket, 2026; Cordesman and Toukan, 2012; Baley and Veldkamp, 2021; Caldara and Iacoviello, 2022). The second block maps each scenario into a time path for net oil-supply disruption after inventory and reserve buffers, building on evidence that the macroeconomic effects of oil shocks depend crucially on inventories and the speed of adjustment (Kilian, 2009; Kilian and Murphy, 2014; Baumeister and Kilian, 2016). The third block translates the resulting scarcity profiles into Brent price paths via a reduced-form but economically interpretable price equation grounded in empirical work on oil supply shocks, news, and risk premia (Kilian, 2009; Kilian and Murphy, 2014;

Figure 1: From geopolitical risk to oil prices: a three-block operational framework.

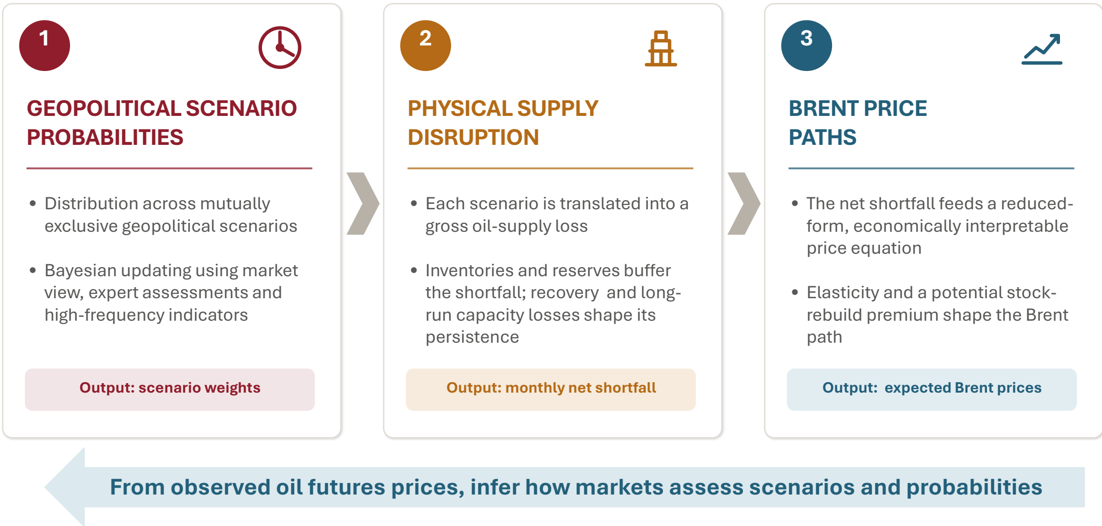

Note: Authors' calculations.

Känzig, 2021; Hamilton, 2009; Kilian et al., 2026). The architecture can be related to statistical reference forecasts to assess whether a given scenario set spans the distribution of macroeconomic outcomes and can also be used in reverse as a diagnostic device, with observed oil prices helping to infer the geopolitical scenario configuration priced by markets. The illustrative application to the June 2026 Eurosystem staff macroeconomic projections exercise (BMPE) (European Central Bank, 2026) shows how this can be done in practice: the framework provides a structural complement to the standard BMPE energy scenarios and a disciplined way to update risk assessments when geopolitical conditions shift within the same round. An operational dashboard implements the framework in real time for authorised users, but the analysis in the paper is self-contained and does not require access to that tool.

Although we illustrate the framework using the 2026 U.S.–Iran conflict and the Strait of Hormuz, the architecture is designed to be portable to other geopolitical settings and energy markets, provided comparable indicators and scenario taxonomies are available.

The framework's main policy contribution is to make the links between geopolitical scenarios, their probabilities, and oil prices transparent and traceable in both directions, from assumptions to price outcomes and from observed prices to market-implied scenarios. Building on this policy contribution, the paper makes three contributions to the literature. First, it adapts mixture representations of scenario analysis in the spirit of Adrian, Giannone, Luciani, and West (2025) to a geopolitical oil-market setting, providing a way to organise scenario-specific distributions and compare them with the statistical reference forecasts used in central banks. Second, it operationalises Bayesian updating along the lines of Baley and Veldkamp (2021) by combining market-implied signals, expert priors, and real-time indicator likelihoods into a tractable posterior over conflict states (Caldara and Iacoviello, 2022; Alonso-Alvarez et al., 2025). Third, it shows how the framework can be inverted and used as a diagnostic tool to infer the geopolitical scenario configuration implied by observed oil prices, linking the analysis to structural scenario forecasts in the sense of Kilian and Murphy (2014) and to recent work on the macroeconomic effects of oil-supply news shocks (Känzig, 2021; Kilian et al., 2026).

The rest of the paper is organised as follows. Section 2 develops the geopolitical scenario design and probability assessment framework. It establishes the conceptual foundations of scenario-based uncertainty quantification, defines the scenario taxonomy and the general parameterisation of scenario dynamics, and sets out the probability block. Sections 3 and 4 describe the supply and price blocks, respectively. Section 5 discusses the reverse-mapping exercise that uses prices as a diagnostic signal about implied geopolitical scenarios. Section 6 presents an illustrative application based on the 2026 U.S.–Iran conflict and its role in the June 2026 Eurosystem projection round. Section 7 concludes. Appendix A briefly describes the operational dashboard implementation, while Appendix B documents calibration details and robustness.

2 Geopolitical scenario design and probability assessment

In this section we first establish the conceptual foundations of the procedure, by interpreting scenarios as components of a mixture distribution over economically relevant outcomes. Next, we define a parsimonious taxonomy of mutually exclusive scenarios and parameterises their dynamic supply implications. Finally, we set out the probability assessment framework, combining market-implied and expert-based priors with a grouped Bayesian update based on real-time indicators.

2.1 Conceptual foundations: scenarios as mixture components

The geopolitical environment generates a *continuous* distribution over possible outcomes: the true state of the world, the supply loss it implies, and the resulting oil price all evolve in a high-dimensional space with no sharp boundaries between “scenarios.” Scenario analysis is a disciplined approach to working with this distribution despite limited information and limited communication bandwidth.

There are two complementary ways to formalize this. The first, which is natural for communication, is to view scenarios as a *partition* of the outcome space: each scenario corresponds to a region of the continuous distribution, and the scenario probability is the probability mass assigned to that region. In this view, choosing scenarios is choosing which regions of the distribution to make explicit.

The second, which is more convenient formally and better suited for synthesis with statistical forecasts, is to view scenarios as *components of a mixture*. Following Adrian et al. (2025), let $p_j(y)$ denote the predictive distribution of outcome y under scenario j, and let $\alpha_j \geq 0$ with $\sum_j \alpha_j = 1$ be the scenario weights. Then the joint scenario forecast is the mixture

$$f(y \mid \boldsymbol{\alpha}) \;=\; \sum_j \alpha_j \, p_j(y). \tag{1}$$

This formulation has two advantages. First, the mixture distribution is itself a proper predictive distribution over y, so it can be compared directly to any statistical reference forecast—such as the output of a Growth-at-Risk model or a quantile regression—using standard measures of distributional fit such as the expected misclassification rate (EMR) introduced by Adrian et al. (2025). Second, in a separate synthesis exercise, the mixture weights α_j can be chosen to maximize concordance of the mixture with the reference distribution, providing a principled way to evaluate whether the proposed scenario set is "complete" enough to span the distribution of outcomes.

An important conceptual point concerns the role of the scenario taxonomy. The scenarios should not be treated as atoms of a discrete probability space but as distributions over the *economic consequences* of geopolitical states. Two scenarios that are very different geopolitically may have very similar economic distributions $p_j(y)$ and therefore carry redundant information in the mixture (1). Conversely, a geopolitically subtle distinction—such as the difference between a formal ceasefire and an uneasy stalemate—may map into very different supply and price distributions if it generates different persistence in disruptions. The economically relevant partition of the outcome space is determined by $p_j(y)$, not by the geopolitical label.

Finally, it is useful to link this framework to the structural scenario forecasts of Kilian et al. (2026). In that literature, scenarios correspond to structural shock configurations—for example, a supply disruption of a given size and persistence—and the question is how to assign probabilities to these configurations consistently with observed data. The present framework adopts the same logic but extends it to a richer geopolitical state space and to a real-time updating setting.

2.2 Scenario taxonomy

A central modelling choice is the definition of the geopolitical state space. This subsection sets out the general principles; the U.S.–Iran/Hormuz specification used in the application is described in Section 6 and Appendix B.

One way to think about scenarios is as outcomes in a multimodal distribution. In standard practice, a "central" scenario can be seen as the mode of a single-mode distribution (Ciccarelli, Pariès, Landau, and Sousa, 2025). When uncertainty increases, the distribution can become multimodal, which justifies selecting a set of scenarios rather than a single benchmark. A second way is to view scenarios as a way to reduce the dimensionality of a continuum of possible outcomes into a small set of economically distinct alternatives.

Therefore, a common practice in scenario design is the reduction to a defined set of mutually exclusive alternatives. Mutual exclusivity is not strictly necessary, but it is helpful in order to align the scenario design to a decision tree and to keep probabilities interpretable for communication.

In the Hormuz application discussed later, we implement this logic using four mutually exclusive scenarios based on possible outcomes from geopolitical analysis (Cordesman and Toukan, 2012; Katzman, 2021)[1]:

- *Diplomatic resolution* (S1): negotiated de-escalation and gradual normalization of flows, implying a mild and transient supply loss, declining market-risk pressure, and relatively fast re-anchoring of prices.
- *Prolonged stalemate* (S2): no full resolution but also no major regional escalation, with persistent moderate loss, residual security concerns, and slow improvement in flows.
- *Partial resumption with persistent disruptions* (S3): recurrent strikes or intermittent blockade episodes despite some reopening, generating a larger and longer-lived supply shortfall, stronger backwardation pressure, and delayed normalization.
- *Major escalation* (S4): broad regional conflict with sustained closure threat and possible spillovers, leading to severe and slowly recovering supply loss, elevated tail-risk pressure, and materially higher macro tail risk (Cordesman and Toukan, 2012; Katzman, 2021).

This four-state specification spans the outcome space from a benign scenario, in which a quick and permanent resolution is achieved, to a severe scenario, in which an intensification of the disruption leads to worse outcomes. It also helps to preserve parsimony. The labels are mainly narrative, as the economic consequences depend less on the formal political label of the conflict than on the persistence and severity of constraints on Gulf exports. For example, an uneasy ceasefire with frequent violations but partial reopening of shipping lanes may look different geopolitically from a formal comprehensive agreement, but in economic terms both correspond to states with high Hormuz throughput and declining market-risk pressure.

The key modelling insight is that scenarios should not be treated as one-off labels but as dynamic profiles. Each scenario s is characterized by three production-side parameters:

[1] The four scenarios were selected after reviewing market commentary, policy-institute assessments and internal market-intelligence material on the range of plausible developments around the Strait of Hormuz. The objective was not to enumerate all possible geopolitical outcomes, but to define a small set of economically distinct states that span the relevant oil-supply consequences while remaining tractable for probability assignment, dashboard implementation and policy communication. In practice, the scenario set was deliberately kept below five alternatives to preserve parsimony and avoid over-fragmenting the probability mass across states with similar economic implications.

- χ^s: the post-anchor shock factor, which scales the intensity of the disruption after the empirically observed initial phase;
- H^s: the recovery horizon, which governs how quickly the gross disruption fades;
- ℓ^s: the residual long-run production term, capturing persistent capacity impairment or, in benign scenarios, an overshooting normalization relative to the disrupted baseline.

For the early phase of a disruption, gross shortfalls can be imposed using observed crude oil production data. From the handover month onward, the gross supply path under scenario s is generated by a scenario-specific recovery profile,

$$g_t^s \;=\; g_{\text{anchor}}\, \chi^s\, F_t^s \;+\; \ell^s\, (1 - F_t^s)\,, \tag{2}$$

where g_{anchor} is the gross shortfall at the handover date and F_t^s is a logistic remaining-disruption factor governed by the recovery horizon H^s. The first term captures the acute disruption as it fades over time, while the second term captures the residual long-run production term towards which the path converges. In benign scenarios, the long-run term can allow for a small normalization overshoot relative to the disrupted baseline. Table B.3 in Appendix B reports the baseline calibration of these parameters for each scenario in the Hormuz application and discusses their robustness.

This parameterisation is essential because two scenarios with similar initial disruptions can have very different price implications if one implies rapid recovery while the other entails prolonged impairment and commercial inventory rebuilding.

2.3 The probability set-up: from priors to posteriors

This subsection describes the general probability block of the framework. In the Hormuz application in Section 6, the same structure is instantiated with episode-specific priors and indicators.

The probability block combines three distinct information sources: market-implied signals, expert assessments, and real-time indicators. This makes the scenario distribution both auditable and updateable over time.

2.4 The prior

The prior over the four scenarios, $\pi_0 = (\pi_0^{S_1}, \ldots, \pi_0^{S_4})$, blends a *market-implied* prior P_M with an *expert* prior P_E using fixed weights:

$$\pi_0(S_i) = \frac{\omega_M P_M(S_i) + \omega_E P_E(S_i)}{\sum_j \left[\omega_M P_M(S_j) + \omega_E P_E(S_j)\right]}, \qquad \omega_M = 0.625, \quad \omega_E = 0.375. \tag{3}$$

The renormalization ensures that the resulting vector is a valid probability distribution over the four scenarios. This is useful because both components are obtained from mappings and scoring steps that need not add up exactly to one before normalization.

Market-implied component. Prediction markets aggregate dispersed information through financial incentives and, under appropriate conditions, provide well-calibrated probability forecasts of geopolitical events.

We use $K = 5$ prediction-market contracts traded on platforms such as Polymarket (Polymarket, 2026), linked to diplomatic agreement, ceasefire durability, Hormuz normalization, and renewed military action in the application.

Because such contracts are binary rather than spanning the four-state space $\{S_1, S_2, S_3, S_4\}$ directly, each contract k is mapped through a fixed *signed loading* $\ell_{k,i}$ on scenario S_i; the corresponding loading matrix is reported in Appendix B.1.

Let $y_k \in [0, 1]$ be the contract-k yes-probability. We form a centred signal $s_k = y_k - \frac{1}{2}$, accumulate a scenario score:

$$z_i = \sum_{k=1}^{K} s_k \, \ell_{k,i}, \tag{4}$$

and map scores to probabilities with a temperature-τ_M softmax:

$$P_M(S_i) = \frac{\exp(z_i/\tau_M)}{\sum_j \exp(z_j/\tau_M)}, \qquad \tau_M = 1. \tag{5}$$

Polarity is carried by the signs of s_k and $\ell_{k,i}$: a contract bullish on de-escalation loads positively on the benign scenario and negatively on the escalation scenario. The market information enters the *prior only*; it is not reused in the indicator likelihood.

Expert component. The expert component is based on structured scenario assessments from policy and geopolitical institutes such as CSIS, Chatham House, and the Atlantic Council (Cordesman and Toukan, 2012). Each source is read into a score vector over the four scenarios; P_E is the source-averaged, renormalised vector, $P_E(S_i) \propto \frac{1}{n} \sum_{r=1}^{n} \sigma_{r,i}$, where $\sigma_{r,i}$ is source r's score for S_i. Leave-one-source-out variation is reported as a robustness diagnostic in the application.

These assessments evolve slowly relative to market signals and provide a slower-moving structural anchor. The resulting prior is the starting point for the indicator-based Bayesian update described below. Prediction markets provide fast-moving, event-specific signals, while expert assessments provide a slower-moving structural anchor.

2.5 The likelihood and Bayesian updating

The updating step follows a grouped Bayesian filter in the spirit of Baley and Veldkamp (2021). Let $Z_t = (Z_t^1, \ldots, Z_t^K)$ denote the vector of K high-frequency indicators observed at time t: in the application, these include tanker traffic through the Strait of Hormuz, Brent front-to-back calendar spreads, tanker insurance premia, option-implied skew on near-dated Brent contracts, geopolitical risk indices (e.g., the Caldara and Iacoviello 2022 index), and proxies for military tempo and diplomatic activity.

The likelihood $\mathcal{L}(Z_t \mid S_i)$ measures how consistent the observed data vector Z_t is with scenario S_i. In practice, each indicator k is first converted into a discrete bucket according to pre-defined thresholds. The observed value $z_{k,t}$ is mapped into $b_{k,t} \in \{\text{low}, \text{medium}, \text{high}\}$. The model then uses an expert-calibrated likelihood matrix $P(b_k \mid S_i)$ to evaluate the probability of observing that bucket under each scenario.

These likelihood matrices are fixed calibration objects inspired by historical analogue episodes such as past tanker wars, targeted attacks on critical infrastructure, and recent regional conflicts. They are not regression-estimated coefficients. What changes in daily operation is the raw indicator value, and therefore the bucket assigned to each indicator, while the conditional likelihoods $P(b_k \mid S_i)$ remain fixed.

The indicator groups receive heterogeneous weights in the grouped-likelihood implementation. This grouping is important because several variables are not conditionally independent. For example, transit calls, Brent calendar spreads, and Brent–WTI spreads all partly measure physical oil-market stress. Treating them as fully independent likelihoods would overstate the amount of information contained in the data. Within each evidence block g, indicator likelihoods are therefore combined as a weighted geometric mean:

$$\log L_g(S_i) = \sum_{k \in g} \tilde{w}_{k,g,t} \log P(b_{k,t} \mid S_i), \tag{6}$$

where the normalised within-block weights are freshness-adjusted:

$$\tilde{w}_{k,g,t} = \frac{\phi_{k,t} w_k}{\sum_{j \in g} \phi_{j,t} w_j},$$

with $\phi_{k,t} \in \{1, 0.5, 0.25\}$ denoting the freshness factor.

Blocks combine through a global likelihood-power discount λ and block weights W_g. In the dashboard implementation, W_g is the mean freshness-adjusted indicator weight of block g,

$$W_g = \frac{1}{|g|} \sum_{k \in g} \phi_{k,t} w_k,$$

and is not renormalised to sum to one across blocks. The posterior is formed in log space:

$$\log \tilde{\pi}_t(S_i) = \log \pi_0(S_i) + \lambda \sum_g W_g \, \log L_{g,t}(S_i), \qquad \lambda = 0.35, \tag{7}$$

$$\pi_t(S_i) = (1-\kappa)\frac{\exp[\log \tilde{\pi}_t(S_i)]}{\sum_j \exp[\log \tilde{\pi}_t(S_j)]} + \frac{\kappa}{N}, \qquad \kappa = 0.04, \quad N = 4, \tag{8}$$

followed by a final renormalization. The discount $\lambda < 1$ guards against over-confidence from noisy real-time signals and residual within-block dependence. In the baseline specification, we set $\lambda = 0.35$. The algorithm nevertheless allows this parameter to vary, enabling an assessment of the sensitivity of the posterior scenario probabilities to alternative discount values.

The live dashboard additionally applies the freshness adjustment. If an indicator has not been updated within its expected refresh window, its weight is reduced rather than being treated as fully current or discarded. This is particularly important for variables such as tanker-insurance premia, parallel-market exchange rates, or military-positioning proxies, which may be economically relevant but difficult to update automatically from public sources.

2.6 Properties of the updating scheme

This probabilistic structure is attractive for at least three reasons. First, it recognises that markets and experts do not answer exactly the same question: markets are fast and forward-looking, while expert assessments are slower but richer in structural context. Second, it makes clear that posterior probabilities are endogenous to observable developments and can therefore be updated as the geopolitical environment evolves. Third, it supports communication, because the movement from prior to posterior via the grouped-likelihood update can be explained in intuitive language without abandoning formal discipline.

3 The supply block: from geopolitics to net oil-supply shortfalls

The second block translates geopolitical states into physical oil-market disruptions. This avoids mapping scenarios directly into prices without an intervening quantity mechanism, which is a central criticism of purely narrative scenario approaches.

For each scenario s, the gross supply path g_t^s is given by equation (2). This path is generated by the scenario-specific dynamic profile: the post-shock factor, the recovery horizon, and the long-run residual supply term. Benign scenarios therefore normalise quickly, while severe scenarios peak higher and recover more slowly.

The economically relevant quantity is not the gross shortfall but the *effective* shortfall after accounting for inventory releases, reserve releases, and other available buffers. Let b_t denote the buffer release from commercial stocks or official reserves at time t. These releases are typically front-loaded over the first months and then taper according to a calibrated buffer path. The net shortfall entering the price block is

$$q_t^s = \max(g_t^s - b_t,\ 0). \tag{9}$$

The non-negativity restriction is important, because buffers can offset a deficit but cannot create a net positive supply shock. This also avoids double counting: buffers first narrow the physical gap, while commercial inventory rebuilding later enters the price equation through a distinct premium.

In benign scenarios, the long-run term can allow for a small normalization overshoot relative to the disrupted baseline. Such overshooting is not allowed to enter the net-shortfall calculation as a negative physical shortfall. As a result, benign scenarios do not generate negative scarcity signals, and the cumulative deficit used in the inventory-rebuilding calculation is non-decreasing.

Two features of the buffer b_t are particularly important. *First*, low initial inventory levels limit the ability to absorb shocks, amplifying the price response of a given disruption (Kilian, 2009; Kilian and Murphy, 2014). *Second*, rebuilding inventories can prolong the impact of a disruption even after production recovers, because available barrels are diverted to restocking. This effect is captured separately through the stock-rebuild premium in the price block. Longer-lasting capacity losses under severe conflict are already embedded in the long-run residual term ℓ^s of the supply path.

The framework treats b_t as an *exogenous*, policy-set input: a front-loaded monthly buffer path calibrated from historical reserve-release episodes and reflecting available inventories, official reserve releases, and the feasible pace of coordinated release. The model abstracts from strategic-reserve replenishment as endogenous demand. Strategic reserves can cushion the contemporaneous shortfall, but their subsequent replenishment is a policy decision rather than a market-clearing demand schedule. The stock-rebuild premium should therefore be interpreted as a reduced-form commercial inventory and market-tightness premium.

4 The oil price block: from net shortfalls to Brent price paths

The third block translates the scenario-specific net supply shortfall into a Brent price path. The specification is deliberately parsimonious and builds on reduced-form evidence on the response of oil prices to physical supply disruptions (Kilian, 2009; Kilian and Murphy, 2014; Hamilton, 2009; Känzig, 2021).

The scenario price path consists of a fundamental Brent component and, when activated, a separately identified stock-rebuild premium SR_t^s, which captures temporary market tightness as commercial inventories are replenished.

For each scenario s, the fundamental log-price component is given by

$$lr_t^s = \frac{q_t^s/Q^{pre}}{\varepsilon}, \tag{10}$$

where q_t^s is the net supply shortfall defined in equation (9), Q^{pre} is pre-disruption oil supply, and $\varepsilon > 0$ is the absolute value of the short-run oil-demand elasticity.

Prices respond to the share of global supply that remains uncovered after inventory and reserve releases. A lower value of ε implies a larger price response to a given scarcity share. Shock persistence is captured by the scenario-specific path of q_t^s: slower supply recovery keeps the shortfall elevated for longer and generates a more persistent Brent price response.

The specification separates the contemporaneous physical-scarcity channel from other sources of price pressure. Geopolitical information affects the probability assigned to each scenario, while the fundamental price equation maps the corresponding net physical shortfall into Brent prices.

A separate geopolitical risk-premium term can optionally be activated for sensitivity analysis or reverse-engineering exercises:

$$lr_t^s = \frac{q_t^s/Q^{pre}}{\varepsilon} + \beta_R R_t^s, \tag{11}$$

where $R_t^s = 0$ in the baseline forward paths. The fundamental log-price component is translated into a Brent price path through an anchored price transformation:

$$P_t^{fund,s} = \mathcal{A}_t\left(lr_t^s; P_a, P^{pre}\right), \tag{12}$$

where $\mathcal{A}_t(\cdot)$ denotes the transformation that anchors the scenario path to the latest observed Brent price P_a, while retaining the pre-disruption price benchmark P^{pre}. The exact construction of this transformation, including the treatment of extreme relative price movements, is described in Appendix B.

The model can also include a temporary stock-rebuild premium, SR_t^s, capturing the market tightness associated with the rebuilding of commercial inventories after emergency buffers have been used. When this component is activated, the scenario price is given by

$$P_t^s = P_t^{fund,s} + SR_t^s, \tag{13}$$

where $P_t^{fund,s}$ denotes the Brent path implied by the contemporaneous physical scarcity channel, while SR_t^s captures the temporary pressure associated with replenishing commercial inventories after emergency buffers have been used. Given the limited number of comparable historical episodes, the calibration of the stock-rebuild premium is subject to considerable uncertainty. The premium is therefore treated as an optional add-on that can be included or excluded from the scenario price paths. Its default calibration is discussed in Appendix B.

5 Inverting the framework: from prices to scenarios

The three-block structure can also be used in reverse as a diagnostic exercise. Given observed or externally supplied prices and a calibrated supply and price block, one can infer which physical shortfalls and which geopolitical scenarios are most consistent with those prices. This is most straightforward when the uncertainty is primarily about the state vector S rather than about the parameters of the supply and price equations. It should not be interpreted as a unique structural inversion, since observed oil prices may also reflect non-geopolitical shocks, risk premia, inventory hoarding, or other market conditions outside the scenario space.

The reverse-engineering module separates three related but distinct diagnostics. The first maps external Brent paths into price-implied net shortfalls. The second maps each external Brent price path into the geopolitical scenario space by comparing it with the model-implied scenario price paths. The third maps the model's geopolitical posterior back into the external scenario-path space to obtain external scenario-path weights.

These diagnostics do not re-estimate the Bayesian posterior π_t produced by the probability block in Section 2.3. They are comparison tools that help interpret external price paths in the language of the model and compare the model's posterior with alternative external scenario paths.[2]

Price-implied net shortfall. The first diagnostic translates an external Brent path into an implied physical quantity. This is relevant because it attaches a physical narrative to external price scenarios and helps policymakers assess whether a given price path is consistent with a plausible supply disruption under the model.

To obtain an implied net shortfall from an external Brent path, when activated, the stock-rebuild premium is first removed as an additive price-level term and the anchored log-price move is inverted relative to the path's own anchor month a:

$$\widehat{q}_t = \max\Big(q_a + \varepsilon Q^{pre}\big[\log\tfrac{P_t - SR_t}{P_a - SR_a} - \beta_R(R_t - R_a)\big], 0\Big). \tag{14}$$

Here $\varepsilon > 0$ is the short-run oil-demand elasticity in absolute value, Q^{pre} is pre-disruption

[2] In the empirical application, the external scenario paths correspond to the BMPE oil scenarios provided by the ECB.

oil supply, q_a is the model-implied net shortfall at the anchor, and $\beta_R R_t$ is the optional risk-premium term. In the baseline, $R_t \equiv 0$, so the risk-premium subtraction is zero.

This reconstruction is *anchor-relative* and should not be read as a unique inversion of the full forward price mapping. In particular, it does not invert through the forward branch-selection $\max\{P_a e^{\Delta lr}, P^{pre} e^{lr}\}$ nor through the cap and floor applied to the relative log move. These non-linearities can censor the mapping from prices to physical shortfalls. The reported $\widehat{q}_t$ should therefore be interpreted as a diagnostic price-implied net shortfall under the selected anchor-relative convention, not as a uniquely identified production loss.

Then, the implied gross shortfall is obtained by adding back the buffer used:

$$\widehat{g}_t = \widehat{q}_t + b_t. \tag{15}$$

Price paths to geopolitical mapping. The second diagnostic maps external Brent paths into the geopolitical scenario space. Each external path is compared with each internally generated geopolitical scenario path over the forward window. The baseline mapping is price-only: distances are computed from Brent price paths, while the quantity leg is retained only as a sensitivity. This default is important because the price-implied net shortfall is itself derived from the external price path; using both price and quantity with equal weight would risk double counting the same information.

This diagnostic is useful for assessing how external price scenarios relate to the geopolitical scenarios in the model. It answers whether a given external price path is closer to a benign normalization scenario, an intermediate disruption scenario, or a severe escalation scenario.

Let P_t^u denote external path u, and let P_t^s denote the model-implied Brent path under geopolitical scenario s. The algorithm computes a distance $d_{u,s}$ between each external path and each geopolitical scenario path, and transforms these distances into scenario-conditional weights $M_{u,s}$. These weights answer the question: conditional on external path u, which geopolitical scenario is closest to it? They should be interpreted as diagnostic mapping weights, not as re-estimated geopolitical probabilities.

Geopolitical posterior to external scenario paths. The third diagnostic uses the same mapping matrix in the opposite direction. Given the geopolitical posterior $\pi_s = P_{\text{model}}(S_s)$, the model computes a score for each external Brent path by combining the posterior probabilities with the mapping weights:

$$\text{score}_u = \sum_s \pi_s M_{u,s}. \tag{16}$$

These scores are then normalised across external paths to obtain external scenario-path weights:

$$\omega_u^{\text{ext}} = \frac{\text{score}_u}{\sum_{u'} \text{score}_{u'}}. \tag{17}$$

These weights summarise which external scenario paths are most consistent with the model's geopolitical posterior.

The technical details of the price-implied shortfall calculation and of the distance-based mapping—including the distance metric, normalization rule, softmax temperature and orientation of the mapping matrix—are reported in Appendix B.6.

In practice, these reverse-engineering and mapping diagnostics must be interpreted carefully. First, they treat the price block parameters as known, ignoring parameter uncertainty. Second, they compare external price paths with the limited set of geopolitical scenarios under consideration, whereas observed oil prices may also embed other shocks. Third, the exercises are static and do not exploit the full time-series of prices and indicators. Fourth, when the optional risk-premium overlay is activated, the exercise should be read as a decomposition of prices into physical and non-physical components rather than as a pure recovery of physical shortfalls. The same observed futures curve can be rationalised by larger implied supply losses, by larger weights on adverse scenarios, or by a non-physical premium associated with geopolitical risk, precautionary demand, or inventory hoarding.

Nonetheless, the reverse-engineering module provides a useful diagnostic: when the implied external scenario-path weights diverge sharply from the Bayesian posterior in Section 2.3, the discrepancy itself becomes informative about the difference between market pricing and the model-based geopolitical assessment.

6 An illustration based on the U.S.–Iran conflict of 2026

This section illustrates the framework using the U.S.–Iran conflict of 2026 and the calibration prepared for the June 2026 Eurosystem staff macroeconomic projections exercise (BMPE) (European Central Bank, 2026). The exercise shows how the framework can be used as a structured complement to the standard Eurosystem scenario methodology and as a real-time updating device once geopolitical conditions change.

We proceed in three steps. First, using only the information available at the 9 June 2026 Governing Council cut-off, we estimate the prior and posterior probabilities of the geopolitical scenarios through the probability set-up developed in Section 2. Second, we pass the resulting scenario distribution through the supply and price blocks developed in Sections 3 and 4, obtaining scenario-specific and posterior-weighted Brent price paths and mapping them into the BMPE scenario space. We repeat these two steps using the information available on 15 June and 8 July 2026 [3] to show how changes in market, diplomatic, and real-time indicators translate into revised probabilities, price paths, and BMPE-style diagnostic weights. Third, we complement this forward exercise with the reverse-mapping diagnostic developed in Section 5. Starting from observed oil futures prices, the reverse exercise identifies the combination of geopolitical scenario paths that is most consistent with the observed futures curve. The resulting market-implied scenario weights provide a diagnostic benchmark and do not feed back into the Bayesian posterior.

6.1 Scenario probabilities at the time of the publication of the projections

The Eurosystem June 2026 BMPE report published on 9 June 2026 calibrates its three alternative energy scenarios—Milder, Adverse, and Severe—from option-implied distributions for oil and European gas prices (European Central Bank, 2026). Concretely, the June 2026 BMPE defines these scenarios by reading off selected percentiles of the market-implied density at the cut-off date of 21 May 2026: the 25th percentile for the Milder scenario, the 75th for the Adverse scenario, and the 95th for the Severe scenario. Staff explicitly state that these percentiles "provide a market-based assessment of the current risks surrounding energy prices and implicitly cover alternative evolutions of the war in the Middle East and its implications for energy supply disruptions" (European Central Bank, 2026). The resulting oil-price assumptions span a wide range: USD 88/bbl (Milder), USD 122/bbl (Adverse), and USD 166/bbl (Severe) in Q3 2026, with corresponding deviations from the baseline of approximately −15%, +19%, and +62% respectively; analogous ranges apply to wholesale gas prices (Table 5 in European Central Bank 2026).

This approach has clear virtues: it is market-consistent, replicable, and communicable. However, it has two important limitations for a geopolitically motivated analysis. First, it reads probabilities from the price distribution rather than from a structural model of the underlying conflict: option-implied densities aggregate all market participants' beliefs—including views unrelated to the geopolitical situation itself—and do not separately identify the physical quantities (supply losses, buffer releases, inventory dynamics) that generate those price outcomes. Second, because the percentile mapping defines scenarios by their position in the price distribution rather than by their geopolitical content, there is no transparent link between the assumed energy price path and the specific supply-disruption mechanism that would produce it. The scenarios are thus best understood as a market-implied risk quantification, not as a structural account of how geopolitical events translate into macroeconomic outcomes.

The framework developed in this paper addresses both limitations by moving toward a

[3] These dates were selected because they represent two distinct shifts in the geopolitical information set. By mid-June, market and diplomatic signals had moved markedly towards a durable de-escalation, whereas the early-July update captures the subsequent reassessment associated with the resumption of strikes.

Table 1: From the probabilities of the geopolitical scenarios to assigning weights to given BMPE scenarios

Geopolitical scenario	Governing Council June		Peace announcement		Strikes resumption	
	9 June 2026		15 Jun 2026		8 Jul 2026	
	Prior	Posterior	Prior	Posterior	Prior	Posterior
Diplomatic Resolution (S1)	34	21	58	45	41	34
Prolonged Stalemate (S2)	36	53	24	32	28	28
Partial Resumption (S3)	18	22	11	15	21	29
Major Escalation (S4)	12	4	7	8	10	9
BMPE scenario	Probability (%)		Probability (%)		Probability (%)	
Milder	23		36		30	
Baseline	30		32		29	
Adverse	35		23		26	
Severe	12		9		15	

Note: The top panel reports geopolitical scenario probabilities computed at three specific dates: 9 June 2026, 15 June 2026 and 8 July 2026: 9 June is the date of the June ECB Governing Council and also the publication date of the June 2026 BMPE projections; on 15 June a peace deal was announced; on 8 July 2026 the conflict resumed, despite a peace deal being in place, with new military strikes taking place. The bottom panel reports the corresponding composite BMPE scenario probabilities obtained by mapping the geopolitical posterior into the BMPE scenario space using the distance-based scenario-mapping matrix. The reported results include the stock-rebuild premium. Source: authors' calculations and European Central Bank (2026).

more quantitative geopolitical grounding of the scenario energy price paths. Rather than reading percentiles from an option-implied density, it builds price paths from the bottom up: geopolitical scenarios determine supply-loss profiles (Section 3), which, after netting out inventory and reserve buffers, enter the reduced-form price equation (Section 4) to produce Brent paths that are physically interpretable. The probability assigned to each scenario comes from Bayesian updating on a richer set of signals (Section 2.3) rather than from a single market-implied distribution. The two approaches are therefore *complementary*: the option-implied percentiles provide a market-consistency check on the geopolitically derived price paths, while the diagnostic mapping described in Appendix B.6 shows which geopolitical scenarios are closest to each external benchmark price path.

As of the date of the June 2026 ECB Governing Council meeting, the model posterior assigns the geopolitical scenario probabilities shown in the top panel of Table 1. The posterior is obtained by combining market-implied priors with structured think-tank assessments and updating them using real-time indicators, notably Hormuz shipping traffic, Brent calendar spreads, Brent prices, tanker-insurance premia, and option-implied skew. For example, for the 9 June update, relative to the prior, the posterior reduces the probability of both the benign tail and the most acute escalation state, while shifting probability mass toward the two intermediate scenarios. The distribution is centred on a persistent but not catastrophic disruption: the most likely outcome at the date of the 9 June ECB Governing Council meeting was a prolonged stalemate, followed by partial resumption with continuing frictions. The information available on 9 June was thus neither consistent with a rapid normalization of oil flows nor with a full-scale regional breakdown, but rather with an extended period of tightness, insecurity, and only partial reopening.

The probability configuration at the 9 June date, to follow with the example, is informative in a formal sense. The KL divergence of the posterior $\boldsymbol{\pi} = (0.21, 0.53, 0.22, 0.04)$ from the uniform benchmark $\mathbf{u} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$ is

$$D_{\mathrm{KL}}(\boldsymbol{\pi} \,\|\, \mathbf{u}) = \sum_{s=1}^{4} \pi_s \log \frac{\pi_s}{1/4} \;\approx\; 0.26 \text{ nats } (0.38 \text{ bits}), \tag{18}$$

indicating that the indicator set meaningfully discriminates across states rather than leaving the distribution close to maximum entropy. Economically, most of the mass lies in states that are not catastrophic in a pointwise sense but nonetheless imply persistent supply impairment and therefore sustained price pressure.

Translating the posterior through the supply and price blocks yields a deterministic, posterior-weighted expected Brent level of \$105.06/bbl in December 2027 and \$89.92/bbl in December 2028. Around these central values, a Monte-Carlo diagnostic that perturbs the supply-loss and price-equation parameters ($n = 800$ draws, fixed seed) produces a p_{10}–p_{90} band of \$82–\$244 and \$70–\$219, respectively. The wide bands reflect the uncertainty introduced by the simulated supply-loss parameters and price-equation elasticities, evaluated through the posterior-weighted expected Brent path. The distribution is markedly right-skewed: even a relatively small probability assigned to Major Escalation generates a sizeable upside tail in the expected price distribution.

Comparing these geopolitically derived price paths with the Eurosystem option-implied percentiles provides an implicit calibration check. The framework's expected path (around \$105/bbl in late 2027) lies between the Adverse (USD 90/bbl) and Severe (USD 124/bbl) BMPE scenarios at a similar horizon, consistent with the posterior assigning the bulk of its mass to the two intermediate conflict states rather than to rapid resolution or severe escalation. Figure 2 shows how the geopolitical posterior is translated into BMPE-style diagnostic weights for the June BMPE macro scenarios (Milder, Baseline, Adverse, Severe) at the 9 June date using the price-path mapping diagnostic described in Appendix B.6 (for other dates see also Table 1). Each entry in the matrix is derived from the oil-price-path distance between a geopolitical scenario and an external BMPE benchmark path. The resulting weights summarise which external paths are closest to the posterior-implied distribution of geopolitical Brent paths. They are used as a policy-communication diagnostic and do not feed back into the Bayesian posterior.

We stress that this mapping is an *illustrative diagnostic* rather than a primitive of the model. The matrix is not a fixed, expert-set object: it is recomputed from the price paths implied by the supply and price blocks, via a softmax over the oil-price-path distance between each geopolitical scenario and each BMPE benchmark path. The ex-

Figure 2: From geopolitical scenarios to June 2026 BMPE Brent oil price scenario probabilities.

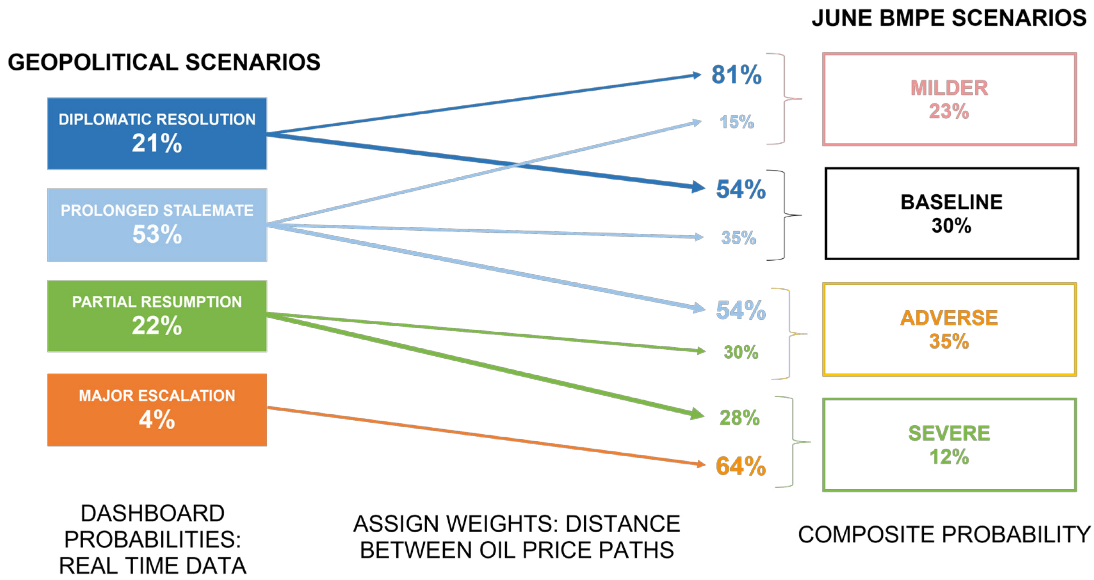

Note: Left block: dashboard posterior probabilities over the four geopolitical scenarios computed on 9 June 2026. Center block: scenario-mapping matrix weights based on oil-price-path distance between geopolitical and BMPE scenarios. Right block: composite BMPE scenario probabilities (Milder 21%, Baseline 28%, Adverse 36%, Severe 15%). The reported results include the stock-rebuild premium. Source: authors' calculations and European Central Bank (2026).

act construction—distance metric, temperature, and horizon weighting—is given in Appendix B.6. The mapping is therefore best read as a bridge that places the geopolitical posterior on the same axis as the official macro scenarios, not as an independent estimate of their probabilities.

6.2 Scenario probabilities after the mid-June de-escalation signals

The exercise is then updated at two subsequent information dates in order to illustrate how the framework reacts to changes in the geopolitical news flow. Relative to the 9 June Governing Council cut-off, the mid-June update incorporates new prediction-market signals and real-time indicators pointing to a more benign geopolitical configuration. The updated prior block uses prediction-market prices that assign very high probabilities to a permanent U.S.–Iran peace deal by both 31 July 2026 and 31 December 2026, while also attaching a sizeable probability to a normalization of Strait of Hormuz traffic by the end of June. The update further incorporates refreshed high-frequency indicators observed at the updated dates, including the Brent–WTI spread, the Brent front-month/12-month spread, a Middle East geopolitical-risk index, the tempo of Iranian missile and drone launches, and the diplomatic activity indicators. In this update, the expert-assessment block is left unchanged, so the movement in probabilities is driven entirely by new market information and real-time signals.

Figure 3: Updated geopolitical scenario probabilities as of 8 July 2026.

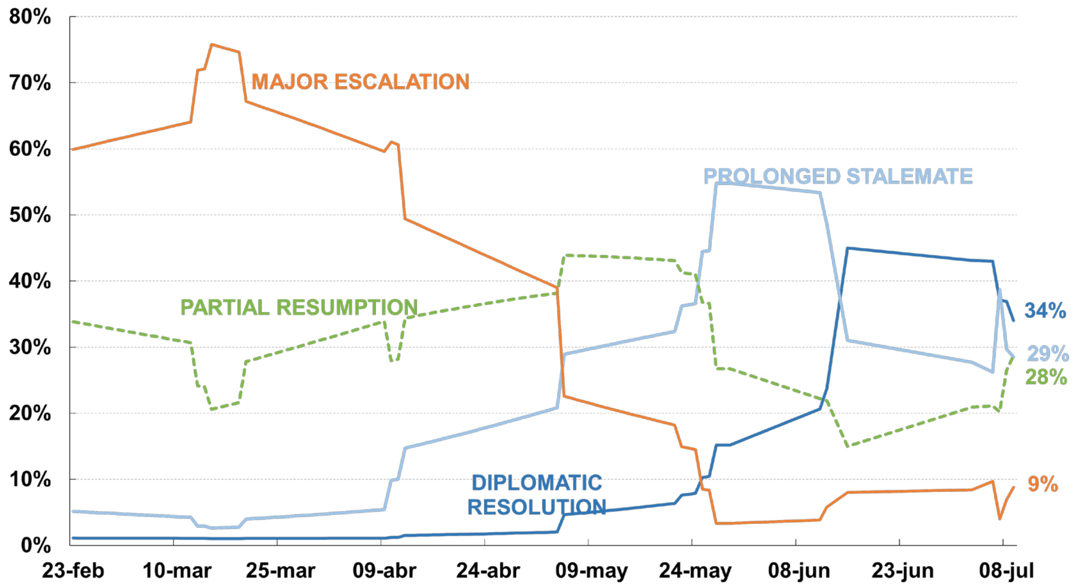

Note: Time-varying probabilities of the four geopolitical scenarios in the 15 June 2026 update. By that date, the posterior assigns roughly 45% to Diplomatic Resolution, 32% to Prolonged Stalemate, 15% to Partial Resumption, and 8% to Major Escalation. Source: authors' calculations.

The resulting posterior distribution shifts materially away from the late-May configuration. As shown in Table 1, the 15 June posterior assigns 45% to Diplomatic Resolution, 32% to Prolonged Stalemate, 15% to Partial Resumption and 8% to Major Escalation.

The third update, dated 8 July 2026, illustrates how the same framework responds when hostilities resume. Relative to the peace-announcement update, the probability of Diplomatic Resolution falls, while probability mass shifts back toward Partial Resumption and, to a lesser extent, Major Escalation. The posterior therefore moves from a de-escalation configuration toward a more fragile environment in which renewed disruptions become more relevant, even though the distribution does not return to the severe-tail configuration.

Overall, the comparison across the three dates shows how the framework translates changes in the geopolitical narrative into revised scenario probabilities. The 9 June cut-off is centred on a persistent but not catastrophic disruption, with Prolonged Stalemate as the modal scenario. The 15 June update shifts the distribution toward Diplomatic Resolution after the peace-announcement signals. The 8 July update partially reverses that movement, increasing the weight on Partial Resumption and renewed disruption risks.

The effect on prices is correspondingly large. Figure 4 compares the posterior-weighted Brent path with the corresponding futures curve across four geopolitical information sets:

Figure 4: Brent price paths across geopolitical information updates and futures' markets curves at the same dates

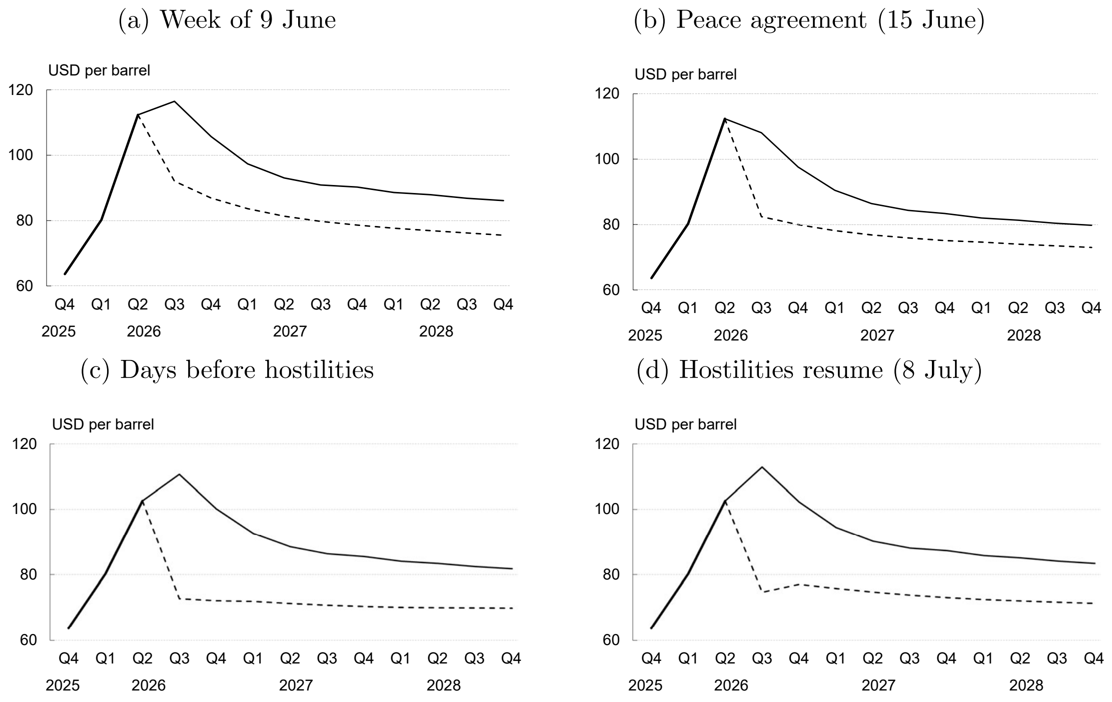

Notes: The figure compares the Bayesian Brent price path with the corresponding futures curve across four geopolitical information sets: the week of 9 June, the peace-agreement update, days before the resumption of hostilities at the beginning of July, and the resumption-of-hostilities update. The solid line reports the Bayesian posterior-weighted Brent path implied by the model at each date, while the dashed line reports the futures curve observed at the corresponding information date. The reported results include the stock-rebuild premium.

the week of the June Governing Council meeting, the peace-agreement update, the days before the resumption of hostilities, and the resumption-of-hostilities update.

The comparison shows how the price implications of the framework move with the information set. The peace-agreement update lowers the posterior-weighted Brent path relative to the earlier configuration, as the posterior shifts toward more benign geopolitical outcomes. By contrast, the two July panels show a partial reversal: renewed geopolitical stress raises the Bayesian path again and moves it away from the more benign futures configuration. This illustrates how the dashboard translates changes in geopolitical information into both probability revisions and price-path revisions.

From the perspective of external scenario-path weights, the sequence of updates implies a rotation in the distribution of risks. The peace-agreement update shifts weight away from the adverse tail and toward the more benign external paths, while the resumption-of-hostilities update partially reverses that movement.

As a robustness check, we also assess the sensitivity of the 8 July composite BMPE scenario probabilities to alternative price-block specifications. The results are broadly robust to a

Table 2: Robustness of composite BMPE scenario probabilities

BMPE scenario	Resumption of strikes (8 July 2026) (%)		
	BdE dashboard Baseline calibration	Higher elasticity $\varepsilon = 0.25$	No stock-rebuild premium $SR_t^s = 0$
Milder	30	31	38
Baseline	29	27	30
Adverse	26	24	16
Severe	15	18	16

Note: Entries are composite BMPE scenario probabilities. The first column reports the results for 8 July 2026 under the baseline BdE dashboard calibration. The second column reports a higher-elasticity specification, setting the short-run elasticity parameter to $\varepsilon = 0.25$. The third column deactivates the stock-rebuild premium, setting $SR_t^s = 0$. Source: authors' calculations.

higher short-run elasticity, while excluding the stock-rebuild premium shifts some probability mass from the Adverse scenario towards the Milder and Baseline scenarios. Table 2 reports the alternative specifications and their effects on the composite BMPE scenario probabilities, while Appendix B provides further details on the underlying calibrations.

Taken together, Table 1, along with Figures 2–4, show how the framework can be used as a real-time updating device within a projection round. At the 9 June Governing Council cut-off, the framework points to a posterior centred on a persistent but not catastrophic disruption. The peace-agreement update shifts the posterior and the Bayesian Brent path in a more benign direction. The July update, in turn, shows how renewed geopolitical stress can quickly rotate the distribution back toward partial resumption and renewed-disruption risks. This is the policy value of the framework in a high-frequency geopolitical environment: it not only quantifies the balance of risks at a point in time, but also explains coherently why that balance changes as new geopolitical and market information arrives.

7 Concluding remarks

Geopolitical uncertainty around the Middle East is best handled through a structured scenario distribution; scenario analysis is most useful when it passes through a physical supply block; and oil-price outcomes depend critically on persistence, inventories, and premia as well as on the initial supply shock. The value of the framework developed in this paper lies in making those links transparent and traceable from assumptions to outcomes.

The analysis has focused on the methodology and on a single illustrative application to the 2026 U.S.–Iran conflict and the June 2026 Eurosystem projection round. The framework shows how market signals, expert judgment, and observed indicators can be integrated coherently via Bayesian updating; how geopolitical scenarios can be translated into physically meaningful supply paths with explicit parameter calibration; and how those supply paths can be mapped into oil prices without collapsing all mechanisms into a single reduced-form coefficient. In doing so, it offers a disciplined and adaptable template for analysing geopolitical oil-market risk in real time and for embedding that analysis in central-bank risk assessments and projection rounds, and it is portable to other regions and energy markets provided comparable indicators and scenario taxonomies are available.

Three directions for future work stand out. First, the scenario synthesis of Adrian et al. (2025) provides a natural bridge between the geopolitical framework developed here and existing density forecasts at central banks and international institutions, and could be used systematically to test the completeness of geopolitical scenario sets used in practice. Second, richer transition matrices $\mathbf{Q}$ estimated from conflict databases would improve the treatment of scenario dynamics in Section 2.3 and allow for a more realistic link between short-run indicators and medium-term conflict paths. Third, the explicit inversion of the framework to recover implied geopolitical scenarios from options-market data and futures curves, sketched in Section 5, deserves formal development and could help quantify when markets are pricing in geopolitical states that diverge from expert or policymaker priors.

References

Adrian, Tobias, Domenico Giannone, Matteo Luciani and Mike West. (2025). "Scenario Synthesis and Macroeconomic Risk". IMF Working Papers, WP/25/105, International Monetary Fund. https://www.imf.org/-/media/files/publications/wp/2025/english/wpiea2025105-print-pdf.pdf

Aiyar, Shekhar, Jiaqian Chen, Christian H. Ebeke, Roberto Garcia-Saltos, Tryggvi Gudmundsson, Anna Ilyina, Alvar Kangur, Tansaya Kunaratskul, Sergio L. Rodriguez, Michele Ruta, Tatjana Schulze, Gabriel Soderberg and Juan P. Trevino. (2023). "Geoeconomic Fragmentation and the Future of Multilateralism". IMF Staff Discussion Notes, SDN/2023/001, International Monetary Fund. https://www.imf.org/en/publications/staff-discussion-notes/issues/2023/01/11/geo-economic-fragmentation-and-the-future-of-multilateralism-527266

Alonso-Alvarez, Irma, Marina Diakonova and Javier J. Pérez. (2025). "Rethinking GPR: The Sources of Geopolitical Risk". Documentos de Trabajo, 2522, Banco de España. https://doi.org/10.53479/39685

Austvik, Ole Gunnar. (2018). "Concepts of Geopolitics and Energy Security". *IAEE Energy Forum,* Second Quarter, pp. 25-28. http://kaldor.no/energy/IAEE-EnergyForum201803-geopolitics-concepts-pubished.pdf

Baley, Isaac, and Laura L. Veldkamp. (2021). "Bayesian Learning: A Survey". NBER Working Paper Series, 29338, National Bureau of Economic Research. https://www.nber.org/papers/w29338

Baumeister, Christian, and Lutz Kilian. (2016). "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us". *Journal of Economic Perspectives,* 30(1), pp. 139-160. https://doi.org/10.1257/jep.30.1.139

Bondarenko, Yevheniia, Vivien Lewis, Matthias Rottner and Yves Schüler. (2024). "Geopolitical Risk Perceptions". *Journal of International Economics,* 152(104005). https://doi.org/10.1016/j.jinteco.2024.104005

Braun, Robin. (2023). "The importance of supply and demand for oil prices: Evidence from non-Gaussianity". *Quantitative Economics,* 14(4), pp. 1163-1198. https://doi.org/10.3982/QE2091

Caballero, Ricardo J., Emmanuel Farhi and Pierre-Olivier Gourinchas. (2008). "An Equilibrium Model of 'Global Imbalances' and Low Interest Rates". *American Economic Review,* 98(1), pp. 358-393. https://doi.org/10.1257/aer.98.1.358

Caldara, Dario, Michele Cavallo and Matteo Iacoviello. (2019). "Oil Price Elasticities and Oil Price Fluctuations". *Journal of Monetary Economics,* 103, pp. 1-20. https://doi.org/10.1016/j.jmoneco.2018.08.004

Caldara, Dario, and Matteo Iacoviello. (2022). "Measuring Geopolitical Risk". *American Economic Review,* 112(4), pp. 1194-1225. https://doi.org/10.1257/aer.20191823

Ciccarelli, Matteo, Matthieu Darracq Pariès, Bettina Landau and João Sousa. (2025). "Exploring an uncertain future with the help of scenarios". *The ECB Blog,* 15 January. https://www.ecb.europa.eu/press/blog/date/2025/html/ecb.blog.20250115~f016f263dd.en.html

Cordesman, Anthony H., and Abdullah Toukan. (2012). "Iran and U.S. Options for Preventive Military Strikes". CSIS report, Center for Strategic and International Studies. https://www.csis.org/analysis/iran-and-us-options-preventive-military-strikes

European Central Bank. (2026). "Eurosystem Staff Macroeconomic Projections for the Euro Area, June 2026". Tech. rep. Technical assumptions and scenario analysis. https://www.ecb.europa.eu/pub/projections/html/index.en.html

Hamilton, James D. (2009). "Causes and Consequences of the Oil Shock of 2007–08". *Brookings Papers on Economic Activity,* 40(1), pp. 215-283. https://www.brookings.edu/wp-content/uploads/2009/03/2009a_bpea_hamilton.pdf

Ioannou, Demosthenes, Javier J. Pérez, Hans Geeroms, Isabel Vansteenkiste and Pierre- François Weber. (2023). "The EU's Open Strategic Autonomy from a Central Banking Perspective: Challenges to the Monetary Policy Landscape from a Changing Geopolitical Environment". ECB Occasional Paper Series, 311, European Central Bank. Report by the International Relations Committee Work stream on Open Strategic Autonomy. https://doi.org/10.2866/602215

Känzig, Diego R. (2021). "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements". *American Economic Review,* 111(4), pp. 1092-1125. https://doi.org/10.1257/aer.20190964

Katzman, Kenneth. (2021). "Iran: Internal Politics and U.S. Policy and Options". Tech. rep., RL32048, Congressional Research Service. Updated 29 July 2021. https://www.congress.gov/crs_external_products/RL/PDF/RL32048/RL32048.365.pdf

Kilian, Lutz. (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market". *American Economic Review,* 99(3), pp. 1053-1069. https://doi.org/10.1257/aer.99.3.1053

Kilian, Lutz, and Daniel P. Murphy. (2014). "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil". *Journal of Applied Econometrics,* 29(3), pp. 454-478. https://doi.org/10.1002/jae.2322

Kilian, Lutz, Michael D. Plante, Alexander W. Richter and Xiaoqing Zhou. (2026). "The Impact of the 2026 Iran War on U.S. Inflation: A Scenario Analysis". Working Papers, 2609, Federal Reserve Bank of Dallas. https://doi.org/10.24149/wp2609

Kuzemko, Caroline, Mathieu Blondeel, Michael Bradshaw, Gavin Bridge, Erika Faigen and Louis Fletcher. (2025). "Rethinking Energy Geopolitics: Towards a Geopolitical Economy of Global Energy Transformation". *Geopolitics,* 30(2), pp. 531-565. https://doi.org/10.1080/14650045.2024.2351075

Liu, Jingye, Fengqi Guo, Ying Shi, Rijia Ding and Zhen Chen. (2025). "Assessing the Security of Crude Oil Supply Chain: The Case of China". *Energy Strategy Reviews,* 59(101733). https://doi.org/10.1016/j.esr.2025.101733

Meza, Abel, Ibrahim Ari, Mohammed Al Sada and Muammer Koç. (2026). "Implications of Interrupting the Hormuz Strait in the LNG Trade". *Journal of Transportation Security,* 19(28). https://doi.org/10.1007/s12198-026-00350-1

Polymarket. (2026). Polymarket Prediction Market Platform. https://polymarket.com/

Appendices

A Operational dashboard

This appendix briefly describes the operational dashboard that implements the three blocks of the framework for practical use. The dashboard is not required to follow the analysis in the main text, but it illustrates how the methodology can be embedded in a policy-oriented tool.

The interface organises the three blocks—scenario probabilities, physical disruptions, and price paths—on a single screen. The top panel displays the four geopolitical scenarios and their current probabilities, together with simple controls that allow the user to adjust the prior weights on market signals and expert assessments, as well as to toggle alternative likelihood specifications for the indicator block. The middle panel reports the implied monthly supply shortfall for each scenario and the probability-weighted expected deficit. The bottom panel shows the corresponding Brent price paths, both by scenario and in expectation, with the latest observed quote used as an anchor.

Figures A.1–A.3 show an illustrative configuration. In this example, the posterior gives substantial weight to a prolonged stalemate, with smaller but non-negligible mass on partial resumption and major escalation. The prior gives a more even weight between diplomatic resolution and stalemate, but the posterior shifts probability as real-time indicators (for example, shipping activity or market spreads) move in favour of more or less persistent disruption. The probability panel also reports the weight assigned to the market-implied channel, which can be varied to examine the sensitivity of the posterior.

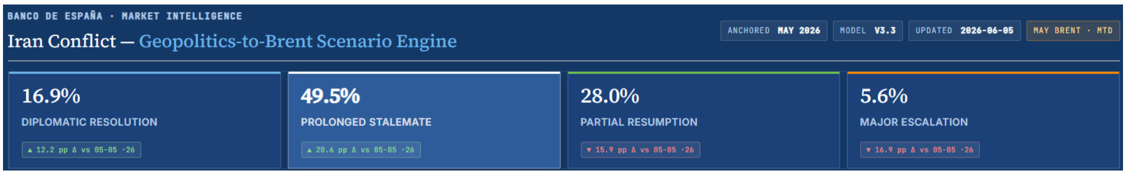

Figure A.1: Example probability panel: scenario labels.

Figure A.3 illustrates the supply and price blocks for the same configuration. Each coloured line in the left panel corresponds to the effective monthly shortfall in million barrels per day under a given scenario, after accounting for inventory and reserve releases. The figure highlights that persistence and inventory depletion can generate a sizeable expected deficit even when the modal scenario does not involve full closure. The right panel presents the implied Brent price paths by scenario and in expectation, together with the latest futures curve. In this example, the futures curve lies between the Diplomatic Resolution and Partial Resumption scenarios in the near term, but the distribution is skewed: a modest probability mass on Major Escalation generates a pronounced upside tail. This representation helps relate changes in beliefs about the geopolitical outlook to shifts in

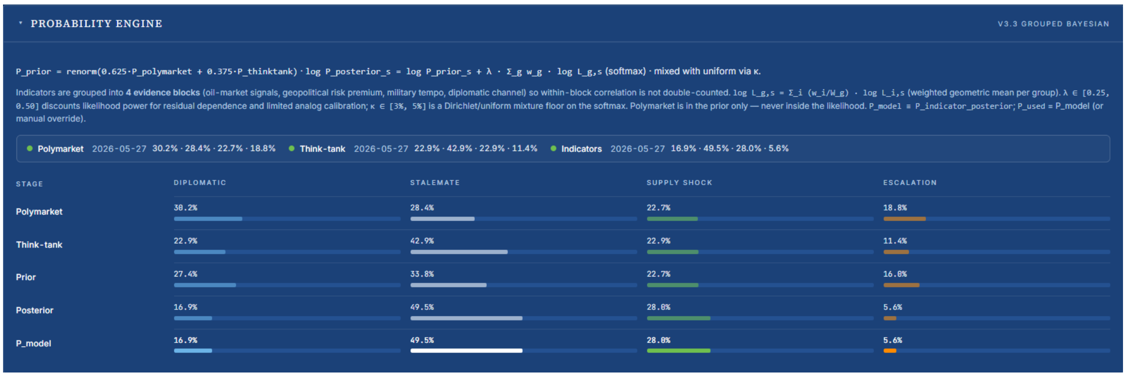

Figure A.2: Example probability panel: posterior probabilities.

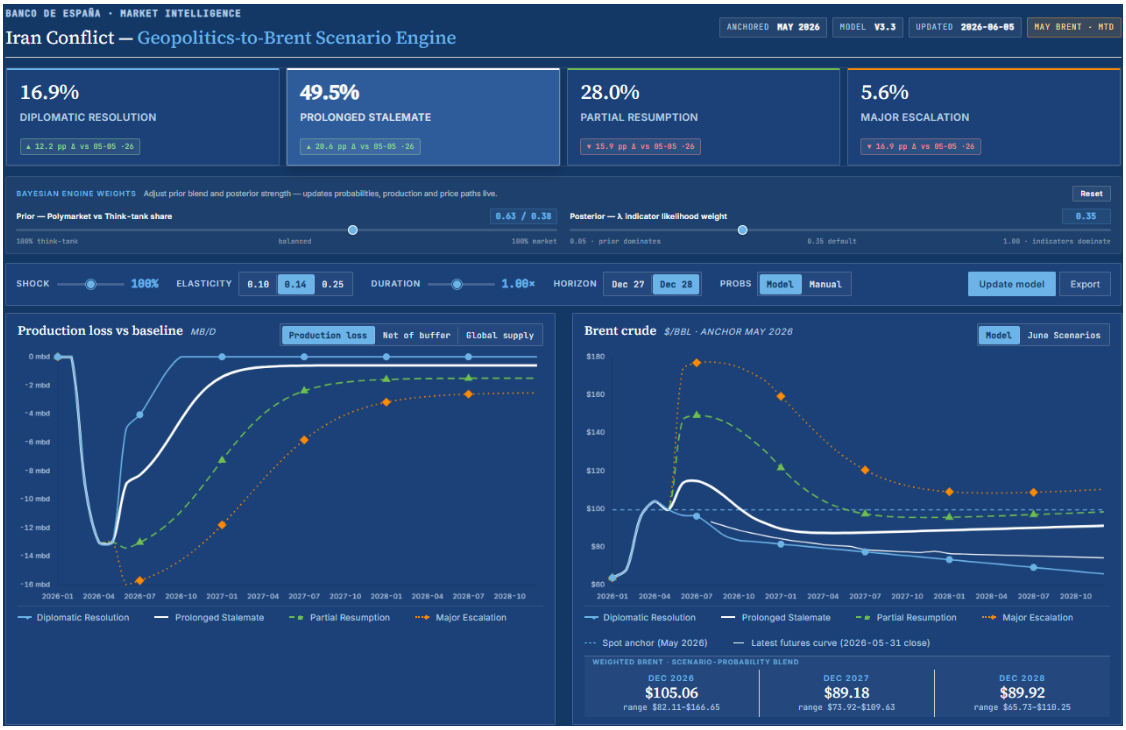

Figure A.3: Example effective supply shortfall and Brent price paths.

the expected price path and its dispersion.

Beyond this static representation, the dashboard is designed as an interactive audit and robustness tool. Each block can be expanded to show the inputs and intermediate objects used by the model. In the probability block, the user can inspect the construction of the prior, the grouped indicator likelihoods, the raw indicator values, their low/medium/high buckets, the original and freshness-adjusted weights, and the source attached to each signal. This makes the movement from prior to posterior traceable: changes in the posterior can be linked to specific evidence blocks, individual indicators, or data-freshness adjustments rather than treated as a black-box update.

The dashboard also allows most calibration objects to be varied directly. Users can edit the market and expert inputs, modify indicator buckets or weights, change the treatment of stale data, alter supply-side parameters and buffer paths, and activate optional pricing overlays such as the non-physical risk premium. Dedicated diagnostic panels report how the scenario probabilities change under alternative assumptions, including headline freshness-adjusted posteriors, no-freshness variants, exclusion of stale indicators, and exclusion of selected evidence blocks. These diagnostics are intended to separate the accepted model output from sensitivity exercises while keeping both visible for policy discussion.

Finally, the lower part of the dashboard implements reverse engineering and mapping diagnostics. Externally supplied Brent paths, such as futures or BMPE-style assumptions, can be inverted through the same supply and price blocks to recover price-implied net shortfalls and the corresponding implied gross shortfalls after adding back the buffer used. In this exercise, the stock-rebuild premium is treated as an additive price-level term, while any optional risk premium is treated as a log-price component. This preserves consistency with the forward price equation and allows the user to distinguish physical shortfall, inventory-rebuilding pressure, and non-physical residual premia when comparing model-implied paths with external price assumptions.

B Baseline calibration and robustness

This appendix documents the baseline calibration of the scenario parameters, buffer paths, price-equation coefficients, the stock-rebuild premium, and the optional risk-premium overlay, and reports robustness checks.

B.1 Prior and indicator inputs for scenario probabilities

The probability block combines three sources of information: market-implied signals, expert assessments, and high-frequency indicators. The first two sources define the prior, while the third source updates that prior through the grouped Bayesian likelihood described in Section 2.3. The separation is useful because the three sources have different strengths. Prediction markets provide fast-moving, event-specific signals; expert assessments provide a slower-moving structural anchor; and live indicators discipline the posterior with observable developments in oil markets, military activity and diplomacy.

B.1.1 Market-implied prior

The market-implied component is based on a small set of liquid Polymarket contracts whose wording is sufficiently close to the geopolitical dimensions of the scenario taxonomy. The contracts are not intended to span the four scenarios one-for-one. Instead, each binary contract is mapped into the four-state space according to its economic content. Contracts related to a peace deal or a durable ceasefire load positively on Diplomatic Resolution and negatively on Major Escalation; contracts related to renewed military action or invasion

Table B.1: Signed loadings of Polymarket contracts on the four geopolitical scenarios

Contract	S1 Diplomatic resolution	S2 Prolonged stalemate	S3 Partial resumption	S4 Major escalation
Peace deal by 31 Jul 2026	1.0	0.2	−0.5	−1.0
Peace deal by 31 Dec 2026	0.7	0.4	−0.3	−0.8
Ceasefire holds by 31 Jul 2026	0.6	0.6	−0.3	−1.0
Hormuz normal by 31 Jul 2026	0.9	0.3	−0.7	−1.0
U.S. invasion / renewed strikes	−1.0	−0.3	0.6	1.0

load in the opposite direction; and contracts on the normalization of Strait of Hormuz traffic are interpreted as direct signals about the persistence of the oil-supply disruption.

In the 15 June 2026 update, the market-implied prior used five contracts:

- a U.S.–Iran permanent peace deal by 31 July 2026;
- a U.S.–Iran permanent peace deal by 31 December 2026;
- whether a Trump-announced U.S.–Iran ceasefire is over by 30 June 2026;
- whether Strait of Hormuz traffic returns to normal by the end of June 2026;
- whether the United States invades Iran before 2027.

The corresponding market probabilities were 0.96, 0.97, 0.18, 0.70 and 0.12, respectively. Where the contract is written in the opposite direction to the model signal, the sign is adjusted before mapping it into scenarios. For example, a contract on the ceasefire being "over" is interpreted as the inverse of ceasefire durability. The market layer is therefore informative about both the probability of de-escalation and the expected duration of the disruption, but it enters only through the prior and never through the indicator likelihood.

The mapping from contracts to scenarios is summarised by the signed loading matrix in Table B.1.

Let $y_k \in [0,1]$ denote the (sign-adjusted) market probability of contract k and $\ell_{k,s}$ its loading on scenario s. The market-implied score for scenario s is $z_s = \sum_k (y_k - \frac{1}{2})\,\ell_{k,s}$, and the market-implied prior is obtained by a softmax with unit temperature, $P_s^M = \exp(z_s)/\sum_{s'} \exp(z_{s'})$, as in Equations (4)–(5). Positive loadings push probability towards the corresponding scenario and negative loadings away from it.

B.1.2 Expert assessments

The expert component is based on structured scenario assessments from policy and geopolitical research institutions, including CSIS, Chatham House and the Atlantic Council. These sources are selected because they provide timely and scenario-oriented analysis of the conflict, with direct implications for Gulf security, Hormuz transit and the duration of oil-market disruption. Their role is different from that of prediction markets. Expert assessments typically update less frequently, but they contain richer qualitative information on military constraints, diplomatic channels, regional spillovers and the plausibility of escalation paths.

In the operational implementation, expert inputs are updated through a weekly source-ingestion pipeline. Source-specific adapters retrieve material from the configured institutions, using the available access mode for each source. Retrieved documents are then normalised into a common format and filtered for relevance to the Iran–oil-supply scenario context. The relevance filter is LLM-assisted by default, but a deterministic keyword-based fallback is used when the language-model step is unavailable or does not provide sufficient evidence.

Relevant documents are then scored across the four geopolitical scenarios using a fixed scenario rubric. This scoring step is also LLM-assisted by default and includes a deterministic fallback. The output is a source-level score vector over the four scenarios. If a source is inaccessible, produces no parseable content, or cannot be scored reliably, the pipeline does not fabricate a new assessment: the previous score for that source is retained and the reason is recorded in the audit trail.

Each source score is subsequently normalised into a probability vector over the four scenarios. The expert component P_E is the source-averaged, renormalised vector,

$$P_E(S_i) \propto \frac{1}{n}\sum_{r=1}^{n} \sigma_{r,i},$$

where $\sigma_{r,i}$ is source r's score for scenario S_i. The expert vector is not used as a stand-alone posterior. Instead, it is blended with the market-implied vector to form the prior used in the Bayesian update.

The pipeline retains audit information on the retrieval method, source accessibility, number of retrieved and relevant documents, scoring mode, fallback use, evidence points, snippets and score changes. The expert component should therefore be interpreted as an auditable calibration input derived from structured source material, not as a regression-estimated probability model. In the baseline implementation, the prior gives a larger weight to the market-implied component and a smaller weight to the expert component. This reflects the higher frequency of market data while retaining the expert layer as a structural anchor against short-lived market fluctuations or contract-specific noise.

B.1.3 Indicator likelihoods

The third information source is a set of live indicators used to update the prior. Each indicator is assigned to one of four evidence blocks:

1. *Oil-market signals*: IMF PortWatch Hormuz transit calls, the Brent–WTI spread,

the Brent front-month to 12-month spread, and the change in the Brent price. These variables capture physical stress in oil markets, transport disruption and backwardation pressure.

2. *Geopolitical-risk premium*: Gulf transit war-risk insurance premia, the Iranian parallel-market exchange rate, and a country-level geopolitical-risk index. These variables capture non-physical risk, financial stress and insurance-market pressure.

3. *Military tempo*: Iranian missile or drone launch tempo and U.S. carrier-strike-group positioning. These variables proxy the kinetic intensity of the conflict and the risk of renewed escalation.

4. *Diplomatic channel*: Track-1 diplomatic meeting cadence. This variable captures the intensity of official or mediated diplomatic engagement and therefore the likelihood of de-escalation.

Each raw indicator value is mapped into a low, medium or high bucket using fixed thresholds. The model then evaluates the probability of observing that bucket under each scenario using a fixed likelihood matrix $P(b_k \mid S_i)$. These matrices are expert-calibrated objects inspired by analogue episodes, including the Iran–Iraq Tanker War, the Abqaiq attack of September 2019 and the Twelve-Day War of June 2025. They are not regression-estimated coefficients. In daily operation, the matrices remain fixed: what changes is the raw indicator value, the bucket assigned to that value, and the freshness weight applied to the indicator.

The grouping of indicators is important because several variables are not conditionally independent. Hormuz transit calls, Brent spreads and backwardation, for example, all contain information about physical oil-market stress. Treating them as independent signals would overstate the amount of information in the data. Within each evidence block, the model therefore combines indicator likelihoods as a weighted geometric mean,

$$\log L_g(S_i) = \sum_{k \in g} \widetilde{w}_{k,g} \log P(b_{k,t} \mid S_i), \tag{19}$$

where $\widetilde{w}_{k,g}$ are weights normalised within the block. The grouped likelihoods are then added to the prior in log space, with a likelihood-power discount to avoid over-confidence from noisy real-time signals and residual within-block dependence. The headline posterior uses freshness-adjusted indicator weights, while diagnostic variants report the posterior without freshness adjustment, excluding stale indicators, or excluding selected evidence blocks. These variants are robustness checks rather than alternative baselines.

The full set of nine indicators, their base within-block weights, their bucket thresholds, and the fixed likelihood matrices $P(b \mid S)$ are reported in Table B.2. For each indicator, the three rows give $P(\text{low} \mid S)$, $P(\text{medium} \mid S)$, and $P(\text{high} \mid S)$ across the four scenarios. For each indicator and scenario, the three bucket probabilities sum to one. The weights shown in the table are base weights before the freshness adjustment; in the live update they are scaled by the freshness factor described in Section 2.5. The block weight W_g used in the posterior update is the mean of the member weights within block g and is not renormalised to sum to one across blocks: $W_{\text{oil}} = 1.233$, $W_{\text{risk}} = 1.033$, $W_{\text{mil}} = 1.000$, and $W_{\text{dip}} = 1.000$.

B.2 Scenario-specific shortfall paths

The economically relevant quantity entering the price block is the net shortfall, defined as the gross supply shortfall after accounting for inventory and reserve buffers. The dashboard reports both objects separately: the gross shortfall assumed under each scenario and the buffer path used to obtain the net shortfall that enters the price equation.

B.3 Scenario-specific shortfall paths

The economically relevant quantity entering the price block is the net shortfall, defined as the gross supply shortfall after accounting for inventory and reserve buffers. The dashboard reports both objects separately: the gross shortfall assumed under each scenario and the buffer path used to obtain the net shortfall that enters the price equation.

B.3.1 Gross shortfall paths

For the first five months of 2026, the gross shortfall path is fixed using the observed early-phase disruption, informed by oil-market monitoring and the IEA Oil Market Report

$$g_t = (0, 0, 9, 13, 13) \text{ mb/d},$$

for January to May 2026. From June 2026 onward, the gross supply path under scenario s is generated by

$$g_t^s = g_{\text{May}} \chi^s F_t^s + \ell^s \left(1 - F_t^s\right), \tag{20}$$

where g_{May} is the May shortfall, χ^s is the post-May shock factor, ℓ^s is the residual long-run production loss, and F_t^s is a logistic recovery factor. Let m_t denote the number of months after May 2026 and let H^s denote the scenario-specific recovery horizon. The recovery factor is

$$F_t^s = 1 - \frac{1}{1 + \exp\left[-\frac{m_t - H^s/2}{\max(H^s/6, 0.5)}\right]}. \tag{21}$$

A larger value of H^s implies a slower normalization of supply and therefore a more persistent price response through the net-shortfall channel.

Table B.2: Indicator weights, thresholds, and fixed likelihood matrices $P(b \mid S)$

Indicator (block)	w_k	Thresholds	S_1	S_2	S_3	S_4
Hormuz transits (oil)	1.5	0.3/0.7				
P(low)			0.05	0.15	0.45	0.70
P(medium)			0.25	0.45	0.40	0.25
P(high)			0.70	0.40	0.15	0.05
Brent–WTI spread (oil)	1.0	3/8				
P(low)			0.60	0.40	0.20	0.10
P(medium)			0.30	0.40	0.40	0.30
P(high)			0.10	0.20	0.40	0.60
Brent backwardation (oil)	1.2	0/4				
P(low)			0.60	0.35	0.20	0.10
P(medium)			0.30	0.45	0.40	0.30
P(high)			0.10	0.20	0.40	0.60
War-risk premium (risk)	1.3	0.1/0.5				
P(low)			0.60	0.40	0.20	0.05
P(medium)			0.30	0.40	0.40	0.25
P(high)			0.10	0.20	0.40	0.70
Iranian rial (parallel) (risk)	0.8	0.3/0.6				
P(low)			0.55	0.40	0.25	0.10
P(medium)			0.30	0.40	0.40	0.30
P(high)			0.15	0.20	0.35	0.60
Israel GPR index (risk)	1.0	0.25/0.5				
P(low)			0.55	0.40	0.25	0.10
P(medium)			0.30	0.40	0.40	0.30
P(high)			0.15	0.20	0.35	0.60
Missile/drone tempo (mil)	1.1	0.2/0.55				
P(low)			0.60	0.40	0.25	0.10
P(medium)			0.30	0.40	0.40	0.30
P(high)			0.10	0.20	0.35	0.60
CSG positioning (mil)	0.9	0.3/0.7				
P(low)			0.55	0.40	0.30	0.15
P(medium)			0.30	0.40	0.40	0.30
P(high)			0.15	0.20	0.30	0.55
Track-1 cadence (dip)	1.0	0.2/0.55				
P(low)			0.10	0.25	0.40	0.60
P(medium)			0.30	0.45	0.40	0.30
P(high)			0.60	0.30	0.20	0.10

Note: The table reports the fixed calibration objects used in the grouped Bayesian update. For each indicator, the header row reports the base indicator weight w_k and the two bucket thresholds used to assign the raw observation to the low, medium or high bucket. The following three rows report the likelihood matrix $P(b \mid S)$. For each indicator and scenario, the probabilities of the three buckets sum to one. The likelihood matrices are expert-calibrated, inspired by analogue episodes, and are not regression-estimated coefficients. The weights reported in the table are base weights before the freshness adjustment; in the live update, stale indicators are down-weighted through the freshness factor before the grouped likelihood is computed.

Table B.3: Baseline calibration of post-May scenario dynamics

Scenario	Post-May shock factor χ^s	Recovery horizon H^s (months)	Long-run loss ℓ^s (mb/d)
Diplomatic Resolution (S1)	0.45	4	−0.8
Prolonged Stalemate (S2)	0.75	7	0.4
Partial Resumption (S3)	1.10	10	1.0
Major Escalation (S4)	1.30	14	1.8

Note: The recovery horizon is measured in months after May 2026, and the long-run loss is measured in million barrels per day (mb/d). A negative long-run loss in the Diplomatic Resolution scenario captures an overshooting normalization relative to the disrupted baseline.

B.3.2 Inventory and reserve buffer

The net shortfall entering the price block is the gross shortfall after inventory and reserve releases:

$$q_t^s = \max\{g_t^s - b_t, 0\}.$$

The baseline buffer path is front-loaded. Monthly releases are calibrated as

$$b_t = (0.0,\ 0.4,\ 1.6,\ 2.2,\ 1.8,\ 1.4,\ 1.0,\ 0.8,\ 0.6,\ 0.4,\ 0.2,\ 0.0)\ \text{mb/d},$$

from January to December 2026, and are zero thereafter unless an alternative path is supplied. This path captures a stylised coordinated inventory and strategic-reserve drawdown that is strongest during the acute phase of the disruption and then tapers as market conditions normalise. The initial months are calibrated using available information on strategic-reserve releases and commercial inventory drawdowns. From mid-2026 onward, the baseline assumes a gradual decline in buffer use. These paths can be changed in the dashboard to incorporate newly observed data or alternative beliefs about the feasible pace of inventory and reserve releases.

B.4 Price-equation calibration

The baseline log-price component is

$$lr_t^s = \frac{q_t^s / Q^{pre}}{\varepsilon},$$

where Q^{pre} is pre-war oil supply and ε is the short-run elasticity parameter. The baseline calibration sets

$$Q^{pre} = 103\ \text{mb/d}, \qquad P^{pre} = 60.75\ \text{USD/bbl}.$$

The central short-run elasticity is

$$\varepsilon = 0.14,$$

which follows the range of short-run oil-market elasticities used in Caldara, Cavallo, and Iacoviello (2019). We consider two sensitivity cases:

$$\varepsilon = 0.10 \quad \text{and} \quad \varepsilon = 0.25.$$

The low-elasticity case, motivated by estimates that place substantial mass on very inelastic short-run oil supply or demand responses (Braun, 2023) , generates a larger price response to a given shortfall. The high-elasticity case is used as a lower-price-response benchmark, consistent with the scenario exercises of Kilian et al. (2026). Throughout, ε is interpreted as the absolute value of the relevant short-run elasticity, so lower values imply larger price responses to a given physical scarcity share.

The model is anchored to the last observed Brent price. Let a denote the anchor month and P_a the observed Brent price in that month. The model computes the relative log move

$$\Delta lr_t^s = lr_t^s - lr_a^s$$

and applies a cap and floor to this relative move:

$$\widetilde{\Delta lr}_t^s = \min\left\{\max\left\{\Delta lr_t^s, \underline{c}\right\}, \overline{c}\right\}, \qquad \underline{c} = \log(0.4), \quad \overline{c} = \log(4.0).$$

The cap is applied to the relative log move from the observed anchor, rather than to the level of lr_t^s. This preserves the non-linearity of the elasticity channel while preventing mechanically extreme jumps or collapses relative to the observed Brent anchor.

B.4.1 Optional risk-premium overlay

The baseline forward paths set the geopolitical risk-premium term to zero. This choice reflects the fact that geopolitical information is already incorporated in the scenario probabilities, while the price block maps the resulting physical scarcity into Brent prices. The framework nonetheless allows an optional non-physical risk-premium overlay:

$$lr_t^s = \frac{q_t^s / Q^{pre}}{\varepsilon} + \beta_R R_t^s.$$

In the baseline, $R_t^s = 0$. When activated, R_t^s can either be specified manually by the user or generated from simple scenario-based presets. A possible preset specification is

$$R_t^s = \bar{R}^s F_t^s,$$

where F_t^s is the same logistic recovery factor used in the production block. The preset values used for sensitivity analysis are

$$\bar{R}^{S1} = 0.00, \quad \bar{R}^{S2} = 0.15, \quad \bar{R}^{S3} = 0.35, \quad \bar{R}^{S4} = 0.60,$$

with $\beta_R = 0.030$. These values are illustrative sensitivity inputs, not estimated structural parameters and not part of the baseline calibration.

B.4.2 Stock-rebuild premium

The stock-rebuild premium is introduced as a reduced-form adjustment for the post-shock cost of replenishing inventories after emergency drawdowns. It is kept separate from the contemporaneous physical scarcity channel. The latter enters the fundamental price equation through the net shortfall,

$$q_t^s = \max\{g_t^s - b_t, 0\}.$$

By contrast, the stock-rebuild premium is based on the cumulative gross deficit, since it is intended to proxy the inventory depletion generated by the disruption rather than the current scarcity after buffers have been released.

Let the cumulative gross deficit in scenario s be

$$C_t^s = \sum_{\tau \le t} \max\{g_\tau^s, 0\} \times 30, \tag{22}$$

measured in million barrels. The effective inventory gap is approximated as the cumulative gross deficit net of an assumed strategic release:

$$G_t^s = \max\{C_t^s - SPR^s, 0\}. \tag{23}$$

Once the gross shortfall has disappeared, this gap declines mechanically at the rebuilding pace κ^{reb}:

$$G_t^s = \max\left\{G_{T_s}^s - \kappa^{reb} \times 30 \times (t - T_s), 0\right\}, \qquad t > T_s, \tag{24}$$

where T_s is the final month with a positive gross shortfall. The baseline calibration sets

$$SPR^s = 400 \text{ million barrels}, \qquad \kappa^{reb} = 1 \text{ mb/d}.$$

The premium is scaled by the peak intensity of the disruption:

$$SR_t^s = \left(\phi \max_t \max\{g_t^s, 0\}\right) \frac{G_t^s}{C_T^s}, \tag{25}$$

where

$$\phi = 2.5 \quad \text{USD/bbl per mb/d of peak deficit.}$$

The parameter ϕ is a reduced-form scaling parameter rather than an estimated structural coefficient. It controls how strongly the cumulative inventory gap is translated into a temporary price-level premium. The baseline value is chosen to generate an inventory-replenishment premium of plausible magnitude relative to historical and market-based assessments of short-lived oil-supply disruptions, while leaving the contemporaneous scarcity response disciplined by the short-run elasticity ε.

The stock-rebuild premium should therefore be interpreted as a reduced-form proxy for post-shock inventory-rebuilding pressure and market tightness. It is based on the gross cumulative deficit, rather than on the contemporaneous net shortfall, because its purpose is to approximate the stock depletion generated by the physical disruption. The net shortfall enters the fundamental price equation, while the stock-rebuild premium captures the subsequent price-level pressure associated with restoring buffers.

This separation avoids double counting. The baseline price block uses the physical scarcity share and the short-run elasticity to discipline the immediate price response, while allowing persistent disruptions to affect prices through two channels: the duration of the net shortfall and the cumulative inventory gap generated by the gross deficit.

B.5 Reverse-engineering and external scenario-path mapping diagnostics

B.6 External scenario-path mapping diagnostic

The mapping from the geopolitical posterior to external scenario paths used in Section 6 and Figure 2, is a dynamic diagnostic rather than a fixed object. For each geopolitical scenario s and each external (or BMPE) Brent scenario path u, define a price-path distance

$$d_{u,s} = w_p \, \mathrm{RMSE}\Big(\log P^{(s)}_{1:H}, \, \log P^{\mathrm{ext},(u)}_{1:H}\Big) + w_q \, \mathrm{RMSE}\Big(\mathrm{share}^{(s)}_{1:H}, \, \mathrm{share}^{(u)}_{1:H}\Big), \tag{26}$$

where $P^{(s)}$ is the Brent path implied by the supply and price blocks, $P^{\mathrm{ext},(u)}$ is the external scenario path u, the RMSE is taken over the common horizon $1{:}H$ with equal horizon weights, and (w_p, w_q) weight the price and quantity-share components. The baseline is *price-only*, $w_p = 1$, $w_q = 0$, while the quantity leg is retained as a sensitivity.

Distances are transformed into scenario-conditional weights using a row-wise softmax

$$M_{u,s} = \frac{\exp(-d_{u,s}/\tau)}{\sum_{s'} \exp(-d_{u,s'}/\tau)}, \qquad \tau = 0.10, \tag{27}$$

For each external path u, the row $M_{u,\cdot}$ sums to one. Thus, $M_{u,s}$ can be read as the diagnostic weight on geopolitical scenario s conditional on external path u.

Given the geopolitical posterior π_s, the external scenario-path weight assigned to path u is

$$\omega^{\mathrm{ext}}_u = \frac{\sum_s \pi_s M_{u,s}}{\sum_{u'} \sum_s \pi_s M_{u',s}}. \tag{28}$$

Because M is recomputed from the current price paths, the mapping adapts automatically as the supply and price calibrations change; it is therefore presented as an illustrative bridge to the official scenario space, not as an independent probability estimate.

BANCO DE ESPAÑA PUBLICATIONS

OCCASIONAL PAPERS

2430 JOSE GONZÁLEZ MÍNGUEZ: The Letta report: a set of proposals for revitalising the European economy. (There is a Spanish version of this edition with the same number).
2431 MARIYA MELNYCHUK and JAVIER MENCÍA: A taxonomy of macro-financial risks and policies to address them.
2432 DMITRY KHAMETSHIN, DAVID LÓPEZ RODRÍGUEZ and LUIS PÉREZ GARCÍA: El mercado del alquiler de vivienda residencial en España: evolución reciente, determinantes e indicadores de esfuerzo.
2433 ANDRÉS LAJER BARON, DAVID LÓPEZ RODRÍGUEZ and LUCIO SAN JUAN: El mercado de la vivienda residencial en España: evolución reciente y comparación internacional.
2434 CARLOS GONZÁLEZ PEDRAZ, ADRIAN VAN RIXTEL and ROBERTO PASCUAL GONZÁLEZ: Navigating the boom and bust of global SPACs.
2435 PATROCINIO TELLO-CASAS: El papel de China como acreedor financiero internacional.
2436 JOSÉ RAMÓN MARTÍNEZ RESANO: CBDCs, banknotes and bank deposits: the financial stability nexus.
2501 PEDRO DEL RÍO, PAULA SÁNCHEZ, MARÍA MÉNDEZ, ANTONIO MILLARUELO, SUSANA MORENO, MANUEL ROJO, JACOPO TIMINI and FRANCESCA VIANI: La ampliación de la Unión Europea hacia el este: situación e implicaciones para la economía española y la Unión Europea.
2502 BANCO DE ESPAÑA: In-person access to banking services in Spain. 2024 monitoring report. (There is a Spanish version of this edition with the same number).
2503 ANDRÁS BORSOS, ADRIAN CARRO, ALDO GLIELMO, MARC HINTERSCHWEIGER, JAGODA KASZOWSKA-MOJSA and ARZU ULUC: Agent-based modeling at central banks: recent developments and new challenges.
2504 ANDRES ALONSO-ROBISCO, ANDRES AZQUETA-GAVALDON, JOSE MANUEL CARBO, JOSE LUIS GONZALEZ, ANA ISABEL HERNAEZ, JOSE LUIS HERRERA, JORGE QUINTANA and JAVIER TARANCON: Empowering financial supervision: a SupTech experiment using machine learning in an early warning system.
2505 JÉSSICA GUEDES, DIEGO TORRES, PAULINO SÁNCHEZ-ESCRIBANO and JOSÉ BOYANO: Incertidumbre en el mercado de bonos: una propuesta para identificar sus narrativas con GDELT.
2506 LAURA JIMENA GONZÁLEZ GÓMEZ, FERNANDO LEÓN, JAIME GUIXERES PROVINCIALE, JOSÉ M. SÁNCHEZ and MARIANO ALCAÑIZ: Evolución de la investigación neurocientífica del efectivo: revisión y perspectivas actuales.
2507 LUIS FERNÁNDEZ LAFUERZA, IRENE ROIBÁS and RAQUEL VEGAS SÁNCHEZ: Indicadores de desequilibrios de precios del mercado inmobiliario comercial.
2508 PANA ALVES and OLIVIER HUBERT: ¿Influye la eficiencia energética en el precio de la vivienda en España?.
2509 ALEJANDRO FERRER and ANA MOLINA: The interaction of liquidity risk and bank solvency via asset monetisation mechanisms. (There is a Spanish version of this edition with the same number).
2510 ISABEL ALCALDE and PATRICIA STUPARIU: Financial education at an early age. (There is a Spanish version of this edition with the same number).
2511 ALEJANDRO GONZÁLEZ FRAGA, AITOR LACUESTA GABARAIN, JOSÉ MARÍA LABEAGA AZCONA, MARÍA DE LOS LLANOS MATEA ROSA, SOLEDAD ROBLES ROMERO, MARÍA VALKOV LORENZO and SERGIO VELA ORTIZ: Estructura del mercado de electrolineras.
2512 FERNANDO ARRANZ GOZALO, CLARA I. GONZÁLEZ MARTÍNEZ and MERCEDES DE LUIS LÓPEZ: Sovereign assets and sustainable and responsible investment: the importance of climate metrics. (There is a Spanish version of this edition with the same number).
2513 IRMA ALONSO-ÁLVAREZ and DANIEL SANTABÁRBARA: Decoding Structural Shocks in the Global Oil Market.
2514 JOSÉ MANUEL CARBÓ, CLAUDIA TOLEDO and ÁNGEL IVÁN MORENO: Hacia un diccionario panhispánico de sentimiento de la estabilidad financiera.
2515 IGNACIO FÉLEZ DE TORRES, CLARA I. GONZÁLEZ MARTÍNEZ and ELENA TRIEBSKORN: The puzzle of forward-looking climate transition risk metrics.
2516 MIGUEL GARCÍA-POSADA: Un análisis de los efectos de la introducción del procedimiento especial de insolvencias para microempresas sobre la propensión a concursar.
2517 ADRIAN VAN RIXTEL: Whatever it takes? Economic policymaking in China in the context of a possible deflationary spiral.
2518 PATRICIA STUPARIU and JUAN RAFAEL RUIZ: Adding up the benefits: education, numeracy and financial literacy. (There is a Spanish version of this edition with the same number).
2519 DAVID CUBERES, AITOR LACUESTA, MARÍA DE LOS LLANOS MATEA and DANIEL OTO-PERALÍAS: El efecto de la regulación sobre el tamaño de las plantas fotovoltaicas.
2520 MARINA GARCÍA GIL and DMITRY KHAMETSHIN: Ayudas directas de la línea COVID-19: los determinantes de la asignación y el efecto sobre el crédito bancario.
2521 LUCÍA CUADRO-SÁEZ, CORINNA GHIRELLI, MAXIMILIANO MORENO-LÓPEZ and JAVIER J. PÉREZ: Monitoring and forecasting food prices in the euro area.

2522 LAURA HOSPIDO, JÚLIA MARTÍ LLOBET and CARLOS SANZ: ¿Qué políticas son efectivas para reducir la exclusión social? Evaluación de cinco proyectos piloto de inclusión social a través de ensayos aleatorizados.

2523 MARIO ALLOZA, MARÍA ELENA CRISTÓBAL RODRÍGUEZ, JULIA GARCÍA-ROYO DÍAZ, BEATRIZ GONZÁLEZ, ALBERTO MARTÍN DEL CAMPO SOLA, ANE MARTÍN UGARTE, ENRIQUE MORAL-BENITO and IRENE PINILLA MELGAREJO: Diagnóstico y consecuencias económicas del grado de competencia en las licitaciones públicas.

2524 ADRIÁN CARRO, JORGE E. GALÁN, ENRIC MARTORELL and RAQUEL VEGAS: A literature review on ex-ante and ex-post analysis of the implications of borrower-based macroprudential measures.

2525 PAULA SEMPERE and RAQUEL PANTOJA: Efectivo y *neuromarketing:* Concepto, funciones y aplicación.

2601 RODOLFO G. CAMPOS, JACOPO TIMINI, FRANCESCA VIANI and ELENA VIDAL: The EU-Mercosur agreement: analysis of its characteristics from a sectoral perspective. (There is a Spanish version of this edition with the same number).

2602 PABLO AGUILAR, CORINNA GHIRELLI and SAMUEL HURTADO: MTBE v2025: new version of the Quarterly Model of the Banco de España.

2603 IRMA ALONSO-ALVAREZ, EKATERINA BUKINA, MARINA DIAKONOVA, NINO KHITARISHVILI, JAVIER J. PÉREZ and PEDRO PIQUERAS: Geopolitical risk: a database of general and bilateral indices.

2604 RUBÉN DOMÍNGUEZ-DÍAZ, MARTA GARCÍA-RODRÍGUEZ, JAVIER QUINTANA and RUBÉN VEIGA-DUARTE: Estimación del crecimiento potencial de la economía española: una revisión metodológica.

2605 CARLES MANERA, FERRAN NAVINÉS, JAVIER FRANCONETTI, MIQUEL QUETGLAS and JOSÉ A. PÉREZ-MONTIEL: Regional outlook on labor productivity in Spain, 2000–2022.

2606 RICARDO BARAHONA: Current trends in performance and flows in the Spanish pension funds industry.

2607 FRANCISCO GONZÁLEZ RODRÍGUEZ, ALBERTO ORTS TORRES, JOSÉ MARÍA SERENA GARRALDA and MIQUEL TARÍ SÁNCHEZ: Issuance yield of MREL-eligible bank debt: a benchmarking model.

2608 MÓNICA CORREA-LÓPEZ, MAR DELGADO-TÉLLEZ and MARTA SUÁREZ-VALERA: Europa en transición energética: respuestas y desafíos tras dos crisis consecutivas.

2609 AITOR LACUESTA, MARÍA VALKOV and MICOLE DE VERA: Un análisis de los mecanismos de protección de los operadores económicos establecidos por la Ley de garantía de la unidad de mercado.

2610 Encuesta Financiera de las Familias (EFF) 2024: métodos, resultados y cambios desde 2022.

2611 ALFONSO CAMBA, MATÍAS COVARRUBIAS, JOSÉ LUÍS RODRÍGUEZ and RAQUEL TÁRREGA: El coste en bienestar de los márgenes en mercados de productos en España: un análisis con empresas heterogéneas.

2612 MICAELA ARIAS, LUCÍA CRIADO, JAVIER GARCÍA-VERDUGO, EDUARDO GUTIÉRREZ, ALMUDENA KESSLER, JOSÉ MANUEL MONTERO, SARA PEREIRA and PAU ROLDÁN-BLANCO: Evolución de la estructura de mercado y de los indicadores de competencia en España en las dos últimas décadas.

2613 ANDRÉS ALONSO-ROBISCO, JOSÉ MANUEL CARBÓ, CARLOS JOSÉ GARCÍA, JORGE QUINTANA and JAVIER TARANCÓN: Retriever or reasoner? Decomposing retrieval-augmented generation performance in external audit supervision.

2614 RUBÉN VEIGA-DUARTE, EDUARDO GUTIÉRREZ, ALBERTO MARCOS-FERNÁNDEZ, ENRIQUE MORAL-BENITO and JAVIER QUINTANA: Las fuentes del crecimiento en España y las principales economías europeas: productividad total de los factores y heterogeneidad sectorial.

2615 MATEO CACHO UZAL, M. DANIELA FILIP y DIANA POSADA RESTREPO: Dinámicas de convergencia regional en las principales economías europeas.

2616 IRMA ALONSO-ALVAREZ, IVÁN KATARYNIUK, DAVID LÓPEZ-SALIDO and JAVIER J. PÉREZ: Geopolitical Oil-Supply Scenarios and Macroeconomic Risks: A Bayesian Toolkit for Policy Analysis.