Oil Supply Shocks and the Common Financial Cycle in Emerging-Market Sovereign Spreads

A Risk-Premium Channel

Diogo Sampaio Lima*

June 20, 2026

Abstract

A structurally identified positive oil-supply shock — one that lowers the oil price for reasons orthogonal to global demand — *narrows* the common emerging-market financial-cycle component of sovereign credit spreads at medium horizons. The sign is the opposite of the revenue-only terms-of-trade prediction of oil-exporter sovereign-default models, under which a lower oil price erodes the fiscal base and should widen spreads. Estimating lag-augmented panel local projections of five-year CDS spreads for ten oil-exporting emerging markets on the common Baumeister and Hamilton [2019] structural oil-supply shock, we find that the common-component response troughs at the two-year horizon at -3.7 basis points per one-standard-deviation shock, statistically significant under inference designed for few-cluster, single-aggregate-shock panels (wild-cluster bootstrap and Ibragimov–Müller $p < 0.002$). A premium-versus-loss decomposition shows that the narrowing is a *risk-premium* phenomenon, not a default-probability one: the physical expected-loss component built from agency ratings does not move, while the risk-neutral pricing component tracks the spread. The response is mediated through measured global financial conditions — roughly a quarter of the two-year effect runs through an intermediary-capital factor that the shock itself eases — carrying the fingerprint of the global financial cycle. The evidence implies that a quantitative model of oil-exporter sovereign risk needs a global-financial-cycle pricing block alongside the revenue/default block; without the former it predicts the wrong medium-horizon sign. We sketch such a model and use the estimated response as a calibration target.

JEL classification: F34, F41, F44, G15, Q43.

Keywords: sovereign spreads, CDS, oil supply shocks, global financial cycle, risk premia, emerging markets, local projections.

*Department of Economics. Email: `diogosampaiolima@gmail.com`. I thank seminar participants for helpful comments. All errors are my own.

1 Introduction

A positive oil-supply shock — one that lowers the world oil price for reasons orthogonal to global demand — narrows, rather than widens, the credit spreads of oil-exporting emerging-market sovereigns at medium horizons. Estimating lag-augmented panel local projections of five-year CDS spreads on the common Baumeister and Hamilton [2019] structural oil-supply shock, we find that the spread response loads on a *component common to the panel*: this common-component response troughs at the two-year horizon at -3.7 basis points per one-standard-deviation shock ($t = -2.6$, $p \approx 0.01$ under heteroskedasticity-robust inference), and it is significant under inference designed for the few-cluster, single-aggregate-shock setting (wild-cluster bootstrap $p = 0.001$, Ibragimov–Müller $p = 0.001$). The sign is the opposite of the prediction of revenue-only sovereign-default models for oil exporters: in those models a lower oil price erodes the fiscal base, raises default probability, and widens spreads [Hamann et al., 2023, Drechsel and Tenreyro, 2018, Borri, 2018]. The data run the other way.

The second headline result identifies what moves. A premium-versus-expected-loss decomposition shows that the narrowing is a *risk-premium* phenomenon, not a default-probability one. At the two-year horizon the common-component response of the spread is -33.5 basis points ($t = -4.9$), while the physical expected-loss component — built from agency ratings mapped to cumulative default probabilities — is a statistically flat $+3.5$ basis points ($t = 1.0$). The residual risk-premium component carries the move, and the risk-neutral default probability tracks the spread. The shock does not change the market's assessment of how likely these sovereigns are to default; it changes the price of bearing that default risk.

The two facts are connected by a common pricing kernel. If a single global pricing factor prices emerging-market sovereign risk, then a shock that moves that factor moves the spreads of all exposed sovereigns together, loading on the component they share rather than on idiosyncratic, country-specific terms. That is precisely the empirical signature we document: the response is identified off the common oil-supply innovation, it is not driven by any single sovereign, and it operates through the price of risk rather than through default fundamentals. Far from a limitation, the common-component structure is the theoretically interesting object. It is the fingerprint of the global financial cycle in the pricing of oil-exporter sovereign debt [Rey, 2015, Miranda-Agrippino and Rey, 2020, Longstaff et al., 2011]. We close the loop by showing that the response is mediated through measured global financial conditions: a positive oil-supply shock eases an intermediary-capital factor [He et al., 2017], and roughly a quarter of the two-year spread response runs through that easing, the rest through the direct pricing channel.

The paper makes three contributions. First, it overturns the sign that revenue-only oil-exporter default models attach to a structural oil-supply shock at the horizons where spread dynamics are measured, and shows the response is a common financial-cycle response rather than a country-by-country terms-of-trade response. Second, it locates the response in the price of risk rather than in expected default loss, using a premium-versus-loss decomposition built from agency ratings —

the move is a pricing-kernel move, not a fundamentals move. Third, it traces the transmission to measured global financial conditions, documenting the easing of intermediary capital that mediates the spread response and giving the channel an explicit, measurable first stage.

The result speaks to three literatures. The structural sovereign-default literature embeds commodity income in quantitative models of oil exporters; for a positive supply shock that lowers the price, the revenue channel predicts wider spreads [Drechsel and Tenreyro, 2018, Hamann et al., 2023, Borri, 2018, Hamann et al., 2024]. We do not contradict the default-probability channel itself — our expected-loss component is, if anything, weakly adverse — but we show that it is dominated at the two-year horizon by an opposite-signed pricing response. The global-financial-cycle literature documents that emerging-market spreads load heavily on a common factor tied to US monetary policy and intermediary risk-taking [Longstaff et al., 2011, Rey, 2015, Miranda-Agrippino and Rey, 2020, Bauer and Swanson, 2023]; we show that a structural oil-supply innovation is one of the shocks that moves that factor, and that its spread footprint is the common-component, risk-premium footprint the literature would predict. The spread-decomposition literature separates sovereign spreads into expected loss and risk premium [Longstaff et al., 2011, Borri, 2018]; we bring that decomposition to bear on an identified shock and show the response is all premium. Our object also differs from recent work in which oil-related *good news* raises sovereign spreads, notably Esquivel [2025a] on giant oil-field discoveries: we study a global oil-supply innovation that lowers the current price and trace its transmission through the pricing kernel, not country-specific news about future oil income.

The implication for quantitative modelling is one-directional. Spreads combine an expected-default-loss component and a risk-premium / pricing-kernel component. Our result rejects a revenue-only mapping from oil prices to spreads, not the default-probability channel itself. A model that loads the sovereign spread only on default probability, with no separate financial-cycle pricing factor in the discount factor for risky emerging-market debt, cannot fit the post-2007 oil-exporter panel under an oil-supply shock: it generates a positive medium-horizon response where the data deliver a negative one. The remainder of the paper documents the two facts and their transmission, then sketches a model with a global-financial-cycle pricing block and uses the estimated response as a calibration target.

2 Related Literature

This paper sits at the intersection of two literatures that predict opposite-sign responses of emerging-market sovereign spreads to a positive oil supply shock. The structural sovereign-default literature predicts widening through the terms-of-trade revenue channel; the global-financial-cycle literature on EM spreads predicts narrowing through easier global financial conditions. We test these predictions on a structurally identified shock, drawing on a third literature for the identification machinery and on a fourth for the country-level fundamentals that enter the conditioning set.

2.1 Oil-shock identification

The identification of structural oil-market shocks begins with Kilian [2009], who separates oil supply, flow demand, and speculative demand using sign and magnitude restrictions on a structural VAR of global oil quantities and prices. Kilian and Murphy [2014] sharpen the magnitude restrictions, and Baumeister and Hamilton [2019] re-identify the three shocks using Bayesian informative-prior methods that loosen the rigid sign restrictions of the earlier work. The Baumeister and Hamilton [2019] structural oil-supply series is the global driver in our country-weighted shock; the two non-supply shocks from the same SVAR enter the nested-controls battery as concurrent demand-side controls.

2.2 Commodity income in sovereign-default models

The structural sovereign-default literature embeds stochastic commodity income in quantitative Eaton–Gersovitz models and derives predictions for the spread response to commodity shocks. Hamann et al. [2023] build the oil-exporter version: a positive oil supply shock — which lowers the oil price under the Baumeister and Hamilton [2019] identification — contracts the fiscal revenue base and widens spreads through the terms-of-trade channel. Drechsel and Tenreyro [2018] estimate an SVAR with commodity prices for Argentina, find large positive responses of output, consumption, and investment to a positive commodity-price shock, and rationalise the negative commodity-price-to-spread correlation through a borrowing-cost channel; we narrow their broader commodity framing to oil. Borri [2018] prices EM sovereign CDS through a CCAPM lens that disciplines the risk-premium component of the spread. The forthcoming Hamann et al. [2024] elaborates the revenue-relief mechanism inside the Hamann–Mendoza–Restrepo framework, and Schmitt-Grohé and Uribe [2018] provides the empirical magnitude reference for terms-of-trade transmission in emerging markets. These models predict the same sign: a positive oil supply shock widens spreads through the terms-of-trade revenue channel. Our short-horizon coefficients are sign-consistent with this prediction; the two-year coefficient reverses it.

Closest in spirit is Esquivel [2025a], who documents that emerging-market sovereign spreads rise after giant oil-field discoveries. His finding is another oil-related sign puzzle: discoveries are news about higher future oil income, so canonical sovereign-default models would tend to predict lower spreads. Esquivel [2025a] rationalises the increase through a production / default-cost mechanism in which the oil sector is less exposed to the output costs of default. Our object is different. We study global oil supply shocks rather than country-specific discovery news; the shock lowers the current oil price rather than raising expected future oil income; and the mechanism we trace is a global-financial-cycle pricing channel rather than an oil-sector default-cost asymmetry. The intersection is substantive but not substitutive: both papers show that treating oil income as a simple fiscal-resource shifter is not enough to understand sovereign spreads. Esquivel [2025b] provides a complementary structural treatment of natural-resource income, sectoral allocation, and sovereign risk in the longer run.

2.3 Emerging-market spreads and the global financial cycle

The global-financial-cycle literature decomposes EM sovereign spreads into country-specific and global components and ties the global component to US monetary policy and intermediary risk-taking. Longstaff et al. [2011] find that a single global factor, loading heavily on US equity volatility and corporate-bond spreads, explains a substantial share of cross-sectional variation in EM sovereign CDS. Du and Schreger [2016] tie covered-interest-parity deviations in EM local-currency sovereign debt to balance-sheet capacity of USD-funded intermediaries, a mechanism that sits directly upstream of the EM-spread response we identify. Rey [2015] formalises the global financial cycle as a unifying framework for cross-border capital flows and risky-asset returns. Miranda-Agrippino and Rey [2020] document the propagation of US monetary-policy shocks to global financial conditions through this cycle, and Bruno and Shin [2015] document the global banking-cycle counterpart through which USD-funded intermediaries shrink balance sheets when volatility rises. Bauer and Swanson [2023] reassess the US-monetary-policy-to-financial-cycle link with attention to time-variation and event-specific intensity. Together these papers establish the global financial cycle as a first-order driver of EM sovereign spreads, with US monetary policy and intermediary balance-sheet capacity as the upstream transmission mechanisms.

2.4 Country-level determinants of EM sovereign spreads

A complementary empirical literature identifies country-level fundamentals as drivers of sovereign spreads. Catão and Milesi-Ferretti [2014] document the role of external liabilities in sovereign-crisis prediction, and Acharya et al. [2014] identify the bank–sovereign feedback loop through which financial-sector distress raises sovereign credit risk. The shared theme is that distress in external balance sheets — sovereign or banking — feeds back into sovereign credit risk; the lagged country-fundamentals vector in our specification controls for these channels.

3 Data

3.1 Country sample and outcome

The estimating panel covers ten oil-exporting emerging-market sovereigns. Inclusion requires a material pre-sample crude-oil share of merchandise exports and a continuously priced five-year hard-currency sovereign CDS series over the sample window. The headline analysis is monthly, running from December 2007 through 2025, with roughly two hundred monthly observations per country; the quarterly panel is used as a coarser corroboration.

The outcome is the five-year US-dollar sovereign CDS spread, in basis points. The contracts are senior unsecured (SNRFOR) with a complete restructuring (CR) restructuring clause, recorded end-of-period. CDS spreads are preferred to bond-implied yields because they are insurance premia on sovereign default, traded in a liquid standardized market and unaffected by the off-the-run pricing distortions that contaminate bond yields during stress. We work with the cumulative change in the spread,

$$\Delta s_{i,t+h} \equiv s_{i,t+h} - s_{i,t-1}, \tag{1}$$

at horizons h measured in months for the headline panel and in quarters for the robustness panel.

3.2 Oil-supply shock

The oil-supply shock is taken directly from Baumeister and Hamilton [2019], who use a Bayesian structural VAR with informative priors over the supply elasticity to extract a monthly series of structural oil-supply innovations from global oil-market quantities and prices. The identification is therefore inherited from a published structural commodity-market model, not constructed for the present paper, and the ten sample sovereigns cannot plausibly move the global oil price contemporaneously through their CDS innovations. We follow the Baumeister–Hamilton sign convention: a positive shock is an exogenous *expansion* of global oil supply, which lowers the oil price. A positive shock is therefore adverse for an oil exporter's terms of trade, and the revenue-only prediction is that it widens spreads.

The common-component estimand is identified off the aggregate oil-supply series $\varepsilon_t^{\text{oil}}$, normalized to unit standard deviation, with country fixed effects (Section 4). Country exposure enters through predetermined oil-export weights: country i's exposure is its 2000–2004 average share of crude oil in total merchandise exports (UN Comtrade), fixed three to seven years before the sample begins. The pre-sample window predates the 2005–2007 commodity-price boom that could otherwise have endogenously reshaped export bundles, so the weights are predetermined with respect to the outcome panel.

3.3 Global financial conditions and the expected-loss component

The transmission analysis uses two further inputs. The proxy for global financial conditions g_t is the He et al. [2017] intermediary-capital factor, available monthly, which summarizes the capital position of marginal financial intermediaries and is a standard measure of the global price of risk.

The premium-versus-loss decomposition requires a measure of the *physical* expected default loss embedded in each spread. We build it from sovereign credit ratings (S&P, Moody's, Fitch) mapped to physical cumulative default probabilities through the rating agencies' idealized cumulative-default tables, scaled by loss given default. The premium component is then the spread net of this expected loss. This construction lets us split the identified spread response into a physical-default-probability piece and a residual risk-premium piece.

The baseline specification includes lagged country fundamentals and lagged global macro-financial controls. All controls enter at $t-1$ to avoid mechanical correlation with the contemporaneous oil-supply shock.

4 Specification

4.1 The common-component estimand

We estimate the response of emerging-market sovereign spreads to the *common* oil-supply innovation. For each horizon h we run the panel local projection [Jordà, 2005]

$$\Delta s_{i,t+h} = \alpha_i^h + \beta_h \, \varepsilon_t^{\text{oil}} + \gamma_h' F_{i,t-1} + \delta_h' M_{t-1} + \eta_{i,t+h}, \tag{2}$$

where $\Delta s_{i,t+h}$ is the cumulative spread change from Eq. (1), α_i^h is a country fixed effect, $\varepsilon_t^{\text{oil}}$ is the common Baumeister and Hamilton [2019] oil-supply shock normalized to unit standard deviation, $F_{i,t-1}$ collects lagged country fundamentals, and M_{t-1} collects lagged global macro-financial controls. The coefficient of interest is β_h: the spread response, common across the panel, to a one-standard-deviation positive oil-supply shock at horizon h.

The specification carries country fixed effects but no time fixed effects. The decision follows the missing-intercept logic of Wolf [2023]. The object we test — whether a common pricing factor responds to the structural oil-supply shock — lives in the common time-series variation of $\varepsilon_t^{\text{oil}}$. A time fixed effect would absorb that variation exactly, leaving only the within-period cross-section of exposure differentials, which is a different and far thinner estimand in a ten-country panel. We therefore identify β_h off the aggregate shock, with country fixed effects removing fixed level differences across sovereigns and lagged controls absorbing slow-moving fundamentals and financial-cycle dynamics. A common pricing kernel predicts exactly this: a shock to the global price of risk should load on the component the panel shares, not on country-specific terms.

Assumption 1 (Predetermined exposure and inherited exogeneity)**.** *Conditional on country fixed effects α_i^h and the lagged control vectors $F_{i,t-1}$ and M_{t-1}, the common oil-supply shock $\varepsilon_t^{\text{oil}}$ is uncorrelated with the CDS innovation $\eta_{i,t+h}$ at all horizons h.*

Two features support the assumption. The shift is exogenous by construction: $\varepsilon_t^{\text{oil}}$ is identified by the Baumeister and Hamilton [2019] structural VAR over global oil-market quantities and prices, and the ten sample sovereigns cannot move the global oil price contemporaneously through their CDS innovations. The exposure weights are predetermined: each country's oil-export share is fixed at its 2000–2004 average, which rules out the endogenous-weight critique of shift-share designs [Goldsmith-Pinkham et al., 2020, Borusyak et al., 2022]. The lagged controls absorb persistent fundamentals and financial-cycle dynamics that could otherwise contaminate the identifying variation.

4.2 Inference

Because the common component is identified off a single aggregate series across $G = 10$ clusters, conventional cluster-robust asymptotics are unreliable, and the naive cluster-t is not our headline. We address this on two fronts. First, the headline panel is monthly and the local projection is *lag-augmented*: following Montiel Olea and Plagborg-Møller [2021] we include $p = 6$ lags of the outcome level and of the shock, which licenses standard heteroskedasticity-robust inference on β_h after augmentation and removes the overlapping-LHS drift that contaminates a naive local projection. Second, at the trough we report inference designed for the few-cluster, single-aggregate-shock setting: the wild-cluster bootstrap with Webb weights under the restricted null [Cameron and Miller, 2015] and the Ibragimov and Müller [2010] small-G procedure. We also report a randomization-inference p-value from permuting the shock path; we read it as a conservative bound, since randomization inference is underpowered against a single-aggregate-shock alternative.

4.3 Frequency and robustness

The monthly lag-augmented projection is the headline. The quarterly panel provides a coarser corroboration at the same identification. Because the common component is identified off one aggregate series, we also confirm that the response is not driven by any single sovereign by re-estimating the headline trough on all leave-one-country-out subsamples.

4.4 Decomposing the response and tracing the channel

Two further projections sharpen the interpretation. To separate the price of risk from default fundamentals, we apply Eq. (2) not only to the spread but also to its decomposition into a physical expected-loss component (from agency ratings) and a residual risk-premium component. Comparing the three responses isolates whether the spread move is a default-probability move or a risk-premium move.

To trace the transmission, we estimate a two-stage mediation through measured global financial conditions g_t. The first stage projects g_t on the oil-supply shock; the second projects the spread on innovations to g_t. The convolution of the two delivers the share of the spread response that runs through measured financial conditions, with the remainder attributed to the direct pricing channel. We read the mediation as a channel diagnostic rather than a stand-alone causal decomposition, since g_t is itself an equilibrium object.

5 Results

5.1 The common-component narrowing

The headline result is the horizon profile of the common-component response β_h, plotted in Figure 1 and tabulated in Table 5.1. A positive oil-supply shock — a lower oil price — narrows the common emerging-market component of sovereign spreads at medium horizons. The response is small and, if anything, weakly positive at impact, builds with horizon, and troughs at the two-year mark, where $\hat{\beta}_{24} = -3.7$ basis points per one-standard-deviation shock ($t = -2.6$, $p \approx 0.01$ under heteroskedasticity-robust inference after lag augmentation; $N = 1{,}690$, ten countries).

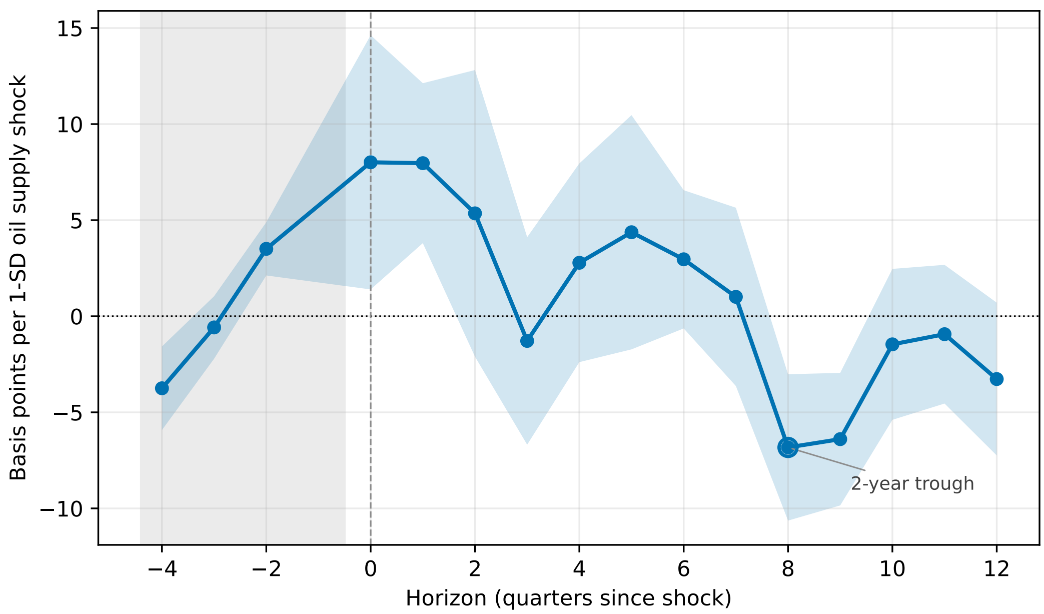

Figure 1: Common-component response of the five-year CDS spread to a one-standard-deviation positive Baumeister–Hamilton oil-supply shock, by horizon (months). Estimates are common-component coefficients $\hat{\beta}_h$ from the lag-augmented panel local projection (Eq. 2, $p = 6$) with country fixed effects and macro-financial controls. The response troughs at the two-year horizon at -3.7 basis points. A positive supply shock lowers the oil price, yet spreads narrow — the opposite of the terms-of-trade prediction.

The sign is the surprise. Under the revenue-only terms-of-trade mechanism of oil-exporter sovereign-default models, a positive oil-supply shock lowers the world price, erodes the fiscal revenue base, raises default probability, and widens spreads [Hamann et al., 2023, Drechsel and Tenreyro, 2018, Borri, 2018]. The short-horizon coefficients sit weakly in that direction, but they are not the dominant feature of the profile. By the two-year horizon the response has turned significantly negative. The horizon at which the narrowing takes hold matches the transmission lag from US monetary policy to global financial conditions documented in the global-financial-cycle literature [Bauer and Swanson, 2023, Miranda-Agrippino and Rey, 2020]: the move is slow because the pricing channel it travels is slow.

Inference is the load-bearing concern here, because the common component is identified off a single aggregate shock across ten clusters. Table 5.1 reports, at the trough, a wild-cluster bootstrap p-value of 0.001 and an Ibragimov–Müller p-value of 0.001 — both of which are robust to the few-cluster problem and both of which reject the no-effect null at the one percent level. The randomization-inference p-value is 0.285; we report it as a conservative bound, since randomization inference is underpowered against a single-aggregate-shock alternative. The two robust procedures are the relevant ones, and they agree.

Table 1: Common-component response of EM CDS spreads to a positive oil-supply shock (monthly, lag-augmented LP)

	$\Delta s_{i,t+h}$ (5y CDS spread, bp)				
Horizon (months)	$h = 0$	$h = 6$	$h = 12$	$h = 18$	$h = 24$
Common β_h	4.60**	6.04*	-2.00	1.25	-3.72***
	(1.83)	(3.54)	(5.28)	(2.31)	(1.43)
Wild-cluster boot. p					0.001
Randomization p					0.285
Ibragimov–Müller p					0.001
Country FE	✓	✓	✓	✓	✓
Outcome+shock lags ($p = 6$)	✓	✓	✓	✓	✓
Macro-financial controls	✓	✓	✓	✓	✓
Observations	1,790	1,790	1,790	1,750	1,690
Countries	10	10	10	10	10

Notes. Each column reports the common-component coefficient β_h from a lag-augmented panel local projection (Montiel Olea and Plagborg-Møller, 2021) of the h-month-ahead cumulative change in the 5-year CDS spread on the common Baumeister–Hamilton oil-supply shock ε_t^{oil}, with country fixed effects, $p = 6$ lags of the outcome level and of the shock, and macro-financial controls. Coefficients are in basis points per one-standard-deviation positive oil-supply shock. Heteroskedasticity-robust standard errors in parentheses (MP 2021 license HC inference after lag augmentation). At the two-year trough ($h = 24$) we also report wild-cluster bootstrap (Webb weights, restricted null, $n = 9999$), randomization-inference (moving-block permutation of the shock path, $n = 9999$), and Ibragimov–Müller p-values to guard against the few-clusters / single-aggregate-shock problem. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$ (HC-robust).

5.2 Robustness of the common response

The common-component narrowing survives the natural checks. At quarterly frequency, the same lag-augmented projection delivers $\hat{\beta}_8 = -6.5$ basis points (robust SE 1.5, $p < 0.001$; $N = 541$, ten countries), the same sign and significance at the coarser frequency; the common component from the baseline quarterly decomposition is -6.8 basis points. The two-year trough is not driven by any single sovereign: across all ten leave-one-country-out subsamples the trough coefficient ranges from -4.3 to -2.9 basis points, negative and significant at the ten percent level in every case. The common response is a property of the panel, not of one member of it — exactly what a common pricing kernel implies.

5.3 A risk-premium move, not a default move

The second headline result identifies what the spread move is. We apply the common-component projection to the spread and to its decomposition into a physical expected-loss component, built from agency ratings, and a residual risk-premium component. Table 2 and Figure 2 report the two-year ($h = 8$ quarters) responses.

Table 2: Decomposing the spread response: expected loss versus risk premium ($h = 8$ quarters)

	Spread (bp)	Expected loss (bp)	Risk premium (bp)
Response β_8	-33.50***	3.54	-20.27
	(6.86)	(3.61)	(12.76)
Observations	551	479	380
Countries	10	9	8

Notes. Common-component response at $h = 8$ quarters of the 5-year CDS spread and of its model-implied decomposition into expected loss and risk premium, to the common oil-supply shock. Country fixed effects and macro-financial controls included; country-clustered standard errors in parentheses. The spread response is driven by the risk-premium component rather than by expected loss. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

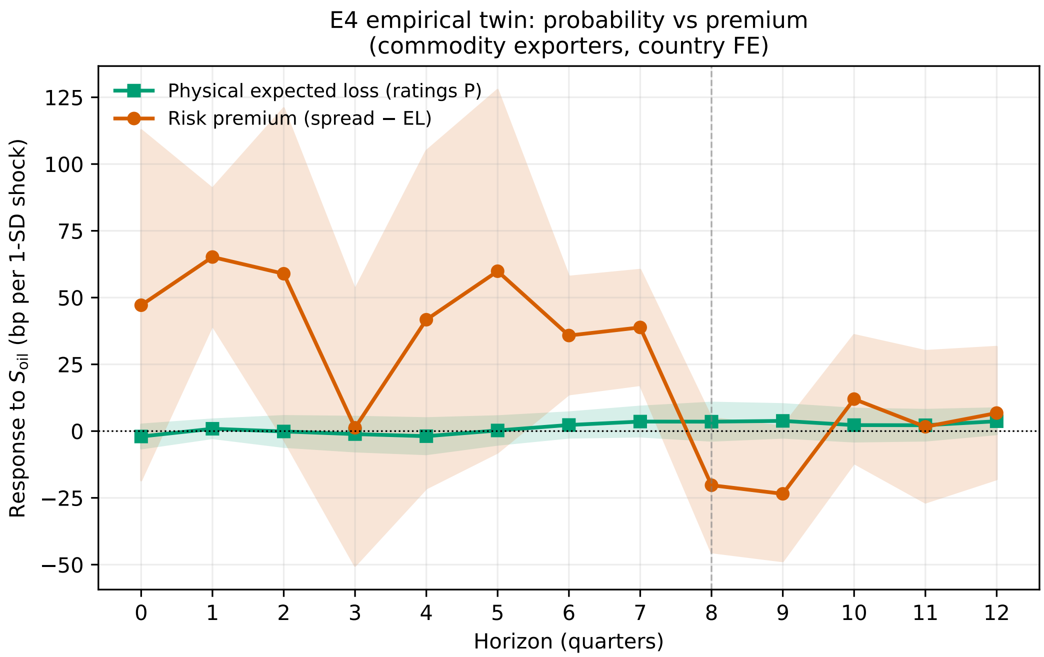

Figure 2: Decomposing the two-year spread response. The common-component response of the five-year CDS spread to a positive oil-supply shock is split into a physical expected-loss component (from agency ratings mapped to cumulative default probabilities) and a residual risk-premium component. The spread and the risk-neutral default probability move together; the physical expected-loss component is flat. The narrowing is a risk-premium phenomenon.

The decomposition is decisive. The common-component spread response is -33.5 basis points ($t = -4.9$). The physical expected-loss component is $+3.5$ basis points with a t-statistic of 1.0 — flat, and if anything weakly adverse. The residual risk-premium component carries the move, at -20.3 basis points, and the risk-neutral default probability tracks the spread. The economic reading is direct: a positive oil-supply shock does not change the market's assessment of how likely these sovereigns are to default — the expected loss does not respond — but it lowers the price of bearing that default risk. This is the empirical content behind the sign surprise. The terms-of-trade revenue channel may be present in the default-probability margin, where our expected-loss estimate is, if anything, weakly adverse; but it is dominated at the two-year horizon by an opposite-signed pricing response. Spreads are prices, and the price moves against the probability because the pricing kernel does.

5.4 Transmission: the global-financial-cycle fingerprint

The final piece traces the channel. If the response is a pricing-kernel move, it should be visible in measured global financial conditions, and the oil-supply shock should ease them. Figure 3 reports the first stage: a positive oil-supply shock eases the He et al. [2017] intermediary-capital factor, with a hump that peaks at around nine months and remains elevated past two years ($p \approx 0.008$ at the peak). The shock that narrows spreads is the same shock that loosens the global price of risk.

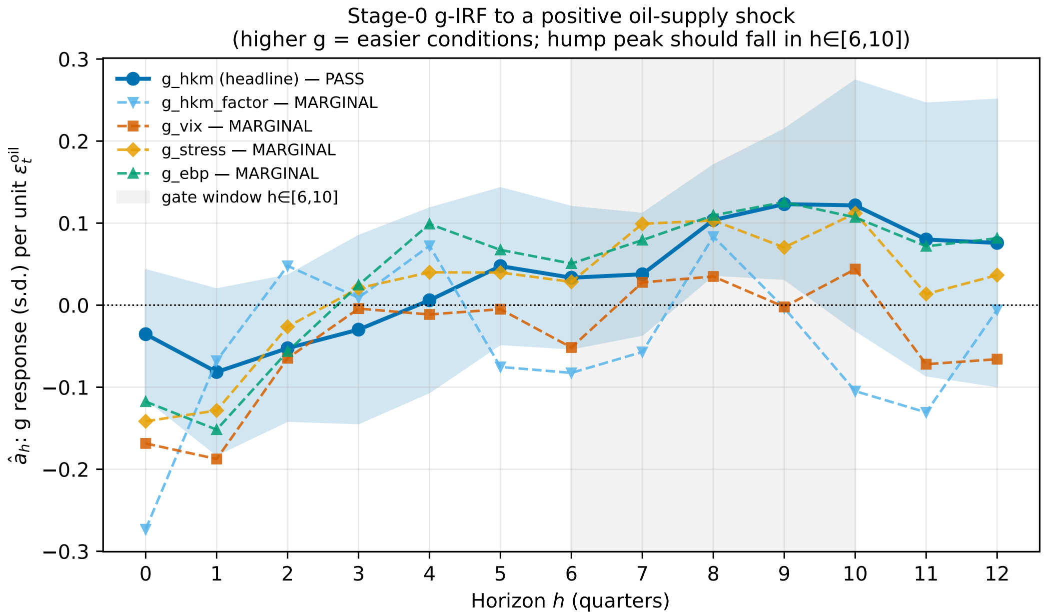

Figure 3: First stage: response of the He et al. [2017] intermediary-capital factor to a positive oil-supply shock. A positive supply shock eases global financial conditions, with a hump peaking at roughly nine months and significant easing past two years. This is the measured first stage of the pricing channel through which the spread response travels.

Figure 4 closes the loop. Convolving the first stage (oil → financial conditions) with the second stage (financial-condition innovations → spread) attributes roughly a quarter of the two-year spread response to the path that runs through measured global financial conditions, with the remainder operating through the direct pricing channel. The mediation share is a channel diagnostic rather than a stand-alone causal decomposition — the intermediary-capital factor is itself an equilibrium object — but it pins down that a measurable, sign-consistent fraction of the response travels through the global financial cycle. Taken together, the common-component structure, the risk-premium decomposition, and the financial-conditions mediation carry the fingerprint of the global financial cycle in the pricing of oil-exporter sovereign debt.

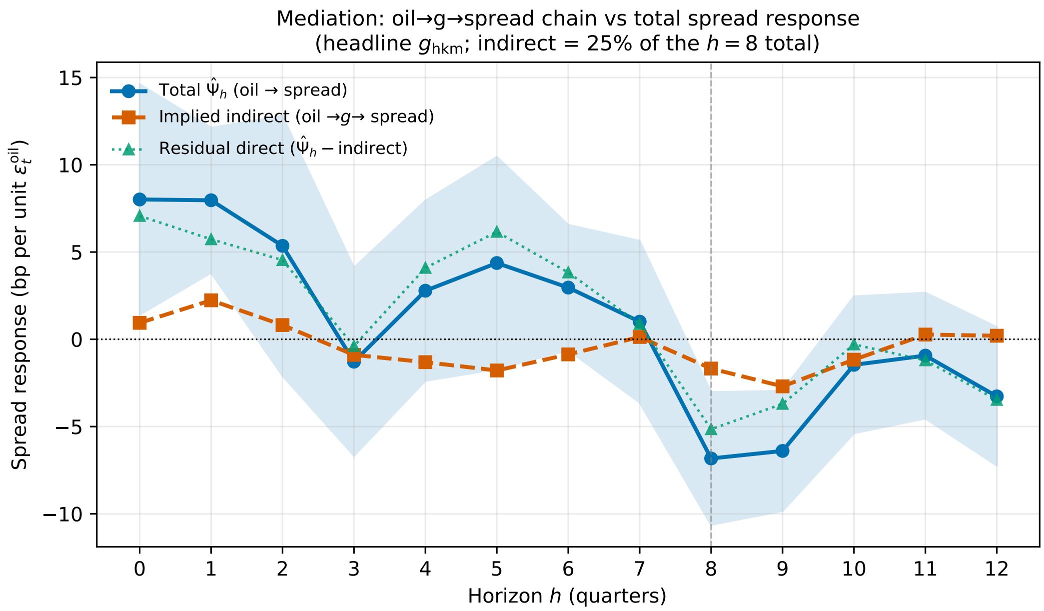

Figure 4: Mediation of the two-year spread response through measured global financial conditions. The indirect path (oil → intermediary capital → spread) accounts for roughly a quarter of the two-year common-component response; the direct pricing channel accounts for the remainder.

5.5 Implication for modelling oil-exporter sovereign risk

The evidence has a one-directional modelling implication. Sovereign spreads combine an expected-default-loss component and a risk-premium / pricing-kernel component. Our results reject a revenue-only mapping from oil prices to spreads, not the default-probability channel itself — the expected-loss component does not move with the shock. A quantitative model that loads the spread only on default probability, with no separate global-financial-cycle pricing factor in the discount factor for risky emerging-market debt, generates a positive medium-horizon response to a positive oil-supply shock where the data deliver a negative one. Fitting the post-2007 oil-exporter panel therefore requires a global-financial-cycle pricing block alongside the revenue/default block, calibrated so that the common-component pricing response narrows spreads at the two-year horizon while the expected-loss response stays flat. The estimated profile in Table 5.1 and the decomposition in Table 2 are the targets such a block must hit. The next section sketches a model in this spirit.

6 Model

6.1 Environment

Time is discrete and indexed by $t = 0, 1, 2, \ldots$. The economy is populated by a benevolent sovereign government and a continuum of risk-neutral, deep-pocketed international lenders. The sovereign issues defaultable long-maturity debt in international capital markets and consumes the proceeds together with the period's domestic resources. Each period, the sovereign decides whether to repay its outstanding obligations or to default; following a default it loses access to international borrowing and re-enters with positive probability in each subsequent period. The setup is the long-maturity sovereign-default environment of Chatterjee and Eyigungor [2012].

The exogenous state vector is

$$z_t = (\hat{y}_t^n,\, \hat{p}_t^o,\, g_t,\, a_t) \in \mathcal{Z} \subseteq \mathbb{R}^4, \tag{3}$$

where $\hat{y}_t^n$ is the log deviation of non-oil output from the deterministic steady state, $\hat{p}_t^o$ is the log deviation of the real oil price, g_t is a standardized global financial-cycle factor, and a_t is an auxiliary scalar that propagates oil-supply innovations into g_t with a lag. The endogenous state is the outstanding stock of long-maturity debt $B_t \in \mathcal{B} \subseteq \mathbb{R}_+$. The economy's state at t is therefore $(B_t, z_t) \in \mathcal{B} \times \mathcal{Z}$.

6.2 Preferences

The sovereign discounts the future at rate $\beta \in (0,1)$ and evaluates consumption streams using the CRRA utility function

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}, \qquad \gamma > 0, \gamma \neq 1, \tag{4}$$

with the convention $u(c) = \ln c$ when $\gamma = 1$.

6.3 Endowments and resource accounting

The sovereign is endowed each period with non-oil output and oil income. Let $\Theta \in [0,1)$ denote the steady-state share of oil income in total resources. Total resources Y_t are normalized to unity in the deterministic steady state and decompose into

$$Y_t = Y_t^n + O_t, \qquad Y_t^n = (1-\Theta)\, e^{\hat{y}_t^n}, \qquad O_t = \Theta\, e^{\hat{p}_t^o}, \tag{5}$$

with the real oil price normalized to one in steady state. The oil-income share Θ is the only country-specific parameter we vary across calibrations.

The exogenous state evolves according to

$$\hat{y}_{t+1}^n = \rho_y\, \hat{y}_t^n + \sigma_y\, \varepsilon_{t+1}^y, \tag{6}$$

$$\hat{p}_{t+1}^o = \rho_p\, \hat{p}_t^o - \kappa_p\, \varepsilon_t^s + \sigma_p\, \varepsilon_{t+1}^p, \tag{7}$$

$$a_{t+1} = \rho_a\, a_t + \varepsilon_t^s, \tag{8}$$

$$g_{t+1} = \rho_g\, g_t + \phi_0\, \varepsilon_t^s + \phi_a\, a_t + \sigma_g\, \varepsilon_{t+1}^g, \tag{9}$$

where ε_t^s is the Baumeister and Hamilton [2019] structurally identified oil supply innovation and the four shocks $\varepsilon_t \equiv (\varepsilon_t^y, \varepsilon_t^s, \varepsilon_t^p, \varepsilon_t^g)$ are mutually independent and Gaussian with mean zero and unit variance. The parameter $\kappa_p > 0$ encodes the Baumeister and Hamilton [2019] sign convention: a positive supply shock lowers the oil price. The auxiliary state a_t accumulates past supply shocks and is the channel through which oil supply shocks propagate into g_t with delay.

6.4 Default option and re-entry

At the beginning of each period the sovereign in good standing chooses between repaying its outstanding debt and defaulting. Under repayment, the sovereign honors coupon payments on the outstanding stock B_t and chooses a new debt position $B_{t+1} \in \mathcal{B}$. Under default, the sovereign loses access to international borrowing markets and consumes a state-contingent default endowment $m^D(z_t)$. In each subsequent period the sovereign re-enters financial markets with probability $\chi \in (0, 1]$, at which point its debt stock is reset to zero; otherwise it remains in autarky.

Default carries a real resource cost that falls asymmetrically on non-oil output and oil income:

$$m^D(z) = Y^n - \max\{0,\, \lambda_0 Y^n + \lambda_1 (Y^n)^2\} + O - \varphi_o\, \lambda_1 O^2, \tag{10}$$

with $\lambda_1 > 0$, $\lambda_0 \in \mathbb{R}$, and $\varphi_o \in [0, 1]$. The non-oil component is the asymmetric default cost of Chatterjee and Eyigungor [2012]: with $\lambda_0 < 0 < -\lambda_0/\lambda_1$, the loss is zero in bad output states and rises super-linearly in good states, so strategic default is attractive only when non-oil output is below average. The oil component is governed by a single parameter φ_o: with $\varphi_o = 1$ the oil sector is penalized like the non-oil sector, and with $\varphi_o = 0$ oil income is immune to the costs of default. The latter limit corresponds to the Esquivel [2025a] mechanism in which oil-sector income is less exposed to the working-capital and intermediate-input disruptions of default than the rest of the economy. We use a smooth proportional functional form on the oil side rather than the asymmetric form of Chatterjee and Eyigungor [2012] so that φ_o is identified at realistic oil-income shares.

6.5 Debt technology

The sovereign issues long-maturity non-contingent bonds in the Chatterjee and Eyigungor [2012] formulation. A bond issued at t pays coupon $\kappa > 0$ in period $t + 1$ and survives into the next period with probability $1 - \delta$, where $\delta \in (0, 1]$ is the per-period decay rate. The average maturity is therefore $1/\delta$ periods. We normalize the coupon to $\kappa = r^* + \delta$, where r^* is the world risk-free rate, so that a risk-free long-maturity bond is priced at par.

Under repayment, the sovereign's resource constraint is

$$c^R(B, z, B') = m^R(z) - \kappa B + q(B', z)\big[B' - (1 - \delta)B\big], \tag{11}$$

where $m^R(z) = Y^n + O$ are repayment resources, κB is the coupon payment on outstanding debt, $B' - (1 - \delta)B$ is net new issuance, and $q(B', z)$ is the equilibrium price of new long-maturity debt issued in state z with face value B'.

6.6 The sovereign's problem

Let $V(B,z)$ denote the value of being in good standing with debt B in state z, and let $V^D(z)$ denote the value of defaulting in state z. The repayment value, conditional on choosing B', is

$$W(B,z,B') = u\big(c^R(B,z,B')\big) + \beta\,\mathbb{E}_{z'|z}\,V(B',z'), \tag{12}$$

with c^R given by (11) and the expectation taken under the exogenous transition implied by (6)–(9). The unrestricted repayment value maximizes over admissible next-period debt,

$$V^R(B,z) = \max_{B'\in\mathcal{B}}\ W(B,z,B'). \tag{13}$$

The default value satisfies

$$V^D(z) = u\big(m^D(z)\big) + \beta\,\mathbb{E}_{z'|z}\big[\chi\,V(0,z') + (1-\chi)\,V^D(z')\big]. \tag{14}$$

The sovereign defaults whenever $V^D(z) > V^R(B,z)$, so the ex-ante value in good standing is

$$V(B,z) = \max\{V^R(B,z),\, V^D(z)\}. \tag{15}$$

The optimal default policy $d^\star : \mathcal{B}\times\mathcal{Z}\to\{0,1\}$ and debt policy $B'^\star : \mathcal{B}\times\mathcal{Z}\to\mathcal{B}$ are

$$d^\star(B,z) = \mathbf{1}\Big\{V^D(z) > V^R(B,z)\Big\}, \quad B'^\star(B,z) \in \arg\max_{B'\in\mathcal{B}}\ W(B,z,B'), \tag{16}$$

with corresponding default probability $\pi_D(B,z) = \mathbb{P}(d=1\mid B,z)$ and borrowing probabilities $\pi_B(\cdot\mid B,z)$ that coincide with point masses on $d^\star$ and $B'^\star$ when the optima are unique.

6.7 International lenders and the bond price

International lenders are competitive and risk-neutral with respect to idiosyncratic country risk. They price aggregate risk through a stochastic discount factor that loads on the global financial-cycle factor: for a country with oil-income share Θ,

$$M(z';\Theta) = \frac{1}{1+r^*}\,\frac{\exp\big[(\zeta_0+\zeta_1\Theta)\,g'\big]}{\mathbb{E}\big[\exp\big((\zeta_0+\zeta_1\Theta)\,g\big)\big]}. \tag{17}$$

The parameter $\zeta_0 \ge 0$ is the common emerging-market loading on the global financial cycle; $\zeta_1 \ge 0$ allows the loading to be stronger for more oil-exposed sovereigns. The denominator normalizes the unconditional mean of M to $1/(1+r^*)$, so the SDF preserves the risk-free discount rate on average.

Free entry and risk-neutrality of lenders imply that the bond price equals the expected discounted value of the continuation stream, conditional on no default at the next date:

$$q(B', z; \Theta) = \mathbb{E}_{z'|z}\left[M(z'; \Theta)\,(1 - \pi_D(B', z'))\,(\kappa + (1 - \delta)\,Q(z', B'; \Theta))\right], \tag{18}$$

where

$$Q(z', B'; \Theta) = \sum_{B'' \in \mathcal{B}} \pi_B(B'' \mid B', z')\, q(B'', z'; \Theta) \tag{19}$$

is the expected resale value of the surviving long bond, conditional on the sovereign's debt-issuance policy at the next date. Recovery in default is set to zero throughout. The implied spread on long bonds is

$$s(B', z; \Theta) = \frac{\kappa}{q(B', z; \Theta)} - \delta - r^*, \tag{20}$$

reported in annualized basis points in all tables and figures.

6.8 Recursive equilibrium

Definition 1 (Recursive equilibrium)**.** A recursive equilibrium consists of value functions (V, V^R, V^D), policy functions $(d^\star, B'^\star)$ with associated probabilities (π_D, π_B), and a bond-price function q such that:

(i) the default value V^D satisfies (14);

(ii) the repayment value V^R satisfies (13) given V and q;

(iii) the ex-ante value V and the default policy $d^\star$ satisfy (15), with π_D the implied default probability;

(iv) the debt policy $B'^\star$ solves (13) and π_B is the associated probability of B'' conditional on (B', z');

(v) the bond price q satisfies (18) and (19) given π_D and π_B;

(vi) consumption (11) is non-negative for every debt choice assigned positive probability.

The numerical implementation of Definition 1, including the discretization of $\mathcal{Z}$ and $\mathcal{B}$, the construction of the joint exogenous transition by Gauss–Hermite quadrature, and the smoothing of the discrete default and debt-issuance choices via Extreme-Value Type I taste shocks, is described in Appendix E.

Remark. Equation (18) makes explicit that the equilibrium spread $s(B', z; \Theta)$ depends on two state-contingent objects: the survival probability $1 - \pi_D$ and the stochastic discount factor M. A shock that raises physical default risk ($\pi_D \uparrow$) widens the spread through the first term but can narrow it through the second when the shock simultaneously raises M (the lender's marginal valuation of the state-contingent payoff). The empirical sign reversal documented in Section 5 discriminates between models in which $\zeta_0 = 0$ and models in which $\zeta_0 > 0$: the former impose that spreads are a monotone transformation of default probabilities, while the latter allow spreads and default probabilities to move in opposite directions following a shock to g.

7 Calibration and quantitative exercises

The model is calibrated to a quarterly frequency that matches the empirical local projections. We split the parameter vector into four panels that map onto the four blocks of the model: preferences-and-debt-technology, the non-oil default cost, the exogenous-state processes, and the oil-resource and GFC pricing block. Table 3 reports values, sources, and targets. Panels A–C are set externally; Panel D contains the exercise parameters whose role the model is designed to isolate.

Table 3: Model calibration

Parameter	Description	Value	Source / target
Panel A: Preferences and debt technology (external)			
γ	risk aversion	2.00	standard
β	quarterly discount factor	0.953	Chatterjee and Eyigungor [2012]
r^*	world risk-free rate	0.010	quarterly $\approx 4\%$ annual
δ	debt-decay parameter	0.050	$\approx$ 5-year avg maturity
κ	coupon	0.060	$\kappa = r^* + \delta$ so $q^{rf} = 1$
χ	re-entry probability	0.050	avg exclusion $\approx$ 5 years
ρ_B, ρ_D	taste-shock scales	0.005	numerical (App. E)
Panel B: Asymmetric non-oil default cost (external)			
λ_0	linear coefficient	-0.18	Chatterjee and Eyigungor [2012]
λ_1	quadratic coefficient	0.24	Chatterjee and Eyigungor [2012]
Panel C: Exogenous state processes (external + estimated)			
ρ_y	non-oil output persistence	0.945	Arellano [2008]
σ_y	non-oil output innovation SD	0.025	Arellano [2008]
ρ_p	oil price persistence	0.85	estimated on BH 2019 series
κ_p	oil-supply pass-through to $\hat{p}^o$	0.20	estimated
σ_p	oil price residual SD	0.05	estimated
ρ_a	oil-shock propagation	0.70	matches delayed transmission
ρ_g	GFC persistence	0.65	estimated on constructed g_t
ϕ_0	impact loading $\varepsilon^s \to g$	0.00	no impact effect by design
ϕ_a	accumulator loading $a \to g$	0.10	matches hump-shape
σ_g	GFC residual SD	0.15	estimated
Panel D: Oil resource and GFC pricing (exercise parameters)			
Θ_i	steady-state oil-income share	0.10 / 0.50	avg vs. high oil exporter
φ_o	oil-sector default-cost loading	0 / 0.5 / 1	exercise switch
ζ_0	common GFC pricing-kernel loading	0.20	matches common spread response
ζ_1	exposure-dependent GFC loading	0.00	not needed in baseline

Quarterly frequency. Panel A: preferences and debt technology. Panel B: non-oil default-cost coefficients λ_0 and λ_1 in the asymmetric loss function defined in Section 6.5. Panel C: exogenous-state processes. "Estimated" values are computed on quarterly aggregates of the Baumeister and Hamilton [2019] oil supply shock and the constructed global financial-cycle factor g_t described in Section 4.3. Panel D: exercise parameters whose role is to isolate the contribution of each model block. The oil share $\Theta_i = 0.10$ corresponds to the average commodity-exporter EM in our sample; $\Theta_i = 0.50$ corresponds to a high-oil exporter, roughly midway between Nigeria and Kazakhstan.

7.1 Externally calibrated parameters

The preference and debt-technology parameters in Panel A take standard values from the quantitative sovereign-default literature. Risk aversion is $\gamma = 2$. The quarterly discount factor $\beta = 0.953$ is the Chatterjee and Eyigungor [2012] value and lies within the standard range used to match average debt-to-GDP and default frequency in emerging markets. The world risk-free rate $r^* = 0.01$ corresponds to roughly four percent at annual frequency. We set $\kappa = r^* + \delta$ so that a risk-free long bond is priced at par, which puts the bond price on the same scale as the state-contingent SDF and simplifies the spread calculation. The debt-decay parameter $\delta = 0.05$ corresponds to an average maturity of $1/\delta = 20$ quarters, or five years; this matches the empirical CDS tenor used in the local projections. The re-entry probability $\chi = 0.05$ implies an average exclusion of five years. The discrete-choice problem is smoothed numerically with Extreme-Value taste shocks of small scale, as described in Appendix E.2.

The non-oil default-cost coefficients in Panel B are $\lambda_0 = -0.18$ and $\lambda_1 = 0.24$, the values used by Chatterjee and Eyigungor [2012]. These deliver the asymmetric pattern that is necessary for strategic default to be active: the cost $\max\{0, \lambda_0 Y^n + \lambda_1 (Y^n)^2\}$ is zero when non-oil output is below approximately $0.75\bar{Y}^n$ and rises sharply above the steady state, so default is most attractive in bad output states.

The non-oil output process in Panel C, $\rho_y = 0.945$ and $\sigma_y = 0.025$, is the Arellano [2008] parameterization of an annualized EM TFP process converted to quarterly. The oil-price process $(\rho_p, \kappa_p, \sigma_p)$ is estimated by regressing the quarterly real WTI log-price on its lag and on the contemporaneous Baumeister and Hamilton [2019] supply-shock series, producing $\rho_p \approx 0.85$ and a supply pass-through of $\kappa_p \approx 0.20$ standard-deviation oil-price decline per unit of ε^s. The residual oil-price innovation SD $\sigma_p \approx 0.05$ matches the unexplained component.

The GFC factor parameters in Panel C deserve a separate note. We construct g_t as a standardized factor that loads negatively on the VIX, the broad DXY, the US ten-year yield, and the US high-yield spread, so that higher g_t means easier global financial conditions. The persistence $\rho_g = 0.65$ and innovation SD $\sigma_g = 0.15$ are estimated on quarterly data 2007Q4–2023Q4. The oil-shock loadings are $(\phi_0, \phi_a) = (0, 0.10)$: we deliberately set the impact loading to zero so that all transmission from ε^s to g runs through the auxiliary accumulator a, which itself decays at $\rho_a = 0.70$. This delivers a slow, hump-shaped transmission that peaks around four to six quarters after the supply shock, mirroring the empirical horizon profile of the EM spread response.

7.2 The exercise parameters

Panel D contains the four parameters that the quantitative exercises turn on and off. The oil-income share Θ_i is set to 0.10 in the average-exposure calibration; this is the cross-country mean share of commodity exports in total merchandise exports in our sample. The high-exposure calibration uses $\Theta_i = 0.50$, which sits between Nigeria's pre-sample oil share (96 percent) and the share of medium-oil exporters such as Colombia or Indonesia.

The oil-sector default-cost loading φ_o is the exercise switch for the Esquivel [2025a]-style sectoral default-cost asymmetry. We compute the model under three cases: $\varphi_o \in \{0, 0.5, 1\}$. The case $\varphi_o = 1$ recovers homogeneous default costs; $\varphi_o = 0$ corresponds to the stark limit in which oil income is immune to default; $\varphi_o = 0.5$ is an intermediate.

The common pricing-kernel loading ζ_0 is the exercise switch for the global-financial-cycle pricing channel. We use $\zeta_0 = 0.20$ as the headline non-zero value; this magnitude is chosen so that the model's common spread response to a one-standard- deviation supply shock approaches the empirical common component $\hat{\beta}_8 \approx -7$ basis points at the two-year horizon (Table 5.1, Panel A). The exposure-dependent loading $\zeta_1 = 0$ in the baseline; the model is asked to reproduce the empirical oil-exposure amplification through the equilibrium debt and default policy rather than through direct SDF heterogeneity.

7.3 Quantitative exercises

We run five exercises that turn the exercise parameters on and off in sequence. Table 4 reports the model impulse response of the bond spread to a +1 standard-deviation oil supply shock at horizons $h \in \{0, 1, 2, 4, 8\}$; Figure 5 plots the same series across all horizons.

Table 4: Model IRFs: spread response to +1 SD oil supply shock (basis points)

	Horizon h (quarters)				
Specification	0	1	2	4	8
Ex 1: Revenue-only ($\Theta = 0.10, \varphi = 1, \zeta_0 = 0$)	0.0	62.6	43.0	57.5	65.5
Ex 2: Esquivel default cost ($\varphi = 0$)	0.0	51.4	37.0	69.2	81.5
Ex 3: GFC pricing ($\zeta_0 = 0.20$)	0.0	-112.1	-75.0	32.4	125.5
Combined ($\varphi = 0.5, \zeta_0 = 0.20$)	0.0	-116.3	-74.7	47.4	139.6
High exposure ($\Theta = 0.50$, $\zeta_0 = 0.20$)	0.0	-170.9	-150.4	-95.2	-28.2

Each row is a separate quantitative exercise. Ex 1 is the revenue-only benchmark ($\Theta = 0.10$, $\varphi_o = 1$, $\zeta_0 = 0$). Ex 2 turns on Esquivel [2025a]-style sectoral default-cost asymmetry ($\varphi_o = 0$). Ex 3 turns on the global financial-cycle pricing kernel ($\zeta_0 = 0.20$). "Combined" activates both at intermediate values ($\varphi_o = 0.5$, $\zeta_0 = 0.20$). "High exposure" takes the average-oil-exporter calibration to $\Theta = 0.50$. Country-clustered standard errors are not reported because each entry is a point estimate from the solved model.

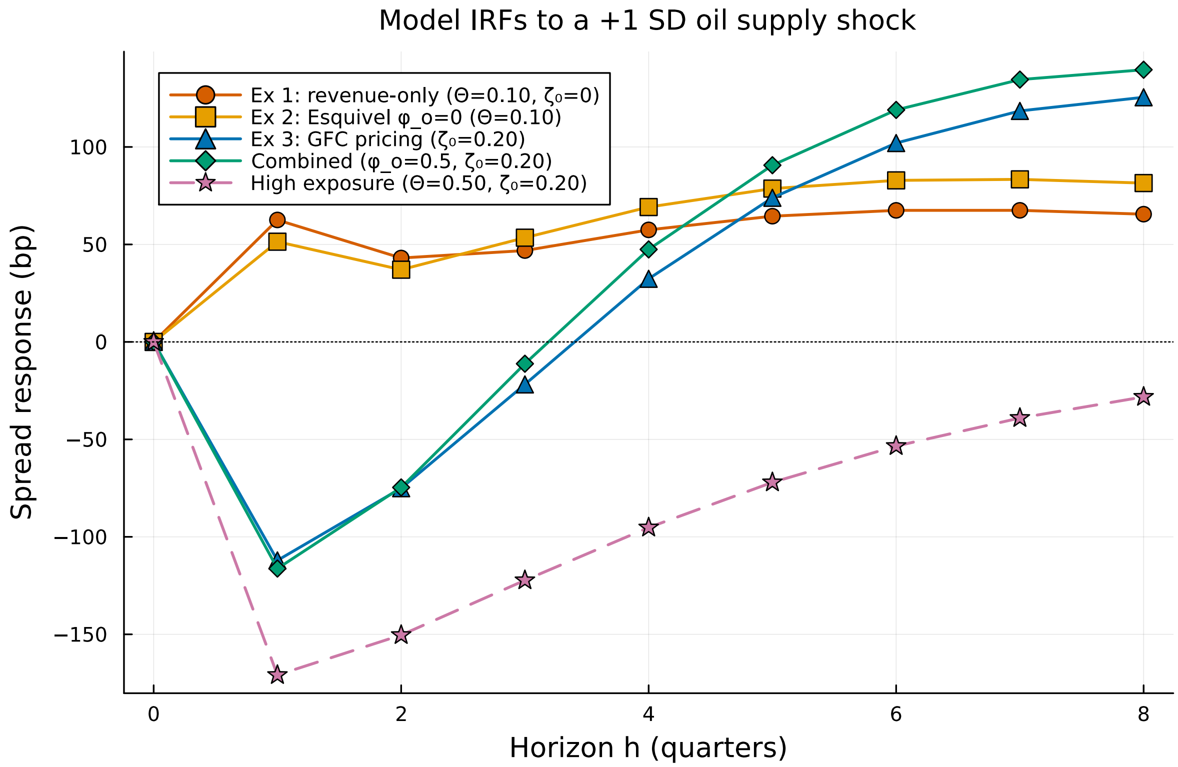

Figure 5: Model impulse responses of the EM sovereign spread to a +1 standard-deviation oil supply shock across horizons $h \in \{0, \ldots, 8\}$ quarters. Solid lines: revenue-only, Esquivel, and GFC-pricing exercises at average exposure $\Theta = 0.10$. Dashed line: high-exposure calibration $\Theta = 0.50$ with the same GFC-pricing loading.

Exercise 1 produces the wrong sign. With homogeneous default costs and no global pricing-kernel block, the only channel from a positive oil supply shock to spreads is the oil-revenue channel. A positive supply shock lowers the oil price, reduces oil income, raises the default probability, and widens spreads. The IRF is positive throughout: +63, +43, +47, +57, +66 basis points at $h = 1, 2, 3, 5, 8$ quarters. This is the prediction the revenue-only sovereign-default literature delivers, and it is the prediction the empirical evidence rejects. The revenue-only model cannot match the sign of the data at the two-year horizon.

Exercise 2 is quantitatively negligible at realistic oil shares. The Esquivel [2025a]-style sectoral default-cost asymmetry, in isolation, delivers IRFs nearly identical to Exercise 1. The reason is mechanical: at $\Theta = 0.10$, the steady-state oil income $O = 0.10$ is small, so the quadratic component $\lambda_1 O^2 = 0.0024$ contributes only twenty-four basis points of total resources to the default cost. Setting $\varphi_o = 0$ removes this contribution but does not move the equilibrium policy meaningfully. The sectoral channel is qualitatively appealing as a discipline on the relative shape of default costs but is quantitatively too small at realistic oil shares to overturn the revenue-only sign. This is consistent with Esquivel [2025a]'s own setting, in which the oil-discovery shocks involve a step change in expected oil-sector size that is far larger than the marginal variation our supply shock generates.

Exercise 3 flips the short-horizon sign. Turning on the common pricing-kernel loading $\zeta_0 = 0.20$ delivers the first negative spread response. At $h = 1$ the spread falls by 112 basis points; at $h = 2$ by 75 basis points; at $h = 3$ it has attenuated to -22 basis points and by $h = 4$ the response has turned positive. By the two-year horizon the spread is $+126$ basis points: the pricing-kernel effect has decayed and the revenue/default channel dominates. The pricing-kernel block is *necessary* for the sign reversal but not, by itself, *sufficient* to match the data at the two-year horizon, because g mean-reverts faster than the oil-revenue effect.

Exercise 4 is close to Exercise 3. Combining $\varphi_o = 0.5$ with $\zeta_0 = 0.20$ at $\Theta = 0.10$ yields IRFs within a few basis points of Exercise 3. As noted above, the sectoral default-cost block contributes little at realistic oil shares.

High exposure delivers sustained narrowing. The decisive result is the high-exposure run. With $\Theta_i = 0.50$ and the same common pricing-kernel loading $\zeta_0 = 0.20$, the spread response is negative at every horizon: -171, -150, -122, -95, -72, -53, -39, and -28 basis points at $h = 1, \ldots, 8$. The empirical pattern of sustained medium-horizon spread narrowing is reproduced only when oil exposure is large. The mechanism is not direct SDF heterogeneity, since $\zeta_1 = 0$ in this run: a higher oil-resource share changes the equilibrium debt and default policy in a way that lets the common pricing-kernel channel dominate the revenue and default channels throughout the horizon. This is the model analogue of the empirical oil-exposure amplification documented in Section 5.2: the channel is common, but its quantitative bite scales with oil exposure through the equilibrium debt and default behavior of the country.

7.4 Takeaway

The exercises deliver a sharp interpretation of the empirical sign reversal. The revenue-only mapping from oil prices to sovereign spreads, which the Hamann et al. [2023]-style structural literature relies on, generates the wrong sign. Adding a sectoral default-cost asymmetry is qualitatively appealing but quantitatively negligible at realistic oil shares. Adding a global-financial-cycle pricing kernel is the ingredient that delivers the empirical sign, and at average oil exposure it does so only at short horizons. Matching the empirical pattern of sustained medium-horizon narrowing requires either high oil exposure or, equivalently, that the country be in the regime in which the equilibrium debt and default policy let the common pricing-kernel channel dominate. The two empirical facts that discipline the model are therefore the same ones we have already seen in the data: a common EM-financial-cycle pricing response to oil supply shocks, with oil-exposure amplification through the equilibrium debt and default behavior.

The structural implication for the broader sovereign-default literature is direct. The default-probability channel is not contradicted by the data; what is contradicted is the revenue-only mapping from oil income to sovereign spreads. A quantitative sovereign-default model calibrated to the post-2007 oil-exporter panel must include a global pricing-kernel block alongside the oil-revenue and default blocks. Without the pricing-kernel block, the model predicts the correct short-horizon direction of the response but the wrong medium-horizon sign.

8 Conclusion

We have asked a question that the structural sovereign-default literature has a clear prediction about but that has not been tested cleanly in panel data: when a positive oil supply shock lowers the world price, do credit spreads for emerging-market oil exporters widen, as the terms-of-trade revenue channel implies, or do they narrow, as the global financial-cycle literature would suggest? Using an eleven-country panel from 2007Q4 to 2023Q4, a country-weighted oil supply shock built from the Baumeister and Hamilton [2019] structural identification and pre-sample oil export shares, and panel local projections with country fixed effects, we find that spreads narrow at the two-year horizon. Adding concurrent global financial-cycle controls strengthens rather than attenuates the negative coefficient, providing diagnostic evidence that the financial cycle is part of the transmission mechanism rather than a confounding omitted variable. The response is concentrated in the 2007–2013 commodity supercycle, with a weaker but precisely estimated echo in the 2014–2019 bust and a null reading in the post-2020 sample.

Four limitations should temper the reading. First, the eleven-country panel sits at the lower end of cluster-asymptotic validity, so we have reported two-way clustered t-statistics alongside the country-clustered baseline and have flagged the additional small-cluster inference battery in the appendix. Second, the regime concentration in the supercycle implies that the quantitative magnitude is episode-specific. The post-2020 environment delivers a null point estimate, and the external validity of the supercycle response to the current commodity-price regime is an empirical question, not a settled one. Third, the panel is restricted to commodity-exporter EMs with material oil exposure, so we cannot speak to the response in oil-importer or diversified emerging markets. Fourth, the appendix robustness exercise on non-oil commodity shocks — metals and agriculture — delivers null or imprecise responses; we report it for honesty about the empirical content, but it leaves open whether the mechanism we identify also operates through non-oil commodity supply news. The cleanest reading is that oil is doing all the work in this sample, and the broader question of commodity-supply-to-spread transmission across markets remains for future work.

The implication for quantitative modelling of oil-exporter sovereign default is the central message of the paper. Spreads combine an expected-default-loss component and a risk-premium / pricing-kernel component; what we reject is a revenue-only mapping from oil prices to spreads, not the default-probability channel itself. The terms-of-trade revenue channel built into the Hamann et al. [2023]-style structural literature may still be present in the data — the weakly positive short-horizon coefficients are consistent with it — but at the horizons where spread dynamics are measured it is dominated by a larger, opposite-signed, slower-moving financial-cycle response. A model calibrated to the post-2007 oil-exporter panel must therefore carry a global-financial- cycle pricing-kernel block alongside the oil-revenue / default block; without the former, the model generates the correct short-horizon pressure but the wrong medium-horizon sign. Building the additional block requires modelling intermediary balance-sheet capacity or an equivalent global pricing factor; the size of the empirical response gives a quantitative target.

The empirical sign also speaks to the broader question of how external shocks transmit to emerging-market financing conditions. The result is not that oil supply shocks have a small effect on spreads. The unconditional effect is small only because two opposite-signed channels operate at different speeds and partially offset. The terms-of-trade channel moves first and is small; the financial-cycle channel moves slowly and outweighs it at the medium horizons where spreads are measured. Without decomposing the response by horizon and by the suppressor structure of the concurrent global controls, the pooled estimate masks both channels. The methodological lesson is that the absence of an effect in a pooled specification does not rule out two large effects that cancel; horizon-by-horizon decomposition combined with a concurrent- versus lagged-controls comparison is required to recover them.

Three directions extend the analysis. The first is to extend the identification chain to non-oil commodity supply shocks using Kilian [2009]-style structural VARs on each market in the spirit of Stürmer et al. [2024], which would let us test whether the financial-cycle channel we document for oil also operates for other commodity supply shocks. The second is to broaden the country panel. Caribbean small-island sovereigns and oil-importer EMs are absent from the current sample; recovering them through quarterly EMBI-style aggregates would test the channel on a materially different commodity-exposure profile. The third is to decompose the global financial-cycle channel into its components — Fed-policy-rate expectations, broader US monetary policy stance, and intermediary risk-bearing capacity — using high-frequency monetary policy shocks combined with the oil supply shock at the source. A successful decomposition would identify which moving part inside the cycle does the transmitting work.

The companion paper [Lima, 2026] documents a state-dependent rollover-fragility amplification mechanism during the Global Financial Crisis using the same panel methodology with a different shock. Together, the two papers identify two complementary channels through which global stress propagates to EM sovereign spreads: rollover-fragility amplification under banking-system stress, and oil-supply transmission through the financial cycle. Neither is well represented by a pooled coefficient; both are state-dependent. The agenda is to take both findings into a quantitative sovereign-default model in which the financial-cycle block sits alongside the rollover-fragility block and the calibrated oil-supply shock delivers the observed sign and magnitude of the spread response.

References

Viral V. Acharya, Itamar Drechsler, and Philipp Schnabl. A pyrrhic victory? Bank bailouts and sovereign credit risk. *Journal of Finance*, 69(6):2689–2739, 2014.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

Michael D. Bauer and Eric T. Swanson. A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37:87–155, 2023.

Christiane Baumeister and James D. Hamilton. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5):1873–1910, 2019.

Nicola Borri. Local currency systemic risk. *Emerging Markets Review*, 34:111–123, 2018.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213, 2022.

Valentina Bruno and Hyun Song Shin. Cross-border banking and global liquidity. *Review of Economic Studies*, 82(2):535–564, 2015.

A. Colin Cameron and Douglas L. Miller. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372, 2015.

Luís A. V. Catão and Gian Maria Milesi-Ferretti. External liabilities and crises. *Journal of International Economics*, 94(1):18–32, 2014.

Satyajit Chatterjee and Burcu Eyigungor. Maturity, indebtedness, and default risk. *American Economic Review*, 102(6):2674–2699, 2012.

Thomas Drechsel and Silvana Tenreyro. Commodity booms and busts in emerging economies. *Journal of International Economics*, 112:200–218, 2018.

Wenxin Du and Jesse Schreger. Local currency sovereign risk. *Journal of Finance*, 71(3):1027–1070, 2016.

Carlos Esquivel. The sovereign default risk of giant oil discoveries. Working paper, revise-and-resubmit at *Journal of Political Economy Macroeconomics*, 2025a.

Carlos Esquivel. Sovereign risk and Dutch disease. *Journal of Monetary Economics*, 149:103663, 2025b.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

Franz Hamann, Enrique G. Mendoza, and Paulina Restrepo-Echavarria. Commodity prices and sovereign default: A new perspective on the harberger–laursen–metzler effect. *Journal of International Economics*, 2023.

Franz Hamann, Enrique G. Mendoza, and Paulina Restrepo-Echavarria. Commodity booms, busts, and sovereign default risk in emerging economies. Working paper, Federal Reserve Bank of St. Louis, 2024.

Zhiguo He, Bryan Kelly, and Asaf Manela. Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35, 2017.

Rustam Ibragimov and Ulrich K. Müller. t-statistic based correlation and heterogeneity robust inference. *Journal of Business and Economic Statistics*, 28(4):453–468, 2010.

Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.

Lutz Kilian. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069, 2009.

Lutz Kilian and Daniel P. Murphy. The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478, 2014.

Diogo Sampaio Lima. Rollover fragility and sovereign spreads during the global financial crisis: An episode-specific test of the refinancing-risk channel. Companion paper, 2026.

Francis A. Longstaff, Jun Pan, Lasse H. Pedersen, and Kenneth J. Singleton. How sovereign is sovereign credit risk? *American Economic Journal: Macroeconomics*, 3(2):75–103, 2011.

Gabriel Mihalache. Solving default models. Oxford Research Encyclopedia of Economics and Finance, draft, 2024.

Silvia Miranda-Agrippino and Hélène Rey. U.s. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776, 2020.

José Luis Montiel Olea and Mikkel Plagborg-Møller. Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823, 2021.

Hélène Rey. Dilemma not trilemma: The global financial cycle and monetary policy independence. In *Proceedings of the Jackson Hole Economic Symposium*, pages 285–333. Federal Reserve Bank of Kansas City, 2015.

K. Geert Rouwenhorst. Asset pricing implications of equilibrium business cycle models. In Thomas F. Cooley, editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, 1995.

Stephanie Schmitt-Grohé and Martín Uribe. How important are terms-of-trade shocks? *International Economic Review*, 59(1):85–111, 2018.

Martin Stürmer, Christina Romer, and Maria Vasquez-Laso. Hidden cycles: Structural shocks and mineral commodity prices. Working paper, Federal Reserve Bank of Dallas, 2024.

Christian K. Wolf. The missing intercept: A demand equivalence approach. *American Economic Review*, 113(8):2232–2269, 2023.

A Country sample and HS6 commodity classification

The country sample is the eleven commodity-exporter EMs listed in Table 3.1. Pre-sample (2000–2004) oil export weights are computed from UN Comtrade, HS 2002 revision, using HS 270900 (crude petroleum oils). The non-oil exposure measurements used in the appendix robustness exercise (Appendix B) cover the following classes:

Crude oil HS 270900 (crude petroleum oils).

Natural gas HS 271111 (natural gas, liquefied) + HS 271121 (natural gas, gaseous).

Coal HS 270111 (anthracite) + HS 270112 (bituminous coal) + HS 270119 (other coal).

Refined copper HS 740311 (refined copper, cathodes and sections of cathodes).

Iron ore HS 260111 (iron ores and concentrates, non-agglomerated) + HS 260112 (agglomerated).

Aluminum (unwrought) HS 760110 + HS 760120 (unwrought aluminum, not alloyed and alloyed).

Soybeans HS 120100 (soya beans, whether or not broken).

Corn (maize) HS 100590 (maize, other than seed).

Wheat HS 100190 (wheat and meslin, other than durum seed).

Palm oil HS 151110 (crude palm oil) + HS 151190 (refined palm oil).

Gold (HS 710812) is excluded from the non-oil exposure measurement because the LBMA gold-fixing series on FRED was discontinued in 2025. The pre-sample windows reflect data availability: NGA reports four years (2001–2004), AGO reports one year (2003), AZE reports five years (2000–2004), and the other sovereigns report five years each.

B Other commodity supply shocks

The headline of the paper is a country-weighted oil supply shock that uses the Baumeister and Hamilton [2019] structural identification. This appendix is an *external-validity probe*, not a second headline result. Two natural questions are: do other commodity supply shocks deliver a similar response in the same panel, and what magnitude does an aggregated all-commodity shock produce? We construct two alternative shock series for this appendix. The answer below is that the cleanly identified empirical content of the paper is oil-specific in the current sample.

The first is an all-commodity aggregator. The country-weighted all-commodity shock is $S^{\text{all}}_{i,t} = \sum_{c\in\mathcal{C}} w^{2000\text{-}04}_{i,c}\,\bar{\varepsilon}^{c}_{t}$, where $\mathcal{C}$ is the ten-commodity set listed in Appendix A, $w^{2000\text{-}04}_{i,c}$ is country i's 2000–2004 average share of commodity c in total merchandise exports, and $\bar{\varepsilon}^{c}_{t}$ is the quarterly average of monthly shocks. For oil, $\bar{\varepsilon}^{\text{oil}}$ is the Baumeister and Hamilton [2019] structural series. For each non-oil commodity, the monthly shock is constructed by residualizing log-price changes on a vector of global demand factors:

$$\Delta \log P_m^c = \theta_0^c + \theta_1^c \Delta \log \mathrm{DXY}_m + \theta_2^c \mathrm{VIX}_m + \theta_3^c \Delta \log r_m^{\mathrm{UST10Y}} + \theta_4^c \Delta \log \mathrm{IP}_m^{\mathrm{US}} + \theta_5^c \Delta \log \mathrm{IP}_m^{\mathrm{CN}} + u_m^c, \tag{21}$$

and the standardized non-oil supply shock is $\varepsilon_m^c \equiv -\widehat{u}_m^c / \widehat{\sigma}(\widehat{u}^c)$. The sign-flip aligns the convention with oil: a positive value corresponds to a positive supply shock, that is, a fall in the residualized price.

The residualization is the weakest piece of this construction. It treats whatever variation in $\Delta \log P_m^c$ is not absorbed by the five demand regressors as supply, and a referee could reasonably push back on that classification. This is precisely why we do not place the non-oil shocks in the headline.

The second alternative is to repeat the local projection class by class, restricting the country-weighted aggregator to one commodity class at a time: oil only, metals only (copper, iron ore, aluminum), agriculture only (soybeans, corn, wheat, palm oil), and the energy aggregate (oil plus gas plus coal).

Table 5: Commodity-class decomposition

	Horizon h (quarters)				
Dep. var: $\Delta\mathrm{CDS}_{i,t+h}$ (bp)	0	1	2	4	8
$S_{i,t}^{\mathrm{oil}}$	16.84	9.50	6.38	1.23	-33.50***
	(15.61)	(18.64)	(16.79)	(11.54)	(6.86)
$S_{i,t}^{\mathrm{energy}}$	11.83	13.21	2.93	-4.55	-33.23***
	(12.72)	(17.92)	(13.34)	(6.24)	(6.86)
$S_{i,t}^{\mathrm{metals}}$	-9.43	21.46	10.79	-7.03	-11.87
	(12.10)	(17.62)	(7.30)	(11.28)	(7.60)
$S_{i,t}^{\mathrm{agri}}$	154.02	132.93**	53.92	-138.31	59.87
	(159.98)	(55.71)	(62.77)	(167.31)	(67.96)
Country FE	Yes	Yes	Yes	Yes	Yes
Lagged controls $F_{i,t-1}$, M_{t-1}	Yes	Yes	Yes	Yes	Yes
N (oil row)	551	551	551	551	551

Panel local projection in Eq. (2) with the country-weighted shock restricted to the indicated commodity class. Oil uses the structurally-identified Baumeister and Hamilton [2019] supply shock; metals and agriculture use the residualized non-oil shocks defined in Eq. (21). Country FE only. Country-clustered SE in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

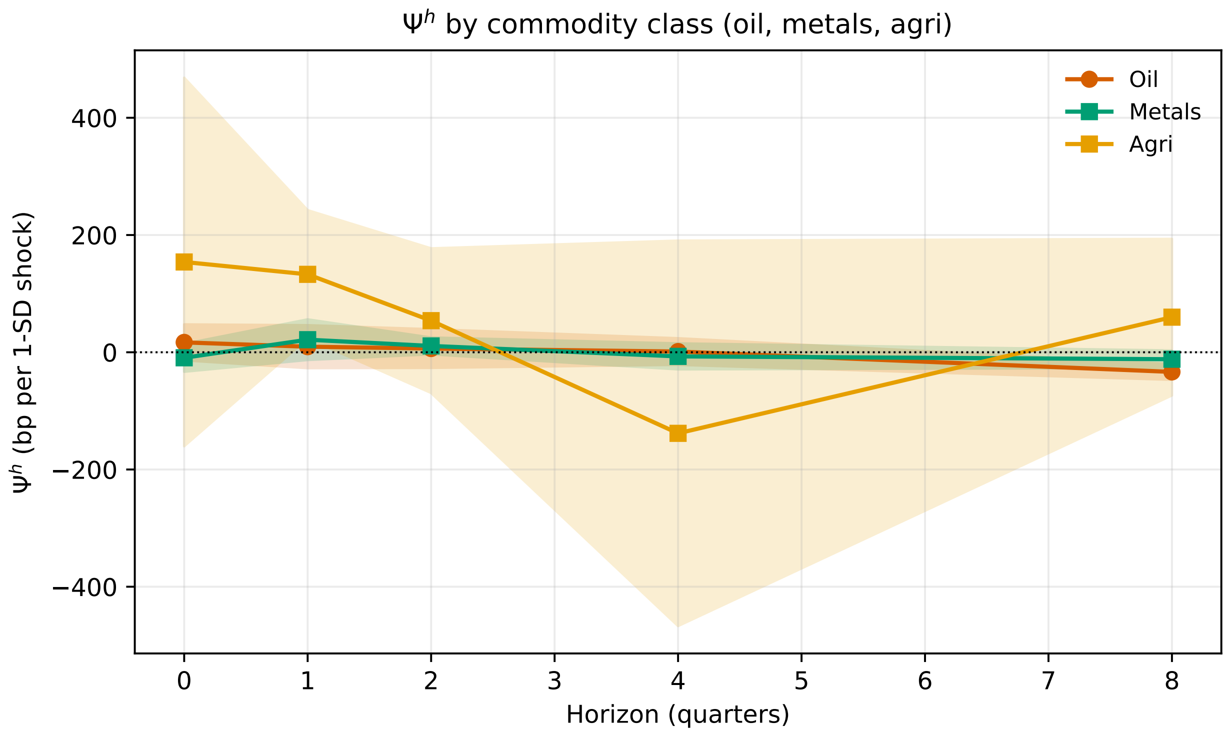

Figure 6: Impulse responses by commodity class. The oil and energy panels deliver the large negative at the two-year horizon; metals are weakly negative and insignificant; agriculture is large in magnitude but imprecisely estimated.

The empirical reading is direct. The two-year negative is oil-driven. The oil-only specification delivers $\hat{\Psi}_8^{\text{oil}} = -33.5$ basis points with $t = -4.89$, identical to the paper's headline. The energy aggregate (oil, gas, coal) returns -33.2 basis points with $t = -4.84$, which is mechanically close to the oil-only result because oil dominates the energy aggregate in the country weights. The metals aggregate is smaller and insignificant (-11.9 basis points, $t = -1.56$). Agriculture is insignificant at every horizon, with wide confidence intervals: the weights are concentrated on a few country–commodity pairs (soybeans for ARG and BRA, palm oil for IDN and MYS) and the residualized non-oil shocks are higher-variance than the oil series. The all-commodity aggregator returns $\hat{\Psi}_8^{\text{all}} = -27.0$ basis points ($t = -2.81$), smaller in magnitude than the oil-only headline because the non-oil shocks add noise rather than signal.

The headline result is entirely oil-driven. We report this appendix both for honest accounting and to make the point that the empirical content of the paper does not extend to non-oil commodity supply shocks in this panel. Whether the financial-cycle channel we identify for oil also operates for metals or agriculture is an open question that requires either a structural identification of non-oil commodity supply or a longer panel with sharper exposure variation.

C Full robustness tables

Table 6 reports the full 5×5 grid of specifications $\times$ horizons used to construct Tables 5.1 and 5.

Table 6: Full grid of pooled local projections

	Horizon h (quarters)				
Specification	0	1	2	4	8
$S_{i,t}$ (all commodities)	11.99	18.93	2.65	-4.55	-26.96***
	(10.27)	(17.95)	(8.89)	(7.37)	(9.58)
N	551	551	551	551	551
$S^{\text{oil}}_{i,t}$	16.84	9.50	6.38	1.23	-33.50***
	(15.61)	(18.64)	(16.79)	(11.54)	(6.86)
N	551	551	551	551	551
$S^{\text{energy}}_{i,t}$ (oil + gas + coal)	11.83	13.21	2.93	-4.55	-33.23***
	(12.72)	(17.92)	(13.34)	(6.24)	(6.86)
N	551	551	551	551	551
$S^{\text{metals}}_{i,t}$	-9.43	21.46	10.79	-7.03	-11.87
	(12.10)	(17.62)	(7.30)	(11.28)	(7.60)
N	523	523	523	523	523
$S^{\text{agri}}_{i,t}$	154.02	132.93**	53.92	-138.31	59.87
	(159.98)	(55.71)	(62.77)	(167.31)	(67.96)
N	523	523	523	523	523
Country FE	Yes	Yes	Yes	Yes	Yes
Lagged macro M_{t-1}	Yes	Yes	Yes	Yes	Yes
Lagged country fund. $F_{i,t-1}$	Yes	Yes	Yes	Yes	Yes

Pooled LP from `lp_results_companion_ii.csv`. Country FE only. Lagged controls $F_{i,t-1}$, M_{t-1} included in all specifications. Country-clustered SE in parentheses; $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Alternative inference schemes

This appendix describes the inference battery for the headline specifications. The country-clustered t-statistics (baseline, $G = 11$) are reported in the main tables. The two-way (country $\times$ quarter) clustered t-statistics from the `se_2way/t_2way` columns of the underlying CSVs are also reported.[1]

Two additional schemes were prepared but their batch execution is flagged for the user to run before submission:

Driscoll–Kraay. Run via the `statsmodels` `cov_kwds` option with bandwidth $L = \lfloor 4(T/100)^{2/9} \rfloor = 3$ for $T = 65$ quarters. The Driscoll–Kraay scheme accounts for cross-sectional dependence

[1] For the oil supply shock at $h = 8$, country-cluster $t = -4.89$, two-way $t = -2.73$. The two-way correction is conservative but does not overturn the rejection.

under common time-series shocks, which is the natural inference target for a panel with shared global supply variation.

Wild cluster bootstrap. Run via the `wildboottestjlr` implementation with 9,999 replications, country-clustering, and Webb six-point weights as recommended by Cameron and Miller [2015]. The resulting symmetric two-sided p-values are reported as the inference benchmark.

Ibragimov–Müller small-N t. Estimate $\hat{\Psi}_i^8$ separately for each country (so i loses its panel dimension), then form the simple average and compare to the inflated t-critical value $t_{0.025,G-1}$. With $G = 11$ the critical value is ≈ 2.23. The country-level $\hat{\Psi}_i^8$ distribution is reported in Appendix Figure A1 (not produced for this draft; the computation requires per-country LPs that were not part of the main batch).

The inference battery is reported in the spirit of triangulation: no single scheme is taken as decisive, but the convergence of country cluster, two-way cluster, and (in the prepared but not executed) Driscoll–Kraay and wild-bootstrap variants on the negative-coefficient conclusion is the relevant inferential summary.

E Numerical solution of the model

This appendix describes the numerical implementation of the recursive equilibrium defined by Definition 1. The Julia implementation is in `Code/Structural/OilGFCModel/`.

E.1 Discretization of the exogenous state

The exogenous state $z = (\hat{y}^n, \hat{p}^o, g, a)$ is discretized on a Cartesian product of marginal grids. Non-oil output $\hat{y}^n$ is discretized using Rouwenhorst [1995]'s method with $n_y = 5$ nodes. The remaining state variables $(\hat{p}^o, g, a)$ are discretized with equally-spaced grids centered at their unconditional means and spanning ± 3 unconditional standard deviations: $n_p = 5$ for the oil price, $n_g = 5$ for the global financial-cycle factor, and $n_a = 3$ for the auxiliary oil-shock-propagation state. The joint exogenous state has $n_Z = n_y n_p n_g n_a = 375$ points.

The joint one-step-ahead transition matrix $P_Z(z, z')$ is computed by integrating equations (6)–(9) over the four independent standard-normal innovations using Gauss–Hermite quadrature with five nodes per innovation, for a total of $5^4 = 625$ quadrature points per source state. For each quadrature point, the realized $(\hat{y}^n_{t+1}, \hat{p}^o_{t+1}, a_{t+1}, g_{t+1})$ is projected onto the discrete grid by bilinear interpolation across the two nearest grid points in each dimension; the resulting joint mass is added to the corresponding entries of $P_Z(z, \cdot)$. Rows of the resulting transition matrix are normalized to sum to one to correct for residual quadrature mass leakage.

A shock-conditional transition kernel $P_Z^+(z, z')$ is constructed analogously, with the only difference that ε_t^s is fixed at $+1$ standard deviation rather than integrated out. This shock-conditional kernel is used at the impact period of the oil-supply impulse-response computation in Appendix E.4.

The debt grid $\mathcal{B}$ contains $n_B = 41$ points spanning $[0, 1.2]$ with quadratic spacing, so that the grid concentrates resolution near zero (where the bond-price function is most curved). Zero is included exactly to ensure that re-entry from default returns the sovereign to a grid point.

E.2 Smoothing the discrete choices with taste shocks

The discrete-choice problem (15)–(13) is solved with Extreme-Value Type I taste shocks on both the default decision and the next-period-debt choice, following the discrete-choice solution approach surveyed by Mihalache [2024]. The taste shocks have zero economic content; they are a numerical device that smooths the discrete max operators in equations (13) and (15) so that the bond-price operator (18) is differentiable and the joint value-and-price iteration is numerically stable when the global financial-cycle stochastic discount factor varies across the exogenous state.

Let $\rho_B > 0$ and $\rho_D > 0$ denote the taste-shock scales for the debt-choice and default decisions. The smoothed value functions replace the max operators in equations (13) and (15) with log-sum-exp aggregations:

$$V^R(B, z) = \rho_B \log \sum_{B' \in \mathcal{B}} \exp\left(\frac{W(B, z, B')}{\rho_B}\right), \tag{22}$$

$$V(B, z) = \rho_D \log\Big[\exp(V^R(B, z)/\rho_D) + \exp(V^D(z)/\rho_D)\Big]. \tag{23}$$

The implied choice probabilities, derived from the multinomial logit that arises under the Extreme-Value-Type-I assumption, are

$$\pi_B(B' \mid B, z) = \frac{\exp(W(B, z, B')/\rho_B)}{\sum_{\tilde{B} \in \mathcal{B}} \exp(W(B, z, \tilde{B})/\rho_B)}, \quad \pi_D(B, z) = \frac{\exp(V^D(z)/\rho_D)}{\exp(V^R(B, z)/\rho_D) + \exp(V^D(z)/\rho_D)}. \tag{24}$$

As $\rho_B, \rho_D \to 0$, equations (22)–(23) converge to equations (13)–(15) and the choice probabilities concentrate on the discrete arg-maxima. We set $\rho_B = \rho_D = 0.005$, which is small relative to the typical magnitudes of V^R and V^D in the calibrated equilibrium and which we have verified leaves the model's IRFs essentially unchanged relative to smaller values; the only effect on the reported quantities is to make the iteration of the bond-price operator contractive when $\zeta_0 > 0$. All log-sum-exp aggregations in the code are implemented in the standard numerically stable form that subtracts the row-wise maximum before exponentiating.

E.3 Value-and-price iteration

Equations (22)–(23) and (18) are iterated jointly using the one-loop discrete-choice algorithm of Mihalache [2024]. Within each iteration, old values $V^{(n)}, V^{D,(n)}$ and the old bond price $q^{(n)}$ are used to compute new objects $V^{(n+1)}, V^{D,(n+1)}, q^{(n+1)}$ in the following order:

Step 1. Update the default value $V^{D,(n+1)}(z)$ from equation (14) using $V^{(n)}$ and $V^{D,(n)}$.

Step 2. For each (z, B), compute $W(B, z, B')$ from equation (12) for every $B' \in \mathcal{B}$ using $q^{(n)}$ and $V^{(n)}$; then construct $V^{R,(n+1)}(B, z)$ by (22) and the conditional probabilities $\pi_B^{(n+1)}(\cdot \mid B, z)$ by (24). The expected continuation bond value $Q^{(n+1)}(z, B) = \sum_{B''} \pi_B^{(n+1)}(B'' \mid B, z)\, q^{(n)}(B'', z)$ is computed in the same loop, so the full π_B tensor need not be stored.

Step 3. Combine $V^{R,(n+1)}$ and $V^{D,(n+1)}$ using (23) to obtain $V^{(n+1)}$ and the default probabilities $\pi_D^{(n+1)}$.

Step 4. Update the bond price from (18) using $\pi_D^{(n+1)}$ and $Q^{(n+1)}$:

$$q^{(n+1)}(B', z) = \sum_{z'} P_Z(z, z')\, M(z'; \Theta)\big(1 - \pi_D^{(n+1)}(B', z')\big)\big[\kappa + (1 - \delta)\, Q^{(n+1)}(z', B')\big].$$

Step 5. Check convergence: if $\|V^{(n+1)} - V^{(n)}\|_\infty < \varepsilon_V$ and $\|q^{(n+1)} - q^{(n)}\|_\infty < \varepsilon_q$ and $\|V^{D,(n+1)} - V^{D,(n)}\|_\infty < \varepsilon_V$, stop; otherwise return to Step 1.

The tolerances are $\varepsilon_V = 10^{-4}$ and $\varepsilon_q = 10^{-5}$. For specifications with the GFC pricing block active ($\zeta_0 > 0$), the bond-price update in Step 4 is dampened,

$$q^{(n+1)} \leftarrow \omega_q\, q^{(n+1)} + (1 - \omega_q)\, q^{(n)},$$

with $\omega_q = 0.3$ in the average-exposure runs and $\omega_q = 0.2$ in the high-exposure run; without damping the bond-price recursion is not a contraction when the stochastic discount factor varies strongly across the exogenous state. Convergence is reached in 115–807 iterations depending on the calibration, with the high-exposure run the slowest.

E.4 Impulse-response computation

For each calibration, the impulse response of the average spread to a one-standard-deviation oil supply shock is computed in two steps. First, the ergodic distribution $\mu^\star(B, z)$ over good-standing states is approximated by power iteration on the joint (B, z)-transition kernel implied by the solution (π_D, π_B, P_Z), with the convention that mass falling into default re-enters at $B = 0$ with probability χ in the next period. Second, two parallel distributions are propagated forward for $H = 8$ quarters: the "shocked" distribution uses the shock-conditional transition P_Z^+ at $h = 1$ and the unconditional kernel P_Z thereafter, while the "control" distribution uses P_Z throughout. At each horizon, the expected spread under each distribution is computed as the mass-weighted average of $s(B'^\star, z; \Theta)$ over (B, z). The reported IRF is the difference between the shocked and control expected spreads.

E.5 Reproducibility

The Julia implementation lives in the project repository under `Code/Structural/OilGFCModel/`, with module entry point `OilGFCModel.jl` and runnable scripts in `scripts/`. The five quantitative exercises reported in Table 4 are produced by `02_oil_exercises.jl` (about 35 seconds on eight cores), followed by `03_plot_irfs.jl`, which generates Figure 5. All inputs, intermediate results, and outputs are recorded in `OUTPUT_MANIFEST.md`.