

Commodity Price Cycles: Commonalities, Heterogeneities, and Drivers

Alain Kabundi* and Hamza Zahid†

September 16, 2026

Abstract. This paper studies common commodity-price cycles and the shocks behind them. We estimate a multilevel dynamic factor model (DFM) for 35 real commodity prices over 1970:M4–2024:M7, separating global, group-specific, and idiosyncratic components of price variation. The global factor accounts for 12.3 percent of average commodity-price variation over the full sample, rising from 8.3 percent in 1970–1999 to 18.6 percent in 2000–2024. We then embed the estimated global commodity factor in a structural factor-augmented vector autoregression (FAVAR) with global economic activity and inflation to identify a global aggregate demand shock, a global aggregate supply shock, and a commodity-market shock. Aggregate-supply shocks generate more persistent commodity-price increases than aggregate-demand or commodity-market shocks, accompanied by higher output and an initial decline in inflation. Historical decompositions attribute the 2000s super-cycle mainly to a sustained sequence of positive aggregate-supply shocks, whereas the post-2020 cycle was less persistent and reflected a changing mix of aggregate-supply, commodity-market, and aggregate-demand shocks. The results show that similar increases in the global commodity factor can come from different shocks and therefore have different implications for output, inflation, and the persistence of commodity booms.

Keywords: Commodity price comovement, commodity super-cycles, global business cycle, dynamic factor model, factor-augmented vector autoregression.

JEL Classification: C51, E53, Q02

*Kabundi: International Monetary Fund; akabundi@imf.org.

†Zahid: World Bank; hzahid@worldbank.org (Corresponding author). We thank M. Ayhan Kose, Kei-Mu Yi, Franziska Ohnsorge, Peter Nagle, John Baffes, Gert Peersman, Bent Sørensen, Valery Charnavoki, Gene Kindberg-Hanlon, Manmohan Kumar, Franz Ulrich Ruch, and seminar participants at the International Monetary Fund, University of Houston, and the World Bank for many useful discussions and comments. The views expressed are those of the authors, and do not necessarily represent the views of the institutions they are affiliated with.

I. Introduction

Commodity prices exhibit recurrent boom–bust cycles with macroeconomic consequences for inflation, trade balances, and fiscal revenues. The most recent episode began in early 2020 and culminated in mid-2022, when a broad set of commodities moved sharply and in tandem. Over this period, the prices of crude oil, wheat, and copper rose by roughly 260%, 150%, and 97%, respectively, before declining across the board as global growth slowed in late 2022. Episodes of synchronized commodity price movements, however, are not unique to the post-pandemic period. In the post–World War II era, at least two other broad-based boom–bust cycles, the 1970s and the 2000s, were characterized by large, widespread increases in commodity price indices followed by reversals.¹

A large literature documents the strong positive correlation between common fluctuations in commodity prices and measures of global economic activity (see, for example, Delle Chiaie et al., 2022). This correlation is informative, but it does not identify the forces behind common commodity-price movements. These common fluctuations can arise from aggregate-demand shocks, such as synchronized macroeconomic stimulus or a broad shift in the propensity to spend; from aggregate-supply shocks, such as productivity growth or trade liberalization, which raise commodity prices by increasing the demand for commodity inputs; or from commodity-market shocks, such as shifts in commodity supply disruptions or storage demand. These forces imply different paths for output and inflation, and therefore different policy tradeoffs. This motivates the central question of the paper: which forces drive the common component of commodity-price fluctuations?

We address this question using a two-step empirical framework. First, we estimate a multilevel dynamic factor model (DFM) for a panel of 35 real commodity prices spanning 1970:M4–2024:M7. The model decomposes each commodity price growth rate into a global factor that captures comovement across all commodities, group factors that capture comovement within commodity groups, and idiosyncratic components. Second, we embed the estimated global commodity factor in a structural factor-augmented vector autoregression (FAVAR) alongside measures of global economic activity and global inflation. The FAVAR identifies three shocks, all normalized to raise commodity prices but differing in their output and inflation responses: a global aggregate demand shock, a global aggregate supply shock, and a commodity-market shock.

The three shocks are identified by sign restrictions on the responses of global activity, global inflation, and the global commodity price factor. An aggregate-demand shock raises activity and inflation; an aggregate-supply shock raises activity and lowers inflation; and a commodity-market

¹See Radetzki (2006) and Baffes & Kabundi (2024) for discussions of major commodity price booms in the post–World War II period.

shock raises inflation and lowers activity. A New Keynesian model with multi-sector commodity inputs provides the rationale for these restrictions, generating each pattern in turn. A preference shock raises spending, output, inflation, and the demand for commodity inputs. A productivity shock outside commodity markets lowers marginal cost and inflation, while the expansion of output raises commodity demand and prices. Inventory-demand shocks and commodity supply disruptions raise commodity prices without expanding productive capacity or final demand; the higher cost of commodity inputs raises marginal cost and inflation, while output falls. The model also implies that aggregate responses do not separate inventory demand from sector-specific supply disruptions, so the identified commodity-market shock is interpreted as a composite of disturbances originating within commodity markets.²

Our first set of results concerns the global commodity factor estimated by the dynamic factor model. The global factor explains, on average, 12.3% of the variation in commodity price growth over the full sample, with larger shares for metals and precious metals than for fertilizers and raw materials. Its average contribution more than doubled across subsamples, from 8.3% in 1970–1999 to 18.6% in 2000–2024. The global factor is associated with measures of global economic activity, but it does not coincide one-for-one with them, and its largest peaks often occur around major commodity-market events.

The structural results show that the three shocks sustain the global commodity factor over different horizons. The response to an aggregate-demand shock peaks within the first year and then unwinds; the response to an aggregate-supply shock builds over about half a year and is the most persistent of the three; and the response to a commodity-market shock rises quickly but remains below it after the first year. This ranking anticipates the historical decompositions, which show that major commodity booms differ in the shocks that drive the global commodity factor.

Under the identifying restrictions, the 2000s commodity boom, widely described as a “super-cycle” in the literature,³ was associated predominantly with a sustained sequence of positive aggregate-supply shocks.³ The post-2020 commodity price surge reflected a changing mix of shocks rather than a single persistent driver. Aggregate-supply and commodity-market shocks supported prices through much of the recovery from mid-2020 onward, aggregate-demand shocks remained a drag until late 2021 and then turned positive as the reopening broadened, and commodity-market shocks intensified around Russia’s invasion of Ukraine. The 1970s boom was shaped primarily by commodity-market shocks.

²The restrictions identify each shock through the sign of its effects on output, inflation, and commodity prices. The model’s preference and productivity shocks are illustrative in this sense: other disturbances with the same effects, including ones the model does not include, are equally admissible.

³Productivity growth, trade liberalization, structural reforms, and rapid industrialization in emerging economies are possible sources because each can increase demand for commodity inputs. The decomposition does not attribute the shocks to particular mechanisms, countries, or policies.

The paper makes three main contributions. First, using a large panel of commodities over more than five decades, we document the global component of commodity-price comovement, its evolution over time, and the group-specific comovement that remains once the global component is removed. Second, we decompose the global commodity cycle into aggregate-demand, aggregate-supply, and commodity-market forces, identifying all three jointly in a single structural framework whose restrictions are motivated by a New Keynesian dynamic stochastic general equilibrium (DSGE) model with multi-sector commodity inputs. Third, we use historical decompositions, reported for the full sample and for major commodity-price episodes, to show how the contribution of each shock changes over time.

This paper relates to a growing body of empirical work on common movements in commodity prices. Since Pindyck & Rotemberg (1990) documented comovement among the prices of largely unrelated commodities, a range of methodologies has been used to investigate the existence and drivers of common price cycles in commodity markets (see, amongst others, Ai et al., 2006; Cuddington & Jerrett, 2008; Lescaroux, 2009; Alquist et al., 2020; Baffes & Kabundi, 2023). A number of studies have applied multilevel dynamic factor models to decompose commodity price variation into global, sectoral, and idiosyncratic components (e.g., Yin & Han, 2015; Delle Chiaie et al., 2022; Bilgin & Ellwanger, 2017), while a related strand has explored the macroeconomic drivers of common commodity price fluctuations using structural vector autoregression frameworks (e.g., Byrne, Sakemoto, & Xu, 2017; Charnavoki & Dolado, 2014; Ha, Kose, & Ohnsorge, 2019). Our contribution is to distinguish aggregate-demand from aggregate-supply forces within the common commodity cycle and to quantify their changing historical contributions alongside commodity-market disturbances. This extends the measurement of global and sectoral comovement in Bilgin & Ellwanger (2017) and Delle Chiaie, Ferrara, & Giannone (2022). It also applies a sign-restriction scheme close to that of Charnavoki & Dolado (2014), who study a small commodity-exporting economy, to a broad commodity factor rather than to the crude-oil price studied by Kilian (2009).⁴

The remainder of the paper is structured as follows. Section II presents the theoretical framework that motivates the identification of structural shocks. Section III describes the data, the dynamic factor model, the FAVAR specification, and the estimation strategy. Section IV reports the empirical results: the estimated global and group factors, the variance decompositions, the impulse responses, and the historical decompositions for major boom–bust episodes. Section V concludes.

⁴Appendix A provides a fuller discussion of the related literature.

II. Theoretical Considerations

A stylized New Keynesian model with multi-sector commodity inputs motivates the empirical identification strategy.⁵ The model shows that aggregate-demand, aggregate-supply, and commodity-market disturbances generate distinct responses of output, inflation, and commodity prices. The preference and productivity shocks below illustrate the first two disturbances, and the inventory-demand and negative commodity-supply shocks the third. The empirical shocks may also reflect other disturbances that produce the same responses. Preference and productivity shocks are standard disturbances in New Keynesian models, while storage and supply disruptions are the disturbances most commonly studied in the commodity-market literature.⁶

Production. A continuum of monopolistically competitive firms produce differentiated final goods using labor and a composite commodity input. Firm j 's production technology is

$$Y_t(j) = A_t \left[\alpha^{1/\mu} M_t(j)^{(\mu-1)/\mu} + (1-\alpha)^{1/\mu} N_t(j)^{(\mu-1)/\mu} \right]^{\mu/(\mu-1)}, \quad (1)$$

where A_t is aggregate total factor productivity (TFP), $M_t(j)$ is the commodity input, $N_t(j)$ is labor, α is the commodity cost share, and μ is the elasticity of substitution between labor and commodities. TFP follows a first-order autoregressive process, denoted AR(1), $\hat{a}_t = \rho_A \hat{a}_{t-1} + \varepsilon_t^A$, where $\hat{x}_t \equiv \log X_t - \log \bar{X}$ denotes the log deviation of a variable from its steady-state value. A positive productivity innovation ($\varepsilon_t^A > 0$) raises the marginal product of all inputs, including commodities. Firms set prices subject to Calvo (1983) staggered price setting, generating a New Keynesian Phillips curve that links inflation to marginal cost.

Households. A representative household maximizes expected lifetime utility over consumption and leisure:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\xi_t \frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right], \quad (2)$$

where β is the discount factor, σ is the coefficient of relative risk aversion, φ is the inverse Frisch elasticity of labor supply, $\chi > 0$ scales the disutility of labor, and ξ_t is a preference shifter that follows $\hat{\xi}_t = \rho_{\xi} \hat{\xi}_{t-1} + \varepsilon_t^{\xi}$. A positive preference shock ($\varepsilon_t^{\xi} > 0$) increases the marginal utility of current consumption, raising aggregate demand and, in equilibrium, output and inflation. Monetary policy follows a Taylor-type rule that responds to inflation and the output gap.

⁵The model combines Calvo (1983) price rigidity with the commodity-input and shock-transmission mechanisms studied by Bodenstein & Guerrieri (2011) and Plante (2014).

⁶The model is intentionally stylized and is not intended as a quantitative account of global commodity markets. The online appendix reports the full model derivation and calibration.

Commodity market. Commodity supply is modeled as an upward-sloping schedule,

$$Q_t = \bar{Q} \exp(\zeta \hat{p}_t^{m,r} + \hat{z}_t), \quad (3)$$

where ζ is the supply elasticity and $\hat{z}_t$ captures exogenous commodity supply disturbances. A negative realization ($\hat{z}_t < 0$) reduces commodity supply and raises commodity prices. In practice, commodity supply disruptions are typically sector-specific, for example, oil production outages due to conflict or supply management by the Organization of the Petroleum Exporting Countries (OPEC), weather-related disruptions to agricultural output, or mining strikes. Because energy, metals, and food enter a constant-elasticity-of-substitution (CES) commodity composite, a sector-specific shock can move the broad commodity price index, especially when short-run substitution across groups is limited.

The model also includes a stylized inventory channel. Inventory holders choose commodity stocks because inventories provide storage services, including precautionary availability and the option to meet future demand. Desired inventories for commodity group g satisfy

$$\hat{s}_{g,t}^* = v_t + v_{g,t} - \eta_S \hat{p}_{g,t}^{m,r}, \quad (4)$$

where v_t is a global inventory-demand state, $v_{g,t}$ is a group-specific inventory-demand state, and $\eta_S > 0$ is the price elasticity of desired stocks. Actual inventories adjust gradually,

$$\hat{s}_{g,t} = \rho_S \hat{s}_{g,t-1} + (1 - \rho_S) \hat{s}_{g,t}^*. \quad (5)$$

Commodity market clearing is physical: production use plus inventory accumulation must equal supply,

$$Q_{g,t} = M_{g,t} + I_{g,t}, \quad I_{g,t} = S_{g,t} - (1 - \delta_S) S_{g,t-1}. \quad (6)$$

A positive inventory-demand innovation raises desired stocks, increases current stockbuilding, and therefore raises commodity prices without directly expanding final demand or productive capacity.

Commodity price equilibrium. The log-linearized equilibrium real commodity price index satisfies

$$\hat{p}_t^{m,r} = \frac{(1 - \lambda_S) \hat{d}_t + \lambda_S \hat{i}_t - \hat{z}_t}{\zeta + (1 - \lambda_S) \mu}, \quad \hat{d}_t \equiv \hat{y}_t + (\mu - 1) \hat{a}_t + \mu \widehat{mc}_t, \quad (7)$$

where $\hat{p}_t^{m,r}$ is the real commodity price index, $\hat{d}_t$ is the commodity-demand shifter implied by the firms' cost-minimizing input demand, $\hat{y}_t$ is output, $\hat{a}_t$ is total factor productivity, $\widehat{mc}_t$ is real marginal cost, $\hat{i}_t$ is inventory accumulation, $\hat{z}_t$ is commodity supply, and λ_S is the steady-state share of inventory accumulation and replacement demand in total physical commodity demand. The denominator

captures the slope of excess demand. Higher substitution between labor and commodities or higher commodity supply elasticity dampens price responses. The direct productivity coefficient in the demand shifter is positive when $\mu > 1$ and zero at the Cobb–Douglas benchmark ($\mu = 1$). At that benchmark, productivity affects commodity prices through the equilibrium responses of output and marginal cost. The coefficient is negative when $\mu < 1$, where complementarity between labor and commodities can offset the equilibrium expansion in commodity demand; the Monte Carlo exercise below samples values on both sides of the benchmark. Since output, marginal cost, and inventory accumulation are endogenous, each structural shock affects commodity prices through several channels. The shocks are distinguished by their joint effects on output, inflation, and commodity prices in general equilibrium.

Impulse responses and identifying restrictions. Figure 1 reports impulse responses from the calibrated DSGE model.⁷ A preference shock raises output, inflation, and commodity prices together. A productivity shock lowers marginal cost, and the equilibrium expansion of output raises the derived demand for commodity inputs, so output and commodity prices rise while inflation falls. Inventory-demand and negative commodity-supply shocks raise commodity prices and inflation while lowering output. The empirical commodity-market shock combines these disturbances, which share the sign pattern in columns 3 and 4 of Figure 1.⁸

The shaded regions in Figure 1 assess sensitivity to calibration choices. Of 500 Monte Carlo calibrations, 43 percent preserve the aggregate sign pattern over the first four quarters, and the bands are conditional on that subset.⁹ They show that the model can generate the identifying signs under these calibrations, not that the signs are independently validated by the data. The quarterly model motivates the direction of the responses; the monthly FAVAR imposes the signs only from impact through month 3 and examines longer restriction horizons separately.

III. Data and Methodology

We use a two-step empirical framework. First, a multilevel dynamic factor model decomposes commodity-price growth into global, group-specific, and idiosyncratic components. Second, a

⁷The model is solved using Dynare. The online appendix reports the full calibration, parameter ranges, and implementation details.

⁸A sector-specific disruption that moves prices across several commodity groups contributes to the common component that the global factor captures, while comovement confined to one group is captured by that group’s factor. The correspondence between the model’s price indices and the estimated factors is interpretive.

⁹A calibration is admissible only if every shock produces the required signs for output, inflation, and commodity prices in each of the first four quarters. The ranges include values of μ below one, where labor and commodities are complements; the reported bands use the calibrations that satisfy the joint sign restrictions.

structural FAVAR uses the estimated global factor, world industrial production, and global inflation to identify the shocks driving the global commodity cycle.

III.1 Data

The dataset includes monthly series of commodity prices, global economic activity, and global inflation. The commodity price panel used in the DFM spans 1970:M4–2024:M7. The common sample for world industrial production, the global commodity factor, and global inflation is 1971:M1–2023:M11, which defines the FAVAR estimation period. The commodity price data is obtained from the World Bank’s Commodity Price database (henceforth, the *Pink Sheet*). The Pink Sheet contains data on 55 commodity prices and indices. We exclude series that are either aggregates or close substitutes of other series to avoid introducing price comovement by construction.¹⁰ This leaves us with 35 commodities, which are divided into 9 distinct groups based on the Pink Sheet’s classification. The groups include energy, metals, precious metals, fertilizers, grains, food oils, other foods, beverages, and raw materials.¹¹ All series are deflated using the U.S. Consumer Price Index for all urban consumers (CPI) published monthly by the Bureau of Labor Statistics (BLS) and transformed by taking first differences of logs. We measure global economic activity based on the world industrial production index developed by Baumeister & Hamilton (2019). The global inflation rate is proxied by the inflation rate for the Organisation for Economic Co-operation and Development (OECD).

III.2 Dynamic Factor Model

Consider a panel of N commodity price series, each of length T , which can be categorized into R distinct commodity groups. Our goal is to model common movements across all series and common movements within specific commodity groups. Suppose $y_{i,t}$ represents the growth rate of the price of commodity i at time period t . We assume that $y_{i,t}$ can be decomposed into three components: (1) a common global component (or global factor, $f_{g,t}$) that affects all commodity price series, (2) a common group component (or group factor, $f_{r,t}$) that affects only commodity prices within group r , and (3) an idiosyncratic disturbance, $\varepsilon_{i,t}$, that only affects an individual commodity price series i .

¹⁰For example, we include ‘Logs, Malaysian’ but exclude ‘Sawnwood, Malaysian’ since the latter is produced by cutting the logs longitudinally. The two products are vertically linked along the same production chain, so changes in the price of sawnwood closely follow the fluctuations in the price of logs. Thus, including both series will generate price comovement by construction. Similarly, we keep one monthly price series for each of the following: crude oil, natural gas, coffee, tea, soybeans, sugar, and phosphate.

¹¹Table 1 reports all 35 commodities together with their group classification and variance decompositions.

The growth rate of commodity price series i at time period t can, thus, be written as

$$y_{i,t} = b_{0,i} + b_{g,i}f_{g,t} + b_{r,i}f_{r,t} + \varepsilon_{i,t}, \quad (8)$$

where $i = 1, \dots, N$, $r = 1, \dots, R$ with $R < N$ and $t = 1, \dots, T$. The coefficients $b_{g,i}$ and $b_{r,i}$ are factor loadings that reflect the degree to which fluctuations in $y_{i,t}$ can be accounted for by the global and group-specific factors, respectively. The idiosyncratic disturbance, $\varepsilon_{i,t}$, follows an $AR(q)$ process

$$\varepsilon_{i,t} = \psi_i(L)\varepsilon_{i,t-1} + \zeta_{i,t} \quad (9)$$

where $\psi_i(L)$ is a lag polynomial operator, $\zeta_{i,t} \sim N(0, \sigma_i^2)$ and $E(\zeta_{it}\zeta'_{jt}) = 0$ for all $i \neq j$.

The global and group-specific factors can be collected in the vector $F_t = [f_{g,t}, f_{1,t}, \dots, f_{R,t}]$. F_t follows an $AR(p)$ process such that

$$\mathbf{F}_t = \Phi(L)\mathbf{F}_{t-1} + v_t, \quad (10)$$

where $\Phi(L)$ is a diagonal lag polynomial matrix, $v_t \sim N(0, \sigma_r^2 \mathbf{I}_{R+1})$ with σ_r^2 normalized to 1 and $E(\zeta_{it}v_{rt}) = 0$. Since $\Phi(L)$ is a diagonal matrix, the latent factors evolve independently within the state equation. Also, since innovations $\zeta_{i,t}$ and $v_{r,t}$ are contemporaneously uncorrelated, all comovement in commodity prices is captured by the common factor structure rather than by the idiosyncratic disturbances.

In the baseline specification, $p = 3$ and $q = 2$: each latent factor follows an autoregressive process of order 3, $AR(3)$, and each idiosyncratic component follows an autoregressive process of order 2, $AR(2)$.

Identification of Global and Group-specific factors

The factor loadings can be collected in an $N \times K$ matrix, $\mathbf{B}$. We impose zero-restrictions on $\mathbf{B}$ to identify global and group-specific factors. For instance, in a 2-commodity case ($\{i, j\} = \{Crude\ Oil, Copper\}$), where each commodity belongs to a distinct commodity sector, zero restrictions on the loading matrix can be expressed as follows:

$$\begin{bmatrix} b_{g,i} & b_{Energy,i} & 0 \\ b_{g,j} & 0 & b_{Metals,j} \end{bmatrix},$$

where $b_{g,i}$ and $b_{g,j}$ are the loadings for the growth rates in commodity prices for commodities i and j on the global factor, respectively. On the other hand, $b_{Energy,i}$ and $b_{Metals,j}$ are the loadings for price

growth rates of commodity i and commodity j on group-specific factors, respectively. Thus, the growth rate of each commodity price is driven by a global factor that affects all commodity prices, group factors that affect commodity prices within their respective groups, and an idiosyncratic disturbance.

This decomposition separates broad comovement across commodity prices from residual within-group comovement. The zero restrictions on $\mathbf{B}$ do not imply that cross-group propagation is absent in the data. When a sectoral shock, such as an oil market disturbance, moves prices across several commodity groups, the shared component is captured by the global factor. The group factors then capture the within-group comovement that remains after this common cross-commodity component is removed.¹²

III.3 Factor-Augmented Vector Autoregression Model

A factor-augmented vector autoregression provides a structural decomposition of the global commodity factor estimated by the DFM described by equations (8)–(10). The endogenous vector contains world industrial production, q_t , global inflation, π_t , and the estimated global commodity price factor.

We employ the following general form of a structural FAVAR model

$$\mathbf{B}_0 \mathbf{y}_t = \sum_{i=1}^p \mathbf{B}_i \mathbf{y}_{t-i} + \mathbf{w}_t, \quad (11)$$

where $\mathbf{y}_t$ is a $K \times 1$ vector that contains the endogenous model variables, $\mathbf{w}_t$ is the $K \times 1$ vector of mutually uncorrelated structural shocks, $\mathbf{B}_0$ is the contemporaneous structural coefficient matrix that provides a description of the contemporaneous relationships among the model variables, $\mathbf{B}_i$ is the matrix of coefficients and p is the lag length, set to 12 in the baseline. The reduced form errors can be written as $\mathbf{u}_t = \mathbf{B}_0^{-1} \mathbf{w}_t = \mathbf{C}_0 \mathbf{w}_t$, where $\mathbf{C}_0 \equiv \mathbf{B}_0^{-1}$ is the structural impact matrix and

$$\mathbf{u}_t = \mathbf{y}_t - \sum_{i=1}^p \mathbf{A}_i \mathbf{y}_{t-i}, \quad (12)$$

¹²Appendix D4 assesses this interpretation by re-estimating the model over the same sample with the energy factor allowed to load on metals and grains prices. The estimated global factor is virtually unchanged: its monthly posterior median has a correlation of 0.997 with the baseline factor (0.999 for 12-month moving sums), and the commodity-level global variance shares have a correlation of 0.993 across the two specifications. The zero cross-group restrictions therefore do not materially determine the estimated common component.

$\mathbf{A}_i = \mathbf{B}_0^{-1} \mathbf{B}_i$ and $E(\mathbf{u}_t \mathbf{u}_t') = \Sigma_u = \mathbf{C}_0 \mathbf{C}_0'$ is the $K \times K$ reduced form variance-covariance matrix. Thus, $\mathbf{C}_0$ maps the mutually uncorrelated structural shocks, $\mathbf{w}_t$, into the reduced-form innovations. Since this impact matrix is not uniquely identified from Σ_u , we require further identifying restrictions to estimate $\mathbf{C}_0$.

Identification of shocks

We impose sign restrictions on the structural impact matrix $\mathbf{C}_0$ to identify a global aggregate demand shock, a global aggregate supply shock, and a commodity-market shock. The baseline sign restrictions can be described as follows

$$\mathbf{u}_t = \begin{pmatrix} u_t^q \\ u_t^\pi \\ u_t^f \end{pmatrix} = \begin{bmatrix} + & + & - \\ + & - & + \\ + & + & + \end{bmatrix} \begin{pmatrix} w_t^{\text{Global Aggregate Demand}} \\ w_t^{\text{Global Aggregate Supply}} \\ w_t^{\text{Commodity-Market}} \end{pmatrix} = \mathbf{C}_0 \mathbf{w}_t = \mathbf{B}_0^{-1} \mathbf{w}_t,$$

where signs are imposed on the elements of the structural impact matrix, $\mathbf{C}_0 = \mathbf{B}_0^{-1}$, and all shocks are normalized to increase the commodity price factor.

The matrix above summarizes the on-impact sign pattern. In the baseline empirical specification, these restrictions are imposed over horizons $h = 0, \dots, 3$: for industrial production and the commodity price factor, the cumulative responses must preserve the required sign from impact through month 3, while for inflation the point-wise response must preserve the required sign over the same window. We adopt $H = 3$ to constrain the initial response paths and assess sensitivity to shorter and longer horizons separately.¹³

These restrictions are motivated by the model in Section II and are consistent with related sign-restriction studies.¹⁴ A positive aggregate-demand shock raises activity, inflation, and commodity prices together. A positive aggregate-supply shock raises activity and commodity prices but lowers inflation. A positive commodity-market shock raises commodity prices and inflation while lowering activity.

The commodity-market shock captures disturbances that originate within commodity markets and are orthogonal to aggregate-demand and aggregate-supply forces. Examples include precautionary or speculative storage demand, including shifts associated with geopolitical risk, and sector-specific

¹³ Appendix C2.1 reports robustness to the impact-only case ($H = 0$) and to tighter horizons ($H = 6, 12$).

¹⁴ Charnavoki & Dolado (2014) use a closely related scheme to identify aggregate-demand, aggregate-supply, and commodity-market disturbances in a small commodity-exporting economy, while Ha, Kose, & Ohnsorge (2019) study analogous forces behind global inflation.

supply disruptions that spill over to the broad commodity price index.¹⁵ As Section II shows, aggregate sign restrictions cannot separately distinguish demand-side commodity-market disturbances from commodity-sector supply disruptions when both imply higher commodity prices, higher inflation, and lower output. The identified commodity-market shock should therefore be read as a reduced-form composite of commodity-market-specific forces rather than as a pure commodity demand or pure supply shock.

III.4 Estimation

Dynamic Factor Model. Equations (8)–(10) form a state-space system, with (8) as the observation equation and (9)–(10) as the transition equations. We estimate this system in a Bayesian framework following Kim & Nelson (1998). Inference proceeds by Gibbs sampling, alternating between draws of the latent factors and draws of the model parameters. The factors are drawn using the forward-filtering, backward-sampling algorithm of Carter and Kohn (1994), while parameters without closed-form conditional posteriors are sampled with a Metropolis–Hastings step (see Chib & Greenberg, 1994; Otrok & Whiteman, 1998).¹⁶

FAVAR with sign restrictions. To construct the structural impulse responses, we follow Rubio-Ramírez, Waggoner, & Zha (2010). For each candidate rotation, we draw a $K \times K$ matrix $\mathbf{X}$ with independent $N(0, 1)$ entries and obtain an orthogonal matrix $\mathbf{Q}$ from a sign-normalized QR decomposition of $\mathbf{X}$, with the diagonal elements of $\mathbf{R}$ normalized to be positive. Let $\mathbf{P}$ be the lower-triangular Cholesky factor of the estimated reduced-form residual covariance matrix, so that $\mathbf{P}\mathbf{P}' = \hat{\Sigma}_u$. The candidate structural impact matrix is $\tilde{\mathbf{C}}_0 = \mathbf{P}\mathbf{Q}$, which satisfies $\tilde{\mathbf{C}}_0\tilde{\mathbf{C}}_0' = \hat{\Sigma}_u$. We construct impulse responses from each candidate and retain those satisfying the maintained sign-and-path restrictions.

We incorporate parameter uncertainty using the diffuse prior and Gaussian–inverse-Wishart posterior of Kilian & Murphy (2012), without coefficient shrinkage. We draw the VAR coefficients and residual covariance matrix 600 times and evaluate 300,000 rotations per draw. Figure 4 reports posterior medians and 68- and 90-percent credible bands under the baseline $H = 3$ restrictions. The bands incorporate parameter and rotation uncertainty, conditional on the estimated DFM factor and its specification.¹⁷

¹⁵Storage or speculative demand shifts can raise commodity prices without an immediate increase in current production demand; see Fattouh et al. (2013).

¹⁶Appendix C1 reports the priors and simulation length.

¹⁷Appendix C2 gives the posterior distributions and treatment of unstable draws. Posterior summaries also depend on the rotation prior, since a uniform measure over orthogonal matrices does not imply a flat prior over impulse responses

For the restriction-horizon comparisons and the representative historical decompositions, we follow the point-estimate approach of Kilian & Murphy (2014): we hold reduced-form parameters at their ordinary least squares (OLS) estimates and evaluate 5,000,000 candidate rotations. Evaluating this many rotations for every posterior draw is not feasible. The decompositions and recovered shock series use the Fry & Pagan (2011) median-target rotation, the admissible model whose standardized responses are closest to the pointwise median. This choice preserves the reduced-form covariance matrix and yields mutually uncorrelated shocks.¹⁸

IV. Results

IV.1 Common Cycles in Commodity Prices

The global factor summarizes the variation common to all 35 commodity prices in the panel. Because the factor is unobserved and is recovered only from its statistical relationship with observed prices, what economic variation it represents has to be established rather than assumed.

Figure 2 plots the median of the posterior distribution of the global commodity factor, $f_{g,t}$, estimated from the DFM (8) - (10) over the period 1970–2024, together with 68-percent (16th–84th percentile) and 90-percent (5th–95th percentile) Bayesian credible bands. Panel A reports the factor in 12-month moving sums (i.e., year-over-year growth rates of the underlying commodity prices), while Panel B plots the cumulated factor alongside linearly detrended world industrial production, both standardized to zero mean and unit standard deviation for visual comparability. Panel B provides a level-like representation of the global commodity price cycle and its co-movement with global economic activity.¹⁹

Panel A first shows that the 68- and 90-percent credible bands remain narrow throughout the sample period, indicating limited posterior uncertainty under the maintained DFM specification. Some of the largest peaks in common commodity-price growth coincide with large oil-price spikes. Oil is an input into the production and transportation of many commodities, so oil-market disturbances can generate simultaneous price movements across commodity markets (Baffes, 2007; Bilgin & Ellwanger, 2017; Alquist et al., 2020). Mechanically, the positive loading on crude oil prices also

(Baumeister & Hamilton, 2015); Appendix C2.2 discusses this limitation and the distinction between sign and magnitude restrictions.

¹⁸Appendix C3 describes the selection. Appendices D1 and D5 report the near-zero impact bound and the uncertainty of the historical decompositions across rotations and posterior parameter draws.

¹⁹Since the DFM is estimated on month-on-month growth rates, peaks in Panel A correspond to periods of highest year-over-year commodity-price growth and naturally precede peaks in the price level shown in Panel B.

means that an oil-price increase raises the estimated global factor. The crude-oil loading, however, is 0.30, which is not unusually large relative to other commodities.²⁰

Panel B relates the cumulated commodity factor to global economic activity. Prior work shows that common components of commodity prices are closely related to global economic activity (Alquist et al., 2020; Delle Chiaie et al., 2022). Stronger aggregate activity raises demand for commodities used as production inputs, putting upward pressure on prices (Kilian, 2009; Alquist et al., 2020). To examine this relationship visually, Panel B plots the cumulated commodity factor (solid line) alongside linearly detrended world industrial production (dashed line).²¹ The two series broadly co-move, capturing major business cycle episodes: declines during the 1973–1975 and early-1980s recessions, the expansion of 2003–2007, the sharp fall during the global financial crisis and subsequent recovery, the collapse in early 2020 during the COVID-19 recession, and the rapid rebound in 2021–2022 as global demand recovered.²²

However, Panel B also reveals divergences. The commodity factor rises well above world industrial production during the 1970s oil shocks, outpaces it during the 2000s super-cycle, and surges again in 2021–2022 while world industrial production remains comparatively subdued. Conversely, commodity prices decline more sharply than world industrial production during the mid-1980s and post-2011 commodity price downturns. The global factor increases during the early 1990s recession, while its decline during the mid-1970s recession and the 2007–09 global financial crisis is observed with a lag.²³ More precisely, the correlations between the global factor and two measures of global economic activity, the Kilian index and world industrial production, are 0.52 and 0.66, respectively.²⁴ Thus, changes in global economic activity are a key source of shocks to commodity markets. However, the global commodity factor is an equilibrium outcome that reflects a mixture of structural forces, and the strong aggregate correlation with economic activity may mask differences in the nature and persistence of the underlying shocks, with distinct implications for inflation, commodity price persistence, and policy responses.

The variance decomposition in Table 1 quantifies how much of commodity-price variation is associated with the global and group factors. Over the full sample, 1970–2024, the global factor

²⁰For example, the factor loadings for copper, gold, soybeans, and rubber are 0.45, 0.33, 0.36, and 0.48, respectively. The average factor loading is 0.24 and other than tobacco (for which we estimate a small negative loading), all factor loadings are positive.

²¹Blue-shaded bars in Figure 2 depict U.S. recessions as a proxy for slowdowns in global growth.

²²See Kose, Sugawara, & Terrones (2020) for an identification and analysis of turning points of the global business cycle.

²³For example, while global growth slowed down during 2007–08, Frankel (2008) observed that “commodity prices have found their second wind over precisely this period! Up some 25% or more since August 2007, by a number of indices. So much for the growth explanation.”

²⁴See Kilian (2009) for details regarding the Kilian index; the world industrial production index is from Baumeister & Hamilton (2019). The correlations are reported for year-over-year growth rates.

accounts for 12.3 percent of commodity-price variation on average. Its contribution differs sharply across groups. It accounts for 23.9 percent of metals price variation and 23.4 percent of precious-metals price variation, compared with 3.5 percent for fertilizers and 4.0 percent for other foods. This cross-group pattern is consistent with the stronger links between metals demand and global activity; copper, for example, is often viewed as a barometer of global economic conditions (Hamilton, 2015; Bernanke, 2016).

The same table shows that commodity-price comovement has increased since 2000. The global factor accounts for 18.6 percent of commodity-price variation on average over 2000–2024, compared with 8.3 percent over 1970–1999. The increase is broad-based: the number of commodities for which the global factor dominates the group factor rises, while the average idiosyncratic share falls.

The group factors capture residual within-group comovement after the global component is removed.²⁵ They matter mainly for selected commodity groups. Over the full sample, the broad-based within-group criterion is satisfied by precious metals and food oils, for which the group factors account for 45.4 and 48.1 percent of average price variation, respectively.²⁶ The food-oils result pertains to a two-commodity group consisting of coconut and palm oil. During 1970–1999, the criterion is satisfied by energy, metals, precious metals, and food oils, with average group-factor shares of 38.5, 25.7, 52.7, and 54.7 percent, respectively. During 2000–2024, only precious metals and food oils satisfy the criterion, with corresponding shares of 40.8 and 41.0 percent. Thus, precious metals and food oils exhibit group-wide comovement in both subsamples, whereas energy and metals cease to satisfy the criterion after 2000. Fertilizers, beverages, and many agricultural commodities remain dominated by idiosyncratic disturbances.

These results show that common cross-commodity variation has become more pronounced since 2000, but they do not identify the structural forces behind the global commodity factor. The next subsection uses the FAVAR to decompose the global commodity cycle into aggregate-demand, aggregate-supply, and commodity-market shocks.

²⁵ Appendix Figure 6 reports representative group-factor paths.

²⁶ We define broad-based within-group synchronization as an estimated group factor accounting for at least 10 percent of the variation in the price of each commodity within the group. While the average group-factor share is non-trivial for metals (22.3 percent) and grains (30.7 percent), these groups do not meet the broad-based criterion because tin (8.6 percent) and rice (0.3 percent), respectively, fall below the 10 percent threshold.

IV.2 Structural Drivers of the Global Commodity Cycle

The structural results are reported in three parts: the shock series and their relation to externally identified shocks, the impulse responses of output, inflation, and the commodity factor, and the historical decomposition of the global commodity cycle.

Structural Shocks. Before turning to the impulse responses, we compare the shocks recovered from the estimated model with the structural shocks from the oil-market VAR of Baumeister & Hamilton (2019).²⁷ These comparisons provide supporting evidence rather than uniquely establishing the source of each shock. The mapping is necessarily imperfect because their shocks are identified in a crude-oil VAR, whereas our shocks are recovered from a FAVAR for the global commodity factor. The closest mapping is between the Baumeister & Hamilton (2019) “economic activity” shock and our aggregate-demand and aggregate-supply shocks, and between their oil-market demand shocks and our commodity-market shock. The correlations are consistent with this mapping: computed on the monthly shock series, our aggregate-demand and aggregate-supply shocks correlate positively with the Baumeister & Hamilton (2019) economic activity shock ($r = 0.602$ and $r = 0.319$, respectively), while our commodity-market shock is positively correlated with the combined Baumeister & Hamilton (2019) oil consumption and inventory demand shocks ($r = 0.293$). The commodity-market correlation is lower, as expected, because our shock is identified from a broad cross-section of commodity prices and captures disturbances beyond crude oil, including storage-demand and sector-specific supply disturbances.

Figure 3 plots 12-month moving averages of the paired shock series; correlations computed on these smoothed series are 0.493, 0.555, and 0.282, respectively, as smoothing emphasizes the lower-frequency comovement visible in the figure. The two macroeconomic shock comparisons line up more closely around major global downturns and recoveries, while the commodity-market comparison is more modest.²⁸

Impulse Responses. Figure 4 reports posterior-median responses under the baseline $H = 3$ restrictions, with 68- and 90-percent Bayesian credible bands. The bands incorporate parameter and rotation uncertainty, conditional on the estimated DFM factor, its specification, and the identifying

²⁷The structural shocks are recovered as $\hat{e}_t = \hat{C}_0^{-1}\hat{u}_t$, where $\hat{C}_0$ is the impact matrix of the Fry & Pagan (2011) median-target model and $\hat{u}_t$ are the reduced-form VAR residuals. We plot the cumulative series $\sum_{s=1}^t \hat{e}_{i,s}$ for each shock i in Appendix Figure 7. These summarize accumulated shock realizations; historical decompositions combine them with impulse responses to measure contributions to the factor.

²⁸Appendix D2.2 provides additional details, including subperiod correlations with the Baumeister & Hamilton (2019) shocks and a supplementary diagnostic comparing the aggregate-supply shock with China’s TFP growth during the 2000s super-cycle.

restrictions.²⁹ Each shock's responses are divided by its median commodity-factor impact, so the median impact equals one index point rather than one standard deviation.³⁰

The restrictions apply from impact through month 3, at horizons 0:3; the subsequent paths describe the model's estimated propagation of each shock.³¹ Commodity-price increases caused by aggregate-demand shocks generate a hump-shaped expansion in output and a persistent rise in inflation. The commodity-factor response peaks within the first year and then gradually unwinds. Aggregate-supply shocks have a different dynamic profile. Output rises gradually and remains elevated, inflation falls on impact, reaches its trough after two months, and is close to zero after about one year, and the commodity-factor response becomes the largest after roughly half a year and remains the most persistent across the three shocks. This persistence is consistent with the historical decomposition below, where aggregate-supply shocks account for much of the sustained increase in the global commodity factor during the 2000s super-cycle.

Commodity-market shocks generate the clearest contrast. Even though all shocks are normalized to raise commodity prices on impact, the commodity-market shock lowers output and raises inflation, consistent with a cost-push disturbance. The output decline remains negative throughout the horizon shown, while the inflation response builds gradually and then levels off. The commodity-factor response rises quickly and remains positive, but it remains below the response to aggregate-supply shocks after the first year.³²

Historical Decomposition. While the impulse responses in Figure 4 characterize the average propagation of each structural shock, commodity cycles are shaped by sequences of shocks of varying signs and magnitudes. To understand how each shock has contributed to the evolution of the global commodity price factor over time, we turn to historical decompositions.³³ We focus on three major boom-bust episodes: the 1972–1975 commodity boom, the 2003–2008 super-cycle and global financial crisis, and the post-2020 cycle encompassing COVID-19, the recovery, and Russia's invasion of Ukraine. These, together with the 1950–51 boom, are among the largest episodes in post–World War II commodity markets (Radetzki, 2006; Baffes & Haniotis, 2010). Figure 5 presents counterfactual historical decompositions for each episode, constructed following Kilian & Lewis

²⁹Appendix D1 compares this unrestricted Bayesian baseline with the explicit near-zero negative impact restriction and reports the corresponding credible bands for the restricted specification.

³⁰The normalization makes the median profiles comparable across shocks. The same divisor is applied to every draw, so the bands still reflect variation in the impact across draws.

³¹Baseline identification uses multi-period sign restrictions over horizons 0:3 months; see Section III.3 and Appendix C2.1.

³²For a discussion of the increasing role of aggregate-demand shocks relative to aggregate-supply shocks and oil-price shocks in driving global inflation variability, see Ha, Kose, & Ohnsorge (2019).

³³As Kilian & Lütkepohl (2017) emphasizes, neither impulse response functions nor forecast error variance decompositions are sufficient to assess the role of individual structural shocks in explaining specific historical episodes; historical decompositions address this limitation directly. Full-sample counterfactual paths are reported in Appendix Figure 9.

(2011). Since the figure reports the cumulated estimated global factor rather than a commodity price index, magnitudes are stated in points on the cumulated-factor scale.

The 1972–1975 commodity boom. The first episode covers the commodity price boom of the early 1970s and its aftermath. Figure 5 (top row) shows that, by early 1974, the cumulated commodity factor had reached approximately 30 points on this scale. During the initial phase of the boom (1972–1973), aggregate-supply shocks were the primary driver, with positive aggregate-demand shocks also supporting commodity prices. From late 1973 onward, however, commodity market shocks became the dominant force. Their cumulative contribution rose steadily, reaching about 14 points by early 1975, the largest contribution among the three shocks. This finding is consistent with the role of the OPEC oil embargo and commodity-specific supply disruptions, as well as evidence of speculative or inflation-hedging demand for commodities during this period (see Cooper & Lawrence, 1975; C. A. Carter et al., 2011). The subsequent decline in 1974–1975 reflected a joint reversal of aggregate-supply and aggregate-demand contributions, consistent with the recession and associated slowdown in global activity. Commodity-market shocks remained elevated, preventing a sharper collapse. By the end of 1975, commodity market shocks still accounted for most of the remaining increase, whereas aggregate-supply and aggregate-demand contributions had largely unwound.

The 2003–2008 super-cycle and global financial crisis. The second episode covers the commodity “super-cycle” of the 2000s and the subsequent collapse during the global financial crisis (GFC). Between 2003 and mid-2008, the cumulated commodity factor rose from approximately 3 to 42 points (Figure 5, middle row). Under the maintained identification, the boom as a whole was led by aggregate-supply forces: by the June 2008 peak, the cumulative contribution of aggregate-supply shocks had risen to about 25 points, accounting for most of the increase in the global commodity factor. Rapid productivity growth, trade integration, structural reforms, and improvements in production efficiency in China and other emerging markets are plausible contributors (Zhu, 2012; Buyuksahin et al., 2016), but the decomposition does not distinguish among them. In the theoretical model, a positive TFP innovation outside commodity sectors provides one mechanism: it raises output and the derived demand for commodity inputs, pushing commodity prices upward. Related studies model this mechanism as a shock to the representative firm’s total factor productivity (see, for example, Bodenstein & Guerrieri, 2011; Plante, 2014).³⁴

³⁴More broadly, Buyuksahin et al. (2016) argue that the commodity price supercycle began in the mid- to late 1990s, coinciding with a series of reforms in China, including its eventual accession to the World Trade Organization in 2001. Similarly, India’s economic growth during this period has been linked to productivity gains from trade liberalization (see Topalova & Khandelwal, 2011; Kotwal et al., 2011). See also Baffes & Hanjotis (2010), World Bank (2015), and Alquist et al. (2020) on the role of emerging-market growth in the post-2000 commodity boom.

At the same time, the distinction between levels and changes matters here. In cumulative-contribution levels, aggregate-demand shocks remained negative throughout the boom, implying that aggregate-demand conditions alone would not have sustained the super-cycle. Commodity-market shocks were not the main driver of the boom, but they contributed to the late-stage increase before the peak. Their cumulative contribution rose steadily from around zero in early 2006, and they account for much of the final acceleration into the June 2008 peak, consistent with evidence of speculative activity and financialization in commodity markets during this period (see Cheng & Xiong, 2014; Caballero, Farhi, & Gourinchas, 2008).³⁵

Between the June 2008 peak and the February 2009 trough, the commodity factor fell by roughly 20 points. The aggregate-demand and aggregate-supply contributions each accounted for about half of the decline, while the commodity-market contribution changed little. At the trough, the supply contribution was still positive and the demand contribution deeply negative. Thus, supply shocks continued to support the level of commodity prices, even as their reversal drove much of the fall. The collapse reflects a joint reversal of aggregate-demand and aggregate-supply forces rather than a commodity-market event.

The post-2020 cycle: COVID-19, recovery, and Russia's invasion of Ukraine. The third episode spans mid-2019 through late 2023, covering the COVID-19 recession, the subsequent recovery, and the commodity price spike triggered by the Russian invasion of Ukraine. On the eve of the pandemic (February 2020), the cumulated commodity factor stood at approximately 28 points (Figure 5, bottom row). The initial COVID shock (February–April 2020) reduced the factor by roughly 4 points, with aggregate demand accounting for the larger share of the decline and aggregate supply for the remainder. Commodity-market shocks, by contrast, were modestly positive during this period, likely reflecting OPEC+ production cuts that supported commodity prices.

The recovery from mid-2020 through early 2022 was the largest movement in this episode, with the commodity factor surging to over 45 points by March 2022, its highest level within this episode (though below the full-sample peak reached in 2011). This rebound was not driven by a single shock. In cumulative-contribution levels, aggregate-demand shocks remained a drag on the commodity factor until late 2021, reflecting the depth of the initial pandemic collapse. Over the same period, aggregate-supply shocks provided some support, consistent with a partial rebound in global supply conditions, while commodity-market shocks also supported prices, partly reflecting OPEC+ production cuts. From late 2021 into early 2022, aggregate demand turned positive as the reopening broadened, while commodity-market shocks intensified sharply around the Ukraine invasion. From the COVID trough in April 2020 to the March 2022 peak, however, the largest improvement came

³⁵Replacing world industrial production with the Kilian index also yields an aggregate-supply-led boom, but the commodity-market contribution during 2006–2008 turns negative. Appendix D3 discusses this sensitivity to the activity measure.

from aggregate demand, which accounted for roughly four-fifths of the rebound. Commodity market shocks added a small further contribution, while aggregate-supply shocks made little net contribution because their earlier positive effect had largely unwound by early 2022.

The subsequent decline from mid-2022 onward was driven by two forces. First, commodity market shocks reversed sharply, with their cumulative contribution falling from approximately 6 to -2 points as the Ukraine price premium unwound and commodity-specific disruptions eased. Second, aggregate-demand shocks weakened as the global economy decelerated. Aggregate-supply shocks no longer supported prices and turned negative through much of 2022 before easing later, consistent with the fading of the earlier recovery impulse and weaker global industrial momentum.³⁶

These episode conclusions hold across the admissible models rather than resting on the single median-target rotation. Aggregate-supply shocks lead the 2003–2008 increase in most accepted models, aggregate-demand and aggregate-supply shocks jointly account for the global financial crisis collapse, and aggregate-demand shocks lead the post-COVID rebound.³⁷

V. Conclusion

The commodity booms of the 2000s and early 2020s both involved broad increases across many prices, but they were driven by different shocks. Common movements in commodity prices are therefore not, by themselves, informative about the forces behind them. To draw the distinction, we estimate a multilevel dynamic factor model for commodity prices and then embed the estimated global commodity factor in a structural FAVAR with global economic activity and inflation.

The factor estimates show that the global commodity cycle has become more important since 2000. The global factor accounts for 12.3 percent of commodity-price variation over the full sample, rising from 8.3 percent in 1970–1999 to 18.6 percent in 2000–2024. This increase is broad-based and especially pronounced for metals, and the global share for energy roughly tripled, consistent with the close connection of industrial commodities to global activity. Group factors capture additional within-sector comovement in selected markets, but the main increase in cross-commodity synchronization is associated with the global factor.

³⁶This reading is consistent with contemporaneous commodity-market commentary linking weaker metal demand after the 2022 peak to concerns about China’s slowdown, especially in property and heavy industry (e.g., World Bank, 2022, 2024). The decomposition does not identify country-specific shocks, so this evidence is suggestive rather than direct.

³⁷Appendix D5 assesses historical-decomposition uncertainty over admissible rotations and posterior reduced-form parameters, conditional on the estimated DFM factor and maintained specifications.

Under the maintained identification, the 2000s super-cycle is attributed mainly to a sustained sequence of positive aggregate-supply shocks. These represent broad expansionary, disinflationary forces; productivity gains, trade integration, and structural reforms are possible sources. The post-2020 cycle was less persistent and reflected a changing mix of shocks: aggregate supply and commodity-market forces supported prices during much of the recovery, aggregate demand turned positive only late in 2021, and commodity-market shocks intensified around Russia's invasion of Ukraine before unwinding. Commodity-market shocks help explain sharp movements around peaks and troughs, but do not by themselves account for sustained global commodity cycles.

Because the two booms were driven by different shocks, they carried different inflation implications. The 2000s super-cycle was led by expansionary, disinflationary forces whose contribution was the most persistent in our sample. The post-2020 episode involved a changing mix in which commodity-market shocks played a larger role, particularly around the invasion of Ukraine, raising inflation while weakening activity and forcing a choice between stabilizing prices and supporting activity.

AI-use disclosure. During revisions between March and September 2026, the authors used OpenAI Codex and Anthropic Claude for linguistic editing, improved readability, and coding assistance. Specific model versions were not systematically recorded. The authors reviewed and edited the content and take full responsibility for the publication.

References

- Ai, C., Chatrath, A., & Song, F. (2006). On the comovement of commodity prices. *American Journal of Agricultural Economics*, 88(3), 574-588. Retrieved from <https://EconPapers.repec.org/RePEc:oup:ajagec:v:88:y:2006:i:3:p:574-588>
- Alquist, R., Bhattarai, S., & Coibion, O. (2020). Commodity-price comovement and global economic activity. *Journal of Monetary Economics*, 112, 41 - 56. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0304393219300431> doi: <https://doi.org/10.1016/j.jmoneco.2019.02.004>
- Baffes, J. (2007). Oil spills on other commodities. *Resources Policy*, 32(3), 126 - 134. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0301420707000542> doi: <https://doi.org/10.1016/j.resourpol.2007.08.004>
- Baffes, J., & Haniotis, T. (2010, July). *Placing the 2006/08 commodity price boom into perspective* (Policy Research Working Paper Series No. 5371). The World Bank. Retrieved from <https://ideas.repec.org/p/wbk/wbrwps/5371.html>
- Baffes, J., & Kabundi, A. (2023). Commodity price shocks: Order within chaos? *Resources Policy*, 83, 103640. doi: 10.1016/j.resourpol.2023.103640
- Baffes, J., & Kabundi, A. (2024). Do supercycles dominate commodity price movements? *Economics Letters*, 237, 111646.
- Baumeister, C., & Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83(5), 1963–1999. doi: 10.3982/ECTA12356
- Baumeister, C., & Hamilton, J. D. (2019, May). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review*, 109(5), 1873-1910. doi: 10.1257/aer.20151569
- Bernanke, B. S. (2016, February 19). *The relationship between stocks and oil prices*. Retrieved from <https://www.brookings.edu/articles/the-relationship-between-stocks-and-oil-prices/>
- Bernanke, B. S., & Boivin, J. (2003). Monetary policy in a data-rich environment. *Journal of Monetary Economics*, 50(3), 525–546. doi: 10.1016/S0304-3932(03)00024-2
- Bilgin, D., & Ellwanger, R. (2017). A dynamic factor model for commodity prices. *Bank of Canada Staff Analytical Notes*.

- Blake, A. P., & Mumtaz, H. (2012). *Applied Bayesian econometrics for central bankers* (No. 4). Centre for Central Banking Studies, Bank of England. Retrieved from <https://ideas.repec.org/b/ccb/tbooks/4.html>
- Bodenstein, M., & Guerrieri, L. (2011). Oil efficiency, demand, and prices: A tale of ups and downs. *International Finance Discussion Paper, 2011*, 1-49. doi: 10.17016/IFDP.2011.1031
- Buyuksahin, B., Mo, K., & Zmitrowicz, K. (2016). Commodity Price Supercycles: What Are They and What Lies Ahead? *Bank of Canada Review, 2016*(Autumn), 35-46. Retrieved from <https://ideas.repec.org/a/bca/bcarev/v2016y2016iautumn16p35-46.html>
- Byrne, J., Fazio, G., & Fiess, N. (2013). Primary commodity prices: Co-movements, common factors and fundamentals. *Journal of Development Economics, 101*, 16-26.
- Byrne, J., Sakemoto, R., & Xu, B. (2017). *Commodity price co-movement: Heterogeneity and the time varying impact of fundamentals* (MPRA Paper). University Library of Munich, Germany. Retrieved from <https://EconPapers.repec.org/RePEc:pra:mprapa:80791>
- Caballero, R. J., Farhi, E., & Gourinchas, P.-O. (2008, December). *Financial crash, commodity prices and global imbalances* (Working Paper No. 14521). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w14521> doi: 10.3386/w14521
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics, 12*(3), 383–398.
- Carter, C. A., Rausser, G. C., & Smith, A. (2011). Commodity booms and busts. *Annual Review of Resource Economics, 3*(1), 87-118. doi: 10.1146/annurev.resource.012809.104220
- Carter, C. K., & Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika, 81*(3), 541–553. Retrieved from <http://www.jstor.org/stable/2337125>
- Cashin, P., Mohaddes, K., Raissi, M., & Raissi, M. (2014). The differential effects of oil demand and supply shocks on the global economy. *Energy Economics, 44*, 113 - 134. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0140988314000619> doi: <https://doi.org/10.1016/j.eneco.2014.03.014>
- Charnavoki, V., & Dolado, J. J. (2014, April). The effects of global shocks on small commodity-exporting economies: Lessons from Canada. *American Economic Journal: Macroeconomics, 6*(2), 207-37. doi: 10.1257/mac.6.2.207
- Cheng, I.-H., & Xiong, W. (2014). Financialization of Commodity Markets. *Annual Review of Financial Economics, 6*, 419–441. doi: 10.1146/annurev-financial-110613-034432
- Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review: Papers & Proceedings, 101*(3), 471–475.
- Chib, S., & Greenberg, E. (1994). Bayes inference in regression models with ARMA(p, q) errors. *Journal of Econometrics, 64*(1-2), 183-206. Retrieved from <https://EconPapers.repec.org/RePEc:eee:econom:v:64:y:1994:i:1-2:p:183-206>

- Cooper, R. N., & Lawrence, R. Z. (1975). The 1972-75 commodity boom. *Brookings Papers on Economic Activity*, 1975(3), 671–723. Retrieved from <http://www.jstor.org/stable/2534151>
- Cuddington, J. T., & Jerrett, D. (2008, Dec 01). Super cycles in real metals prices? *IMF Staff Papers*, 55(4), 541–565. doi: 10.1057/imfsp.2008.19
- Delle Chiaie, S., Ferrara, L., & Giannone, D. (2022). Common factors of commodity prices. *Journal of Applied Econometrics*, 37(3), 461–476.
- Fattouh, B., Kilian, L., & Mahadeva, L. (2013). The role of speculation in oil markets: What have we learned so far? *The Energy Journal*, 34(3), 7–33.
- Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the Penn World Table. *American Economic Review*, 105(10), 3150–3182. doi: 10.1257/aer.20130954
- Frankel, J. (2008, March 25). *An explanation for soaring commodity prices*. Retrieved from <https://cepr.org/voxeu/columns/explanation-soaring-commodity-prices>
- Fry, R., & Pagan, A. (2011, December). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4), 938–60. doi: 10.1257/jel.49.4.938
- Ha, J., Kose, M. A., & Ohnsorge, F. (Eds.). (2019). *Inflation in emerging and developing economies: Evolution, drivers, and policies*. Washington, DC: World Bank. doi: 10.1596/978-1-4648-1375-7
- Ha, J., Kose, M. A., Otrok, C., & Prasad, E. (2017). Global macro-financial cycles and spillovers. *Paper presented at the 18th Jacques Polak Annual Research Conference, International Monetary Fund, November 2–3, Washington, DC*.
- Hamilton, J. D. (2009, May). *Causes and consequences of the oil shock of 2007-08* (Working Paper No. 15002). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w15002> doi: 10.3386/w15002
- Hamilton, J. D. (2015, January 25). *What's driving the price of oil down?* Retrieved from <https://econbrowser.com/archives/2015/01/whats-driving-the-price-of-oil-down-2>
- Hamilton, J. D. (2021). Measuring global economic activity. *Journal of Applied Econometrics*, 36(3), 293–303.
- Jackson, L. E., Kose, M. A., Otrok, C., & Owyang, M. T. (2015, August). *Specification and Estimation of Bayesian Dynamic Factor Models: A Monte Carlo Analysis with an Application to Global House Price Comovement* (Working Papers No. 2015-31). Federal Reserve Bank of St. Louis. Retrieved from <https://ideas.repec.org/p/fip/fedlwp/2015-031.html>
- Kilian, L. (2009, June). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3), 1053–69. doi: 10.1257/aer.99.3.1053
- Kilian, L., & Lewis, L. T. (2011). Does the fed respond to oil price shocks? *The Economic Journal*, 121(555), 1047–1072. doi: 10.1111/j.1468-0297.2011.02437.x

- Kilian, L., & Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press. Retrieved from <https://books.google.com/books?id=gYdbAQAACAAJ>
- Kilian, L., & Murphy, D. P. (2012). Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market var models. *Journal of the European Economic Association*, 10(5), 1166–1188. doi: 10.1111/j.1542-4774.2012.01080.x
- Kilian, L., & Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3), 454–478. doi: 10.1002/jae.2322
- Kilian, L., & Zhou, X. (2018). Modeling fluctuations in the global demand for commodities. *Journal of International Money and Finance*, 88(C), 54–78. Retrieved from <https://ideas.repec.org/a/eee/jimfin/v88y2018icp54-78.html> doi: 10.1016/j.jimonfin.2018.07.001
- Kim, C.-J., & Nelson, C. (1998). Business cycle turning points, a new coincident index, and tests of duration dependence based on a dynamic factor model with regime switching. *The Review of Economics and Statistics*, 80(2), 188–201. Retrieved from <https://EconPapers.repec.org/RePEc:tp:restat:v:80:y:1998:i:2:p:188-201>
- Kim, C.-J., & Nelson, C. R. (1999). *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications* (Vol. 1) (No. 0262112388). The MIT Press. Retrieved from <https://ideas.repec.org/b/mtp/titles/0262112388.html>
- Kose, M. A., Otrok, C., & Whiteman, C. H. (2003, September). International business cycles: World, region, and country-specific factors. *American Economic Review*, 93(4), 1216–1239. Retrieved from <https://ideas.repec.org/a/aea/aecrev/v93y2003i4p1216-1239.html>
- Kose, M. A., Sugawara, N., & Terrones, M. E. (2020). Global recessions. *Policy Research Working Paper, No. 9172*. World Bank, Washington, DC.
- Kotwal, A., Ramaswami, B., & Wadhwa, W. (2011, December). Economic liberalization and indian economic growth: What's the evidence? *Journal of Economic Literature*, 49(4), 1152–99. doi: 10.1257/jel.49.4.1152
- Lescaroux, F. (2009, October). On the excess co-movement of commodity prices—A note about the role of fundamental factors in short-run dynamics. *Energy Policy*, 37(10), 3906–3913. Retrieved from <https://ideas.repec.org/a/eee/enepol/v37y2009i10p3906-3913.html>
- Lombardi, M. J., Osbat, C., & Schnatz, B. (2012). Global commodity cycles and linkages: a FAVAR approach. *Empirical Economics*, 43(2), 651–670. doi: 10.1007/s00181-011-0494-8
- McKay, A., & Wolf, C. K. (2023). What can time-series regressions tell us about policy counterfactuals? *Econometrica*, 91(5), 1695–1725. doi: 10.3982/ECTA21045
- Miranda-Agrippino, S., & Rey, H. (2020). U.S. Monetary Policy and the Global Financial Cycle. *The Review of Economic Studies*, 87(6), 2754–2776. doi: 10.1093/restud/rdaa019
- Nakamura, E., & Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4), 1415–1464.

- Otrok, C., & Whiteman, C. (1998). Bayesian leading indicators: Measuring and predicting economic conditions in Iowa. *International Economic Review*, 39(4), 997-1014. Retrieved from <https://EconPapers.repec.org/RePEc:ier:iecrev:v:39:y:1998:i:4:p:997-1014>
- Peersman, G., & Van Robays, I. (2012). Cross-country differences in the effects of oil shocks. *Energy Economics*, 34(5), 1532-1547. Retrieved from <https://ideas.repec.org/a/eee/eneeco/v34y2012i5p1532-1547.html> doi: 10.1016/j.eneco.2011.11.010
- Pindyck, R., & Rotemberg, J. (1990). The excess co-movement of commodity prices. *Economic Journal*, 100(403), 1173-89. Retrieved from <https://EconPapers.repec.org/RePEc:ecj:econjl:v:100:y:1990:i:403:p:1173-89>
- Plante, M. (2014). How should monetary policy respond to changes in the relative price of oil? considering supply and demand shocks. *Journal of Economic Dynamics and Control*, 44, 1 - 19. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0165188914000839> doi: <https://doi.org/10.1016/j.jedc.2014.04.002>
- Poncela, P., Senra, E., & Sierra, L. P. (2014, October). Common dynamics of nonenergy commodity prices and their relation to uncertainty. *Applied Economics*, 46(30), 3724-3735. Retrieved from <https://ideas.repec.org/a/taf/applec/v46y2014i30p3724-3735.html> doi: 10.1080/00036846.2014.939377
- Radetzki, M. (2006). The anatomy of three commodity booms. *Resources Policy*, 31(1), 56-64. Retrieved from <https://EconPapers.repec.org/RePEc:eee:jrpoli:v:31:y:2006:i:1:p:56-64>
- Rubio-Ramírez, J., Waggoner, D., & Zha, T. (2010). Structural vector autoregressions: Theory of identification and algorithms for inference. *The Review of Economic Studies*, 77(2), 665–696. Retrieved from <http://www.jstor.org/stable/40587642>
- Sarno, L., Tsiakas, I., & Ulloa, B. (2016). What drives international portfolio flows? *Journal of International Money and Finance*, 60, 53–72. doi: 10.1016/j.jimonfin.2015.03.006
- Stock, J. H., & Watson, M. W. (1998, August). *Diffusion indexes* (Working Paper No. 6702). National Bureau of Economic Research. Retrieved from <http://www.nber.org/papers/w6702> doi: 10.3386/w6702
- Topalova, P., & Khandelwal, A. (2011, August). Trade Liberalization and Firm Productivity: The Case of India. *The Review of Economics and Statistics*, 93(3), 995-1009. Retrieved from <https://ideas.repec.org/a/tpr/restat/v93y2011i3p995-1009.html>
- Vansteenkiste, I. (2009, July). *How important are common factors in driving non-fuel commodity prices? A dynamic factor analysis* (Working Paper Series No. 1072). European Central Bank. Retrieved from <https://ideas.repec.org/p/ecb/ecbwps/20091072.html>
- World Bank. (2015, July). *Commodity markets outlook* (Tech. Rep.). Washington, DC: World Bank. Retrieved from https://www.worldbank.org/content/dam/Worldbank/GEP/GEPcommodities/GEP2015c_commodity_Jul2015.pdf

World Bank. (2022, October). *Commodity markets outlook* (Tech. Rep.). Washington, DC: World Bank. Retrieved from <https://openknowledge.worldbank.org/bitstream/handle/10986/38160/CMO-October-2022.pdf>

World Bank. (2024, April). *Commodity markets outlook* (Tech. Rep.). Washington, DC: World Bank. Retrieved from <https://thedocs.worldbank.org/en/doc/464523804694923b6e827ed9d01e247f-0050012024/original/CMO-April-2024.pdf>

Yin, L., & Han, L. (2015). Co-movements in commodity prices: Global, sectoral and commodity-specific factors. *Economics Letters*, 126, 96–100.

Zhu, X. (2012, November). Understanding China's growth: Past, present, and future. *Journal of Economic Perspectives*, 26(4), 103-24. doi: 10.1257/jep.26.4.103




FIGURE 1: DSGE Model: Impulse Response Functions with Calibration Uncertainty Bands

Note: Impulse responses to one-percent structural shocks in the calibrated model of Section II. Column 1: preference shock (ϵ^{ξ}). Column 2: productivity shock (ϵ^A). Column 3: inventory-demand shock (ϵ^V). Column 4: negative energy supply shock ($-\epsilon^{z,E}$). Solid lines: baseline calibration. Shaded regions: 68 and 90 percent bands computed from the admissible subset of 500 Monte Carlo calibrations that preserves the model-implied aggregate sign pattern over the first four quarters. Output and the real commodity price index are in percentage deviations from steady state; inflation is in annualized percentage-point deviations.




FIGURE 2: GLOBAL COMMODITY PRICE FACTOR: 1970:M4–2024:M7

Note: Panel A reports the global commodity price factor in 12-month moving sums (year-over-year growth rates). The solid line is the posterior median estimated using DFM (8)–(10); dashed envelope lines trace the 68-percent Bayesian credible band (16th–84th percentile); and lighter dotted envelope lines trace the 90-percent credible band (5th–95th percentile). Panel B plots the cumulated posterior median of the global factor alongside linearly detrended log world industrial production (OECD plus 6 non-member economies). Both series are standardized to zero mean and unit standard deviation over the common sample for visual comparability. Shaded vertical intervals mark periods of oil price spikes as identified in Hamilton (2009) (amber shading) and U.S. recession periods dated by the National Bureau of Economic Research (NBER, blue-grey shading).



FIGURE 3: STRUCTURAL SHOCK COMPARISON WITH BAUMEISTER & HAMILTON (2019) (12-MONTH MOVING AVERAGE)

Note: 12-month centered moving averages of structural shocks, standardized to unit variance over the overlap sample. Top: estimated global aggregate demand shock against the Baumeister & Hamilton (2019) economic activity shock. Middle: estimated global aggregate supply shock against the Baumeister & Hamilton (2019) economic activity shock. Bottom: estimated commodity-market shock against the Baumeister & Hamilton (2019) combined oil consumption and inventory demand shocks. The comparisons are diagnostic because the external series do not uniquely validate the structural labels. The bottom-panel comparison is not exact: the estimated commodity-market shock is identified from the global commodity factor, while the Baumeister & Hamilton (2019) series is specific to crude-oil demand and inventory demand. Correlations (r) shown in each panel are computed on the smoothed series plotted; the corresponding correlations for the underlying monthly series are 0.602, 0.319, and 0.293 (see Section IV.2). Overlap sample: 1975:M2–2023:M11.




FIGURE 4: IMPULSE RESPONSES SCALED TO UNIT MEDIAN COMMODITY-FACTOR IMPACT

Note: Posterior-median impulse responses under the baseline $H = 3$ sign-and-path restrictions. For each shock, all accepted draws are divided by the posterior-median on-impact response of the commodity price factor, so that the median normalized commodity-factor impact equals one; individual impacts remain dispersed. The darker and lighter regions are the 68% and 90% Bayesian credible bands. The 3,277,292 accepted parameter-rotation draws come from 600 posterior parameter draws and 300,000 candidate rotations per draw. The bands combine reduced-form parameter uncertainty and admissible-rotation uncertainty, conditional on the estimated DFM factor and maintained specifications. Industrial production and commodity factor responses are cumulated.




FIGURE 5: COUNTERFACTUAL HISTORICAL DECOMPOSITION: KEY EPISODES

Note: Solid line: cumulated global commodity price factor, reported on the cumulated-factor scale. Dashed line: counterfactual path without the indicated structural shock. Shaded area: cumulative contribution of the shock. Each row corresponds to a historical episode; columns correspond to the three structural shocks. Counterfactuals are constructed following Kilian & Lewis (2011); see Appendix C3 for details.

TABLE 1: Variance Decompositions for Commodity Prices

Group	Commodity	1970–2024			1970–1999			2000–2024		
		Global	Group	Idio.	Global	Group	Idio.	Global	Group	Idio.
Energy	Avg.	8.5	35.8	55.6	5.2	38.5	56.3	16.2	32.2	51.6
	<i>Crude Oil Average</i>	16.7	4.9	78.4	8.7	14.0	77.3	31.4	1.4	67.2
	<i>Coal Australian</i>	8.1	6.8	85.1	1.2	14.4	84.3	15.8	5.6	78.7
	<i>Natural Gas Index</i>	0.8	95.9	3.3	5.5	87.2	7.3	1.3	89.7	8.9
Metals	Avg.	23.9	22.3	53.8	8.1	25.7	66.2	44.7	18.0	37.2
	<i>Aluminum</i>	23.9	19.0	57.1	5.2	23.9	70.8	49.9	13.5	36.6
	<i>Copper</i>	33.4	27.3	39.3	16.4	34.8	48.8	55.8	18.4	25.8
	<i>Lead</i>	18.1	26.5	55.3	12.8	28.1	59.1	30.5	20.8	48.7
	<i>Tin</i>	26.8	8.6	64.7	5.3	11.1	83.6	50.9	3.5	45.6
	<i>Nickel</i>	21.5	10.9	67.5	3.6	12.7	83.7	41.6	7.1	51.3
	<i>Zinc</i>	19.7	41.3	38.9	5.4	43.5	51.2	39.8	44.8	15.4
Precious Metals	Avg.	23.4	45.4	31.2	17.8	52.7	29.6	31.3	40.8	27.8
	<i>Gold</i>	16.8	37.9	45.3	22.2	44.5	33.3	13.6	47.3	39.1
	<i>Platinum</i>	30.5	23.1	46.4	11.7	52.8	35.4	47.7	10.9	41.5
	<i>Silver</i>	22.8	75.2	2.0	19.3	60.7	20.0	32.7	64.3	3.0
Fertilizers	Avg.	3.5	37.2	59.3	8.0	35.7	56.2	0.9	38.4	60.6
	<i>DAP</i>	4.1	45.3	50.6	8.3	44.8	46.9	0.6	53.2	46.2
	<i>TSP</i>	5.9	90.6	3.5	9.0	89.1	1.9	1.9	82.4	15.7
	<i>Urea</i>	3.2	11.4	85.4	11.6	7.4	81.0	0.6	16.4	82.9
	<i>Potassium Chloride</i>	0.8	1.7	97.6	3.2	1.7	95.1	0.7	1.7	97.7
Grains	Avg.	10.1	30.7	59.2	8.6	32.2	59.1	14.6	27.9	57.5
	<i>Maize</i>	11.4	86.3	2.4	9.3	88.8	1.9	15.7	80.8	3.5
	<i>Rice Thai 5%</i>	2.2	0.3	97.5	1.4	0.9	97.7	3.1	0.3	96.6
	<i>Wheat U.S. HRW</i>	9.7	16.7	73.6	8.2	16.9	74.9	12.4	16.1	71.5
	<i>Soybeans</i>	17.0	19.5	63.5	15.6	22.4	62.0	27.3	14.5	58.3
Food Oils	Avg.	21.6	48.1	30.4	15.0	54.7	30.3	29.5	41.0	29.5
	<i>Coconut Oil</i>	19.1	22.9	58.0	13.8	84.1	2.2	25.1	21.0	53.9
	<i>Palm Oil</i>	24.0	73.3	2.7	16.2	25.4	58.4	33.9	61.0	5.1
Other Foods	Avg.	4.0	17.3	78.7	3.8	19.1	77.1	5.0	19.5	75.5
	<i>Banana U.S.</i>	1.0	0.3	98.7	0.7	92.8	6.4	1.0	0.3	98.7
	<i>Orange</i>	0.4	0.2	99.4	0.8	1.5	97.7	3.0	0.3	96.7
	<i>Beef</i>	4.8	0.5	94.7	4.1	0.3	95.6	4.7	1.0	94.3
	<i>Chicken</i>	1.9	0.3	97.8	0.9	0.5	98.6	2.5	95.2	2.3
	<i>Sugar World</i>	11.7	85.3	3.0	12.3	0.4	87.3	13.9	0.5	85.6
Beverages	Avg.	5.5	30.0	64.5	6.6	28.0	65.5	7.7	28.8	63.5
	<i>Cocoa</i>	8.5	89.0	2.5	13.6	83.0	3.4	8.0	84.7	7.3
	<i>Coffee Arabica</i>	5.1	0.5	94.5	3.7	0.3	96.0	9.5	1.2	89.3
	<i>Tea Avg 3 Auctions</i>	3.0	0.4	96.7	2.4	0.6	97.0	5.6	0.4	94.0
Raw Materials	Avg.	11.2	18.7	70.0	7.9	14.9	77.2	14.3	18.2	67.6
	<i>Cotton A Index</i>	13.7	0.2	86.1	7.0	0.3	92.7	16.7	0.8	82.6
	<i>Logs Cameroon</i>	6.3	0.5	93.2	8.1	0.6	91.3	15.1	80.9	4.0
	<i>Logs Malaysian</i>	1.4	0.2	98.4	2.8	0.8	96.4	0.6	7.8	91.6
	<i>Rubber RSS3</i>	34.4	0.9	64.7	20.9	72.0	7.1	37.4	0.9	61.6
	<i>Tobacco U.S. Import</i>	0.2	92.0	7.8	0.6	0.7	98.7	1.5	0.5	98.0
Overall Avg.		12.3	29.0	58.7	8.3	30.4	61.3	18.6	27.1	54.3

Note: Entries report the percentage of the variance of monthly real commodity-price growth explained by the global factor, the commodity-group factor, and idiosyncratic disturbances. Bold rows report simple within-group averages; the final row reports the overall average across all 35 commodities. DAP denotes diammonium phosphate; TSP, triple superphosphate; HRW, hard red winter wheat; and RSS3, ribbed smoked sheet grade 3 rubber.

Commodity Price Cycles: Commonalities, Heterogeneities, and Drivers

Alain Kabundi and Hamza Zahid

ONLINE APPENDIX

Supplementary Material

Contents

Appendix A. Literature Review

Appendix B. Theoretical Model

Appendix C. Estimation, Identification, and Decomposition Details

Appendix D. Robustness and Diagnostic Exercises

Appendix E. Additional Figures and Tables

A Literature Review

This paper relates to empirical work on common movements in commodity prices. Pindyck & Rotemberg (1990) study comovement among largely unrelated commodities and show that their joint price movements exceed what can be explained by common macroeconomic shocks. They label this pattern “excess” comovement. The broad commodity-price upswing of the early 2000s renewed interest in this question, with subsequent work using a range of methods to study the existence and drivers of common price cycles in commodity markets (see, among others, Ai et al., 2006; Cuddington & Jerrett, 2008; Lescaroux, 2009; Alquist et al., 2020; Baffes & Kabundi, 2023).

Factor models provide one way to measure commodity-price comovement. They extract a small number of unobserved factors from a large panel of prices, and the estimated factors summarize the common variation in the data. Together with the factor loadings, these estimates can be used to assess the extent of synchronization in commodity price fluctuations. This approach is useful for large datasets (see Stock & Watson, 1998; Bernanke & Boivin, 2003). When the data have a natural group structure, multilevel dynamic factor models can separate overall comovement from within-group comovement.³⁸ This structure is well suited to commodity markets because prices can be organized into sectors such as energy, metals, and grains.

Only a few studies use multilevel dynamic factor models to account explicitly for sectoral heterogeneity in commodity-price comovement. Yin & Han (2015) use a monthly dataset of 24 commodities over 1991–2014 to decompose commodity-price variation into global, sectoral, and idiosyncratic factors. They find that global and group factors account, on average, for 16.9 and 32.6 percent of commodity-price variation, respectively. Similarly, Delle Chiaie, Ferrara, & Giannone (2022) extract global and block-specific factors from a monthly dataset of 52 commodities over 1980–2015 and find that one global factor summarizes much of the variation in commodity prices.

A related strand of the literature studies the drivers of common commodity-price fluctuations. Byrne, Sakemoto, & Xu (2017) estimate global and sector-specific factors using a hierarchical factor model for 38 non-energy commodity prices over 1974–2014. They include the estimated factors in factor-augmented vector autoregression (FAVAR) models with selected macroeconomic variables to study how different shocks affect common movements in non-energy commodity prices. They find responses to innovations in the global business cycle, real interest rates, and uncertainty, with

³⁸For example, Kose, Otrok, & Whiteman (2003) use a multilevel dynamic factor model to study business-cycle fluctuations across 60 countries and 7 regions of the world. Exploiting the block structure of the data, they interpret the estimated factors as world, region-specific, and country-specific factors. For other applications of multilevel dynamic factor models, see Jackson et al. (2015), Miranda-Agrippino & Rey (2020), Sarno, Tsiakas, & Ulloa (2016), and Ha et al. (2017).

heterogeneous responses across sectors.³⁹ Bilgin & Ellwanger (2017) estimate global and block-specific factors in a sample of 42 commodity price series over 1981–2017. In line with Delle Chiaie et al. (2022), they compare the estimated global factor with measures of global economic activity, including the Baltic Dry Index and global industrial production, and document a positive association between the global factor and global activity.

B Theoretical Model

This appendix presents a stylized New Keynesian dynamic stochastic general equilibrium (DSGE) model that motivates the sign restrictions used in the empirical factor-augmented vector autoregression (FAVAR). Preference and productivity shocks illustrate aggregate-demand and aggregate-supply disturbances, while inventory-demand and commodity-supply shocks illustrate disturbances originating within commodity markets. Other mechanisms may produce the same empirical sign patterns.

We work with a representative world economy. Final goods producers use labor and a composite commodity input, prices are sticky in the Calvo sense, and commodity markets feature upward-sloping supply and inventories. The model is intentionally stylized and is not meant to provide a full quantitative account of global commodity markets. Its role is narrower: to provide an economically coherent benchmark for the qualitative sign patterns used for identification and to provide a compact interpretation of the global and group factors in the factor model.

B1 Model Environment

B1.1 Households

A representative household maximizes expected discounted utility:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\xi_t \frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right], \quad (13)$$

where C_t is consumption, N_t is labor supply, $\beta \in (0, 1)$ is the discount factor, $\sigma > 0$ is relative risk aversion, $\varphi > 0$ is the inverse Frisch elasticity, $\chi > 0$ scales the disutility of labor, and ξ_t is a

³⁹Other studies use related frameworks to examine drivers of common commodity-price movements, but often focus on a single common factor extracted from a subset of commodity prices (see, among others, Vansteenkiste, 2009; Lombardi, Osbat, & Schnatz, 2012; Byrne, Fazio, & Fiess, 2013; Poncela, Senra, & Sierra, 2014).

preference shifter that follows

$$\log \xi_t = \rho_\xi \log \xi_{t-1} + \varepsilon_t^\xi, \quad \varepsilon_t^\xi \sim N(0, \sigma_\xi^2). \quad (14)$$

The household budget constraint is

$$P_t C_t + B_t = W_t N_t + R_{t-1} B_{t-1} + D_t + T_t, \quad (15)$$

where P_t is the final-goods price level, W_t the nominal wage, B_t one-period nominal bonds, R_t the gross nominal interest rate, D_t firm profits, and T_t lump-sum transfers.

The optimality conditions imply the standard Euler equation and labor supply relation:

$$\xi_t C_t^{-\sigma} = \beta R_t \mathbb{E}_t \left[\xi_{t+1} C_{t+1}^{-\sigma} \frac{P_t}{P_{t+1}} \right], \quad (16)$$

$$\frac{W_t}{P_t} = \chi \frac{N_t^\varphi}{\xi_t C_t^{-\sigma}}. \quad (17)$$

A positive preference shock, $\varepsilon_t^\xi > 0$, raises the marginal utility of current consumption, shifts aggregate demand outward, and increases firms' demand for commodity inputs.

B1.2 Final Goods Producers

A continuum of monopolistically competitive firms indexed by $j \in [0, 1]$ produce differentiated goods with labor and a commodity composite:

$$Y_t(j) = A_t \left[(1 - \alpha)^{\frac{1}{\mu}} N_t(j)^{\frac{\mu-1}{\mu}} + \alpha^{\frac{1}{\mu}} M_t(j)^{\frac{\mu-1}{\mu}} \right]^{\frac{\mu}{\mu-1}}, \quad (18)$$

where A_t is total factor productivity, $N_t(j)$ labor input, $M_t(j)$ the commodity composite, $\alpha \in (0, 1)$ the steady-state commodity share, and $\mu > 0$ the elasticity of substitution between labor and commodities.

Productivity follows

$$\log A_t = \rho_A \log A_{t-1} + \varepsilon_t^A, \quad \varepsilon_t^A \sim N(0, \sigma_A^2). \quad (19)$$

A positive productivity shock raises the marginal product of all inputs, including commodities, while lowering real marginal cost through the supply side of the economy.

B1.3 Commodity Composite and Supply

The commodity composite is a two-level constant-elasticity-of-substitution (CES) aggregate. At the upper level,

$$M_t = \left[\omega_E^{\frac{1}{\eta}} M_{E,t}^{\frac{\eta-1}{\eta}} + \omega_L^{\frac{1}{\eta}} M_{L,t}^{\frac{\eta-1}{\eta}} + \omega_F^{\frac{1}{\eta}} M_{F,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \quad (20)$$

where $g \in \{E, L, F\}$ indexes energy, metals, and food, $\omega_g > 0$ with $\sum_g \omega_g = 1$, and $\eta > 0$ is the elasticity of substitution across groups.

At the lower level,

$$M_{g,t} = \left[\sum_{i \in \mathcal{J}_g} \omega_{i|g}^{\frac{1}{\eta_g}} M_{i,t}^{\frac{\eta_g-1}{\eta_g}} \right]^{\frac{\eta_g}{\eta_g-1}}, \quad g \in \{E, L, F\}, \quad (21)$$

where $\mathcal{J}_g$ is the set of commodities in group g , $\omega_{i|g} > 0$ with $\sum_{i \in \mathcal{J}_g} \omega_{i|g} = 1$, and $\eta_g > 0$ is the within-group elasticity of substitution.

The associated price indices are

$$P_t^M = \left[\omega_E (P_{E,t}^M)^{1-\eta} + \omega_L (P_{L,t}^M)^{1-\eta} + \omega_F (P_{F,t}^M)^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad (22)$$

$$P_{g,t}^M = \left[\sum_{i \in \mathcal{J}_g} \omega_{i|g} P_{i,t}^{1-\eta_g} \right]^{\frac{1}{1-\eta_g}}, \quad g \in \{E, L, F\}. \quad (23)$$

Each individual commodity is supplied according to an upward-sloping schedule:

$$Q_{i,t} = \bar{Q}_i \left(\frac{P_{i,t}}{P_t} \right)^{\zeta_g} \exp(z_{g,t}) \exp(u_{i,t}), \quad (24)$$

where $\bar{Q}_i > 0$ is a supply scale parameter, $\zeta_g > 0$ is the group-specific supply elasticity, $z_{g,t}$ a group shock, and $u_{i,t}$ an idiosyncratic shock. The schedule is written for an individual commodity as an illustration; the implementation specifies supply and market clearing at the group level. The group shock follows

$$z_{g,t} = \rho_{z_g} z_{g,t-1} + \varepsilon_t^{z_g}, \quad \varepsilon_t^{z_g} \sim N(0, \sigma_{z_g}^2). \quad (25)$$

This nested structure also motivates the factor model. Disturbances common to all groups move the broad index P_t^M , group-specific disturbances move $P_{g,t}^M$, and commodity-specific disturbances move individual prices. These channels motivate distinguishing global, group-specific, and idiosyncratic components in the empirical factor model. The correspondence is interpretive: group disturbances

feed back into the aggregate index through the CES aggregator, and a movement in a price index is not equivalent to an estimated factor.

B1.4 Inventories and Commodity-Market Demand

To capture disturbances internal to commodity markets, the model includes commodity inventories. Inventories provide storage services, such as precautionary availability and the option to meet future demand. Desired inventory stocks for commodity group g are

$$\hat{s}_{g,t}^* = v_t + v_{g,t} - \eta_S \hat{p}_{g,t}^{m,r}, \quad (26)$$

where v_t is a global inventory-demand state, $v_{g,t}$ is a group-specific inventory-demand state, and $\eta_S > 0$ is the price elasticity of desired stocks. The global component follows

$$v_t = \rho_v v_{t-1} + \varepsilon_t^v, \quad \varepsilon_t^v \sim N(0, \sigma_v^2). \quad (27)$$

The group-specific component follows

$$v_{g,t} = \rho_{v,g} v_{g,t-1} + \varepsilon_t^{v_g}, \quad \varepsilon_t^{v_g} \sim N(0, \sigma_{v_g}^2). \quad (28)$$

Actual inventory stocks adjust gradually:

$$\hat{s}_{g,t} = \rho_S \hat{s}_{g,t-1} + (1 - \rho_S) \hat{s}_{g,t}^*. \quad (29)$$

Physical inventory accumulation is

$$I_{g,t} = S_{g,t} - (1 - \delta_S) S_{g,t-1}, \quad (30)$$

where δ_S is storage depreciation or replacement demand. Accumulation has steady-state value $\bar{I}_g = \delta_S \bar{S}_g$, so its log deviation satisfies

$$\hat{i}_{g,t} = \frac{\hat{s}_{g,t} - (1 - \delta_S) \hat{s}_{g,t-1}}{\delta_S}. \quad (31)$$

Commodity-market clearing is

$$Q_{g,t} = M_{g,t} + I_{g,t}. \quad (32)$$

A positive realization of v_t raises desired inventories, increases current stockbuilding, and raises commodity prices without directly reflecting stronger final demand or higher productive capacity. In

the empirical model, the identified commodity-market shock should therefore be interpreted as a composite of inventory-demand and supply disturbances originating within commodity markets.

B1.5 Price Setting and Policy

Final goods are a constant-elasticity aggregate of differentiated varieties with elasticity $\varepsilon > 1$, so firm j faces the demand curve

$$Y_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\varepsilon} Y_t, \quad P_t = \left[\int_0^1 P_t(j)^{1-\varepsilon} dj \right]^{\frac{1}{1-\varepsilon}}. \quad (33)$$

Prices are sticky in the Calvo sense. In each period, a fraction $(1 - \theta)$ of firms reset prices, while the remaining fraction θ keep prices unchanged. A firm resetting at t chooses $P_t^*(j)$ to maximize the expected discounted value of profits:

$$\max_{P_t^*(j)} \mathbb{E}_t \sum_{k=0}^{\infty} (\theta\beta)^k \frac{\Lambda_{t+k}}{\Lambda_t} [P_t^*(j)Y_{t+k|t}(j) - \mathcal{C}_{t+k}(Y_{t+k|t}(j))], \quad (34)$$

where $\Lambda_t = \xi_t C_t^{-\sigma} / P_t$ is the marginal utility of nominal income. The price index evolves as $P_t^{1-\varepsilon} = \theta P_{t-1}^{1-\varepsilon} + (1 - \theta)(P_t^*)^{1-\varepsilon}$, and the zero-inflation steady state implies $\overline{MC} = (\varepsilon - 1)/\varepsilon$.

Taking input prices as given, firms minimize cost subject to (18). Real marginal cost is

$$MC_t = \frac{1}{A_t} \left[(1 - \alpha) \left(\frac{W_t}{P_t} \right)^{1-\mu} + \alpha \left(\frac{P_t^M}{P_t} \right)^{1-\mu} \right]^{\frac{1}{1-\mu}}, \quad (35)$$

The model is normalized so that steady-state real input prices are equal, $\bar{w}^r = \bar{p}^{m,r} = \overline{MC}$, with $\overline{MC} = (\varepsilon - 1)/\varepsilon$ and $\bar{A} = 1$. Under this normalization the steady-state commodity expenditure share equals α .

The associated input demands are

$$N_t = (1 - \alpha) \left(\frac{W_t/P_t}{MC_t \cdot A_t} \right)^{-\mu} \frac{Y_t}{A_t}, \quad (36)$$

$$M_t = \alpha \left(\frac{P_t^M/P_t}{MC_t \cdot A_t} \right)^{-\mu} \frac{Y_t}{A_t}. \quad (37)$$

Equations (36)–(37) are firm-level cost-minimizing demands; under Calvo price dispersion they hold in the aggregate to first order around the zero-inflation symmetric steady state. Equation (37) is

the key derived-demand relation: commodity demand rises with output and falls with the relative commodity price. Because input demand is scaled by $1/A_t$, the effect of productivity depends on μ , as set out below.

Monetary policy follows a Taylor rule:

$$R_t = R^{1-\phi_R} R_{t-1}^{\phi_R} \left[\left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}_t} \right)^{\phi_y} \right]^{1-\phi_R} \exp(\varepsilon_t^R), \quad (38)$$

where $\Pi_t = P_t/P_{t-1}$ is gross inflation, $\bar{Y}_t = \bar{Y}$ is steady-state output, $\phi_R \in (0, 1)$ captures inertia, $\phi_\pi > 1$ satisfies the Taylor principle, and $\phi_y \geq 0$ captures the response to the output gap.

B1.6 Market Clearing

Goods market clearing requires

$$Y_t = C_t, \quad (39)$$

and each commodity market clears according to (32).

We assume that commodity suppliers and inventory holders are owned by the representative household, so that their net cash flows, including commodity-sale receipts and expenditures on inventory acquisition, enter the budget constraint (15) through D_t . Extraction and storage are assumed not to use labor or final goods from (18), so goods-market clearing is $Y_t = C_t$ while physical commodity flows clear separately through (32). The supply and inventory schedules are reduced-form specifications rather than solutions to an explicit storage problem.

B2 Log-Linearized Equilibrium

We log-linearize around a zero-inflation steady state. Let $\hat{x}_t \equiv \log X_t - \log \bar{X}$ denote deviations from steady state.

B2.1 Dynamic IS Curve

The Euler equation implies

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\sigma} (\hat{r}_t - \mathbb{E}_t \hat{r}_{t+1}) + \frac{1}{\sigma} (\hat{\xi}_t - \rho_\xi \hat{\xi}_t), \quad (40)$$

where $\hat{y}_t = \hat{c}_t$, $\hat{r}_t$ is the deviation of the nominal rate, and $\hat{\pi}_t = \hat{p}_t - \hat{p}_{t-1}$ is inflation.

B2.2 New Keynesian Phillips Curve

Calvo pricing yields

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa \widehat{mc}_t, \quad (41)$$

with $\kappa = \frac{(1-\theta)(1-\theta\beta)}{\theta}$. Linearizing (35) gives

$$\widehat{mc}_t = -\hat{a}_t + (1-\alpha)\hat{w}_t^r + \alpha\hat{p}_t^{m,r}, \quad (42)$$

where $\hat{w}_t^r = \hat{w}_t - \hat{p}_t$ is the real wage and $\hat{p}_t^{m,r} = \hat{p}_t^m - \hat{p}_t$ is the real commodity price. Under the steady-state normalization above the coefficients are the steady-state cost shares, so the expression is the first-order linearization of (35) for any μ : productivity lowers marginal cost directly, while higher commodity prices raise it.

B2.3 Commodity Price Dynamics

Linearizing commodity demand (37) gives

$$\hat{m}_t = \hat{d}_t - \mu \hat{p}_t^{m,r}, \quad \hat{d}_t \equiv \hat{y}_t + (\mu - 1)\hat{a}_t + \mu \widehat{mc}_t, \quad (43)$$

where $\hat{d}_t$ collects the demand-side terms. Because input demand in (37) is scaled by $1/A_t$, productivity is input-saving at given output and enters with coefficient $(\mu - 1)$, which vanishes at the Cobb–Douglas benchmark; marginal cost enters because firms' input demand rises with the scale of production costs. Aggregate commodity supply satisfies

$$\hat{q}_t = \zeta \hat{p}_t^{m,r} + \hat{z}_t, \quad (44)$$

and market clearing with inventories implies

$$\hat{q}_t = (1 - \lambda_S)\hat{m}_t + \lambda_S\hat{i}_t, \quad (45)$$

where λ_S is the steady-state share of inventory accumulation and replacement demand in total physical commodity demand. Combining these expressions yields the equilibrium real commodity price:

$$\hat{p}_t^{m,r} = \frac{(1 - \lambda_S)\hat{d}_t + \lambda_S\hat{i}_t - \hat{z}_t}{\zeta + (1 - \lambda_S)\mu}. \quad (46)$$

This equation is central for the identification argument. Commodity prices rise with the demand shifter $\hat{d}_t$ and with inventory accumulation, and fall with commodity supply. The shifter rises with output. The direct productivity coefficient is positive when $\mu > 1$, zero at the Cobb–Douglas benchmark, and negative when $\mu < 1$. At the benchmark the explicit productivity term in the shifter vanishes, and the commodity-price response depends on the equilibrium adjustment of output, marginal cost, and inventory accumulation. The calibrated impulse responses in Figure 1 confirm that this equilibrium effect is positive. Since output, marginal cost, and inventory accumulation are endogenous, the shocks are identified by their joint effects on output, inflation, and commodity prices rather than by any single term in (46).

The aggregate and group price indices satisfy

$$\hat{p}_t^m = \sum_g s_g \hat{p}_{g,t}^m, \quad (47)$$

$$\hat{p}_{g,t}^m = \sum_{i \in \mathcal{I}_g} s_{i|g} \hat{p}_{i,t}, \quad (48)$$

where s_g and $s_{i|g}$ are steady-state expenditure shares. At the individual level, the nested structure in (20)–(21) implies the demand relation

$$\hat{m}_{i,t} = \hat{d}_t + (\eta - \mu) \hat{p}_t^{m,r} + (\eta_g - \eta) \hat{p}_{g,t}^{m,r} - \eta_g \hat{p}_{i,t}^r, \quad (49)$$

where $\hat{p}_{g,t}^{m,r}$ is the real group price index, and η and η_g are the across- and within-group substitution elasticities.

The nested demand system links individual commodity demand to economy-wide conditions and to relative prices across and within commodity groups. Together with group-specific and commodity-specific supply disturbances, these channels motivate distinguishing global, group-specific, and idiosyncratic components of commodity-price variation in the empirical factor model. This correspondence is interpretive rather than an exact mapping from the theoretical model to the estimated factors.

The log-linearized Taylor rule, with $\hat{y}_t^{\text{gap}} = \hat{y}_t$ denoting output relative to its steady state, is

$$\hat{r}_t = \phi_R \hat{r}_{t-1} + (1 - \phi_R) (\phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t^{\text{gap}}) + \varepsilon_t^R. \quad (50)$$

B3 Shock Transmission and Sign Restrictions

Figure 1 reports impulse responses from the calibrated model for the four disturbances. The preference, productivity and inventory-demand shocks reproduce the three aggregate sign patterns used for identification, as described in Section II.

Two features of the commodity-market case are specific to the model. First, a positive commodity supply shock raises output and lowers both inflation and commodity prices. Because the identified shocks are normalized to raise commodity prices, such a disturbance enters the decomposition as a negative realization of the commodity-market shock rather than as a separate shock. Second, a negative sector-specific commodity supply disturbance, such as an energy supply disruption, raises commodity prices and inflation and lowers output, which makes it similar to the inventory-demand shock at the aggregate level. The identified commodity-market shock should therefore be interpreted as a composite of disturbances originating within commodity markets rather than as a pure inventory-demand shock.

B4 Calibration

Table 2 reports the baseline calibration used for the impulse-response exercises. The values are conventional for a quarterly New Keynesian model and are chosen to preserve the sign logic emphasized in the paper.

Four calibration choices matter most for the qualitative results. First, the commodity share $\alpha = 0.10$ keeps the output effect of commodity-market shocks small while allowing input-cost effects to matter for inflation. Second, we set $\mu = 1$, implemented in Dynare as 1.001 to approximate the Cobb–Douglas benchmark without the CES singularity at exactly one. Third, we use $\theta = 0.60$, which implies enough price rigidity for productivity shocks to reduce inflation while still allowing output to expand sufficiently for the derived-demand channel to operate. Fourth, the inventory parameters imply gradual stock adjustment and positive replacement demand. This keeps the inventory-demand shock from acting as a one-period demand shift and gives the shock a transparent physical transmission channel through stockbuilding. The steady-state inventory share follows from the stock ratio and depreciation, $\lambda_S = \delta_{SS}/(1 + \delta_{SS})$, which is about 0.074 at the baseline values. The labor-disutility parameter χ scales the disutility of labor and is set to match the steady-state labor supply condition. With steady-state output and consumption normalized to one, $\bar{A} = 1$, and the input-price normalization above, steady-state labor is $\bar{N} = 1 - \alpha$ and $\chi = \overline{MC}/(1 - \alpha)^\varphi$, which is about 0.93 at the baseline values.

The Monte Carlo calibration exercise underlying Figure 1 draws 500 parameterizations independently and uniformly from the following ranges: $\sigma \in [1, 2.5]$, $\varphi \in [0.5, 2]$, $\alpha \in [0.05, 0.15]$, $\mu \in [0.8, 1.3]$, $\varepsilon \in [4, 11]$, $\theta \in [0.4, 0.75]$, $\eta \in [0.1, 0.6]$, $\zeta \in [0.05, 0.5]$, $\phi_R \in [0, 0.8]$, $\phi_\pi \in [1.1, 2.5]$, $\phi_y \in [0, 0.5]$, $\rho_A \in [0.85, 0.99]$, $\rho_\xi \in [0.5, 0.9]$, $\rho_v \in [0.88, 0.995]$, $\rho_S \in [0.7, 0.95]$, $\eta_S \in [0.05, 0.3]$, $\delta_S \in [0.04, 0.18]$, $s_S \in [0.4, 1.2]$, and $\rho_z \in [0.4, 0.9]$. Of the 500 draws, 214 satisfy the model-implied aggregate sign pattern over the first four quarters; the shaded bands in Figure 1 are computed from this admissible subset. The ranges include values of μ below one; the reported bands use the calibrations that satisfy the joint sign restrictions.

C Estimation, Identification, and Decomposition Details

This appendix reports baseline estimation, identification, and decomposition details; robustness and validation exercises are reported separately in Appendix D.

C1 Dynamic Factor Model: Estimation Details

The state-space system defined by equations (8)–(10) is estimated using the Bayesian approach of Kim & Nelson (1998).⁴⁰ The estimation objective is to infer from the observed data: (1) the path of the common factors F_t and (2) all unknown model parameters. The Bayesian approach treats these as random variables, and inference is based on obtaining the joint and marginal posterior distributions of all factors and parameters given the observed commodity price data.

Since the joint posterior distribution is not analytically obtainable, we employ Gibbs sampling. The observation vector at date t is $\mathbf{Y}_t$, and $\mathbf{Y}_{1:T}$ denotes the full sample. The Gibbs sampler proceeds by: (i) drawing from the conditional distribution of the model parameters given the data and the factors $\mathbf{F}_t$, and then (ii) drawing from the conditional distribution of the factors $\mathbf{F}_t$ given the data and the parameter draws. Conditional on $\mathbf{F}_t$, generating the unknown parameters of the observation and state transition equations ($\mathbf{B}$, $\Psi(L)$, $\Phi(L)$, Ω) is a standard application of Bayesian linear regression.⁴¹ The AR coefficients are sampled using a Metropolis-Hastings algorithm (see Chib & Greenberg, 1994; Otrok & Whiteman, 1998). The generation of the factor vector $\mathbf{F}_t$ is based on

⁴⁰The classical approach to state-space modeling, based on maximizing the likelihood function, can be computationally inefficient in large-scale models. The Bayesian approach based on Gibbs sampling works with smaller components of the model by drawing from conditional distributions of the parameters (Blake & Mumtaz, 2012). The parametric state-space approach to estimating factor models also yields more accurate variance decomposition estimates than methods based on principal component analysis (PCA) (Jackson, Kose, Otrok, & Owyang, 2015).

⁴¹We collect the AR coefficients from the observation equation in the diagonal lag polynomial matrix $\Psi(L)$ and the error term variance-covariance matrix as $E(\zeta_t \zeta_t') = \Omega$.

the multimove Gibbs-sampling (forward-backward) algorithm of Carter and Kohn (1994), which draws the whole vector from the joint distribution $p(F_1, \dots, F_T \mid \mathbf{Y}_{1:T})$ using the Kalman filter and a backward recursion.⁴²

The estimation procedure can be summarized as follows:

1. Conditional on $\mathbf{F}_t$, sample $\mathbf{B}$ and Ω from their posterior distributions.
2. Conditional on $\mathbf{F}_t$ and the loadings, sample the idiosyncratic autoregressive coefficients $\Psi(L)$ from the commodity-specific residuals.
3. Conditional on $\mathbf{F}_t$, sample $\Phi(L)$ from its posterior distribution (the state-innovation covariance is normalized to the $K \times K$ identity matrix, I_K , for identification).
4. Conditional on the state-space parameters, sample F_t from its posterior distribution using the Kalman filter and backward recursion.
5. Repeat the preceding steps until convergence.

Prior specification. The autoregressive coefficients in equations (9) and (10) have Gaussian priors with zero mean and covariance $0.1\mathbf{I}$, truncated to the stationary region. The factor loadings in equation (8) have $N(0, 1)$ priors, with the zero restrictions described in Section III.2. The innovation variances have inverted-gamma priors of the form $d_0/\chi_{\nu_0}^2$, with $d_0 = 0.25^2$ and $\nu_0 = \lceil 0.05 T \rceil$, where T is the number of observations in the estimation sample.⁴³ The sampler retains 10,000 draws after a burn-in of 2,500 draws.

Factor orientation and variance shares. The zero restrictions and the normalization of the state-innovation covariance leave the sign of each factor undetermined. We orient the global factor so that the loading of the first commodity, crude oil, is positive, and each group factor so that the loading of the first commodity in the group is positive; the sampler reverses the sign of a factor and its loadings when the reference loading is persistently negative. Table 1 reports posterior means of the model-implied variance shares. For each draw, the share attributed to a factor is its squared loading times the unconditional variance of the factor process, divided by the total implied variance of the commodity's price growth. The full-sample and subsample columns are obtained from separate estimations of the corresponding samples.

⁴²The multimove Gibbs sampling procedure is computationally faster and more efficient than singlemove alternatives (Kim & Nelson, 1999). For detailed expositions, see Kim & Nelson (1999), Blake & Mumtaz (2012) and Jackson et al. (2015).

⁴³Results are stable under tighter prior specifications.

C2 FAVAR: Estimation and Identification Details

Input series. The VAR is estimated on the monthly log growth rate of world industrial production and the OECD total consumer price inflation rate measured as an annual growth rate, both standardized to zero mean and unit variance, together with the posterior-mean global commodity factor from the dynamic factor model, which enters without further transformation.

Posterior distributions. For the reduced-form VAR parameters, we follow the Gaussian–inverse-Wishart sampling convention of Kilian & Murphy (2012), without coefficient shrinkage. Let $n = T - p$ denote the number of usable observations and $\mathbf{W}$ the $d \times n$ regressor matrix, which contains twelve lags of each of the three variables and a constant, so $d = 37$. Let $\hat{\mathbf{A}}$ be the ordinary least squares (OLS) coefficient matrix and $\hat{\Sigma}_u = \hat{\mathbf{U}}\hat{\mathbf{U}}'/(n - d)$ the residual covariance estimate. In the scale–degrees-of-freedom parameterization, each covariance draw satisfies $\Sigma_u^{(r)} \sim \mathcal{IW}(n\hat{\Sigma}_u, n)$, and the conditional coefficient draw satisfies $\text{vec}(\mathbf{A}^{(r)}) \mid \Sigma_u^{(r)}, \mathbf{Y} \sim \mathcal{N}\left[\text{vec}(\hat{\mathbf{A}}), (\mathbf{W}\mathbf{W}')^{-1} \otimes \Sigma_u^{(r)}\right]$.

We evaluate 300,000 candidate rotations for each of 600 posterior parameter draws. Under the baseline $H = 3$ sign-and-path restrictions, 3,277,292 parameter-rotation draws are accepted. Finite-horizon impulse responses retain unstable coefficient draws. By contrast, the historical-decomposition uncertainty exercise in Appendix D5 conditions its reported bands and probabilities on stable models, which account for 98.8 percent of its accepted parameter-rotation draws. Including the remaining unstable models barely changes the episode probabilities.

C2.1 Multi-Period Sign Restrictions

The baseline identification imposes sign restrictions from impact through month 3 ($H = 3$, horizons 0:3). Following Kilian & Murphy (2012), who argue that impact-only restrictions may admit implausible structural models, we assess robustness by comparing this baseline with the impact-only case ($H = 0$) and with tighter restriction horizons $H = 6$ and $H = 12$. Since industrial production and the global commodity factor enter the VAR in month-on-month growth rates, we impose the multi-period restrictions on the cumulative impulse responses for these two variables; that is, the cumulated growth-rate responses must maintain the required sign at every horizon up to H . This ensures that the implied level effects on output and the global commodity factor preserve the expected direction. For inflation, which enters in levels, the restriction is imposed directly on the pointwise response at each horizon.

In the notation below, IP denotes world industrial production, CF the global commodity factor, and π global inflation. The multi-period restriction scheme for horizons $h = 0, \dots, H$ is:

- *Global aggregate demand shock:* $\sum_{s=0}^h \Theta_{IP,1}(s) \geq 0$, $\sum_{s=0}^h \Theta_{CF,1}(s) \geq 0$, $\Theta_{\pi,1}(h) \geq 0$ for all $h \leq H$.
- *Global aggregate supply shock:* $\sum_{s=0}^h \Theta_{IP,2}(s) \geq 0$, $\sum_{s=0}^h \Theta_{CF,2}(s) \geq 0$, $\Theta_{\pi,2}(h) \leq 0$ for all $h \leq H$.
- *Commodity-market shock:* $\sum_{s=0}^h \Theta_{IP,3}(s) \leq 0$, $\sum_{s=0}^h \Theta_{CF,3}(s) \geq 0$, $\Theta_{\pi,3}(h) \geq 0$ for all $h \leq H$.

At $H = 0$ this reduces to the impact-only benchmark.

Table 3 reports the acceptance rates for $H \in \{0, 3, 6, 12\}$ based on 5,000,000 rotation draws (holding the VAR parameters at their ordinary least squares, or OLS, point estimates). The baseline $H = 3$ specification accepts 94,500 draws (1.89%), while the impact-only benchmark accepts 184,433 draws (3.69%). Acceptance rates decline as the restriction horizon tightens beyond the baseline, but the accepted set remains non-empty even at $H = 12$, with 37,751 accepted draws. Figure 8 compares the median impulse responses across restriction horizons. The comparison shows how tightening the restrictions changes the accepted responses. In particular, the inflation response to the aggregate-supply shock remains negative beyond month 12 under $H = 12$, whereas shorter restriction horizons allow a later sign reversal. The informative evidence concerns these responses beyond the restricted window and the change in acceptance rates; signs within each restricted window are imposed by construction.

C2.2 Identification Transparency

Implicit prior. As emphasized by Baumeister & Hamilton (2015), sign-restricted VARs are not prior-free. The sign restrictions define an admissible set of structural models, but posterior summaries within that set, including median impulse responses, credible bands, and historical decompositions, depend on the benchmark measure used to draw admissible rotations. The algorithm of Rubio-Ramírez et al. (2010) uses a uniform measure over orthogonal matrices, which is transparent and standard but does not imply a flat prior over economically meaningful impulse responses.

The impulse responses reported in Figure 4 integrate reduced-form parameter uncertainty by evaluating 300,000 candidate rotations for each of 600 posterior parameter draws. Their posterior medians and 68- and 90-percent credible bands therefore reflect reduced-form parameter uncertainty and variation across admissible rotations, conditional on the estimated DFM factor and maintained DFM specification. By contrast, the restriction-horizon exercise in Appendix C2.1 holds the reduced-form parameters at their OLS estimates in order to isolate the effect of changing the restriction horizon

from the impact-only case to tighter path restrictions over $H = 3, 6, 12$ months. This comparison isolates the sensitivity of the response paths and acceptance rates to the duration of the restrictions.

The credible bands in Figure 4 pool all accepted parameter-rotation draws, so parameter draws with more accepted rotations contribute more observations. The historical-decomposition bands in Appendix D5 instead use ten accepted rotations per retained parameter draw, which gives those draws equal weight.

We interpret the reported impulse responses and historical decompositions as summaries under the benchmark rotation prior, disciplined by the sign restrictions motivated by the theoretical model in Appendix B. The historical decompositions additionally select a single admissible model using the Fry & Pagan (2011) median-target rotation, as described in Appendix C3.

We also separate sign restrictions from impact-magnitude restrictions. The baseline specification does not impose any additional near-zero output-impact bound on commodity-market shocks. Instead, Appendix D1 treats the near-zero band as a robustness stress test. Imposing $-0.10 \leq \Theta_{IP,3}(0) \leq 0$ leaves only a small subset of admissible Bayesian draws, showing that this additional magnitude restriction is highly informative. For this reason, we report it as sensitivity evidence rather than as part of the maintained identification.

C3 Counterfactual Historical Decomposition: Methodology

The historical decompositions in Figure 5 and Appendix Figure 9 follow Kilian & Lewis (2011), using the baseline $H = 3$ FAVAR and the Fry & Pagan (2011) median-target model.

To select that model, we minimize the sum of squared deviations from the pointwise median, dividing each deviation by its standard deviation across the 94,344 accepted OLS rotations. This set comes from an independent simulation of 5,000,000 candidates and therefore differs slightly in size from the count reported in Table 3. The target includes all nine variable–shock responses at horizons $0, \dots, 17$ in the VAR’s original units, before cumulation or commodity-impact scaling. The first minimizer is selected in a tie. Unlike an element-wise median of impact matrices, the selected rotation reproduces the reduced-form covariance matrix exactly and yields mutually uncorrelated shocks.

Combining its structural shocks and impulse responses gives the date- t contribution of shock j to variable k :

$$\hat{y}_{k,t}^{(j)} = \sum_{s=0}^{t-1} \Theta_{k,j}(s) \cdot \hat{e}_{j,t-s} \quad (51)$$

where $\Theta_{k,j}(s)$ is the impulse response of variable k to shock j at horizon s , and $\hat{e}_{j,t}$ is the identified structural shock at date t .

Because industrial production enters the VAR in growth rates and the global commodity factor is estimated from commodity-price growth rates, equation (51) first gives a contribution to the transformed monthly variable. For the cumulated paths in Figure 5 and Appendix Figure 9, we cumulate both the observed transformed series and the shock contribution:

$$Y_{k,t} = \sum_{\tau=1}^t y_{k,\tau}, \quad \hat{Y}_{k,t}^{(j)} = \sum_{\tau=1}^t \hat{y}_{k,\tau}^{(j)}. \quad (52)$$

The counterfactual path “without shock j ” subtracts this contribution from the corresponding observed path:

$$\tilde{Y}_{k,t}^{(\text{no } j)} = Y_{k,t} - \hat{Y}_{k,t}^{(j)}. \quad (53)$$

For variables plotted in non-cumulated units, the same subtraction is applied directly to $y_{k,t}$. In all cases, the exercise removes the realized sequence of one shock type at a time while holding the realized paths of all other shocks, the estimated VAR coefficients, and the initial conditions fixed.

It does not replace a policy rule or treat the identified shocks as policy instruments, and it does not allow for the changes in policy rules or expectation formation that could alter the mapping from shocks to outcomes (McKay & Wolf, 2023). The resulting paths should therefore be interpreted as historical accounting decompositions rather than predictions of how the economy would have evolved had those shocks been absent.

D Robustness and Diagnostic Exercises

This appendix reports robustness and diagnostic exercises around the baseline empirical specification. Appendix D1 examines the on-impact magnitude sensitivity of the commodity-market shock and reports the corresponding impulse responses with 68% and 90% credible bands. Appendix D2 provides external diagnostic comparisons for the structural shocks and examines the relationship between the global commodity factor and oil prices. Appendix D3 re-estimates the FAVAR with Kilian’s index of global real economic activity in place of world industrial production, under the baseline specification. Appendix D4 re-estimates the dynamic factor model (DFM) while relaxing the zero cross-group loading restrictions. Appendix D5 evaluates uncertainty in the historical decompositions across admissible rotations and posterior reduced-form parameter draws.

D1 On-Impact Output Response to the Commodity-Market Shock

The baseline identification restricts the commodity-market shock to have a non-positive contemporaneous effect on industrial production ($\Theta_{\text{IP},3}(0) \leq 0$) and to keep the cumulative industrial production

response non-positive over horizons 0:3. This restriction is consistent with evidence documenting that oil and commodity-price shocks depress economic activity on impact. In the structural VAR literature, Kilian (2009) identifies oil-market-specific demand shocks that reduce U.S. real gross domestic product (GDP); Peersman & Van Robays (2012) document that oil supply disruptions generate an immediate decline in real GDP across net energy-importing countries; and Cashin et al. (2014) find a “long-lived fall in economic activity” in oil-importing economies within a 38-country and region global VAR, with positive effects in some energy exporters. This evidence motivates the non-positive short-run output restriction; it does not by itself establish the monthly, global timing imposed here.

A zero or near-zero impact on economic activity is a natural alternative to the maintained weakly negative output response. To assess this possibility directly, we estimate the Bayesian $H = 3$ specification under an explicit near-zero negative band, $-0.10 \leq \Theta_{IP,3}(0) \leq 0$, while keeping the same path restrictions over horizons 0:3. Industrial-production growth is standardized to unit sample standard deviation, so the bound permits an impact decline of at most 0.10 standard deviations for a unit-variance structural innovation. It is imposed before commodity-factor impact scaling and is not an output-to-commodity-price elasticity. Table 4 compares independent simulations of this specification and the unrestricted Bayesian sign-and-path benchmark.

The unrestricted median impact is -0.552 , with a 5th–95th percentile range of $[-0.829, -0.302]$. The near-zero band yields 445 accepted rotations, compared with 3,277,292 in the unrestricted simulation (about 0.014% as many). The restricted median is -0.087 , with a 5th–95th percentile range of $[-0.099, -0.052]$. Near-zero output impacts are therefore feasible under the $H = 3$ path restrictions with parameter uncertainty, but occupy only a small part of the admissible set.

The 445 rotations come from 74 of the 600 posterior parameter draws. Among those 74 draws, accepted counts range from 1 to 63, with a median of 3; the largest supplies 14.2 percent of the restricted set. Resampling the 600 parameter-draw blocks, including those with no acceptances, 10,000 times gives approximate Monte Carlo standard errors of 0.0003, 0.0012, and 0.0051 for the 5th percentile, median, and 95th percentile of the impact response. These estimates describe precision within the finite simulation and do not rule out unsampled admissible regions.

Figure 10 reports the corresponding responses with 68% and 90% credible bands, conditional on the estimated DFM factor and maintained specifications. Given the sparse accepted set, especially at its upper tail, we interpret these bands as a sensitivity diagnostic. The exercise establishes feasibility under a highly informative magnitude restriction; it does not establish precise tail probabilities or robustness to alternative priors. We retain the weaker non-positive output-impact restriction in the baseline.

D2 External Diagnostic Comparisons

External comparisons assess whether the identified shocks align with related evidence; they do not establish the source of each shock. This section also examines whether the estimated global commodity factor mainly reflects oil-price movements.

D2.1 Global Commodity Factor and Oil Prices

Several of the largest peaks in the global commodity factor coincide with known oil price spikes, raising the question of whether the global factor is primarily capturing energy-driven spillovers rather than broad macroeconomic forces. To assess this, we report the correlation between the estimated global commodity factor and crude oil price growth.

The global factor is positively correlated with nominal crude oil price growth, computed over the dynamic factor model estimation sample, with a correlation of $r = 0.45$ for month-on-month growth rates and $r = 0.60$ for 12-month moving sums. These correlations are expected given that crude oil is included among the commodities in the DFM, but they remain well below unity. This indicates that the global factor captures broad cross-commodity comovement rather than simply serving as an oil price proxy.

D2.2 Structural Shocks vs. External Series

Section IV.2 and Figure 3 present the principal comparison of the identified structural shocks with Baumeister & Hamilton (2019). Additional subperiod correlations are reported below.

The correlations are positive in both subperiods.⁴⁴

Aggregate-supply shock and China's TFP growth. As a supplementary diagnostic of one possible productivity component, we compare the aggregate-supply shock with China's total factor productivity (TFP) growth from the Penn World Table 11.0 (Feenstra et al., 2015). Since country-level TFP is available only annually and is not a direct measure of the aggregate-supply forces represented by the empirical shock, the comparison is necessarily limited. Using annual means of the monthly shocks

⁴⁴Subperiod correlations (1975–2000 / 2001–2023): aggregate demand against the Baumeister & Hamilton (2019) economic activity shock: 0.42 / 0.69; aggregate supply against the Baumeister & Hamilton (2019) economic activity shock: 0.32 / 0.33; commodity-market shock against Baumeister & Hamilton (2019) oil demand plus inventory: 0.22 / 0.34. The correlations are similar or larger in the post-2000 period, indicating that the comparison is not driven by a single subsample. The commodity-market correlations are lower because the comparison pairs a broad commodity-market shock with oil-specific demand and inventory shocks.

over 1973–2023, the correlation is modest ($r = 0.20$), underscoring that the identified shock should not be interpreted as a direct measure of Chinese TFP. During 2002–2008, the aggregate-supply shock is positive in five of seven years while China’s TFP grew at an average rate of 4.8% per year. This correspondence is consistent with productivity growth being one contributor to the super-cycle, but it does not establish productivity as the unique source of the empirically identified shock.

D3 Alternative Measure of Global Activity

As a robustness check, we re-estimate the FAVAR replacing world industrial production with Kilian’s index of global real economic activity.⁴⁵ Because the index is a stationary measure of the global activity cycle, it enters the VAR in levels, and its sign restrictions are imposed pointwise at each horizon, as for inflation; all other elements of the baseline specification are retained, including the 12-lag structure, the path restrictions over horizons 0:3, and the Fry & Pagan (2011) median-target rotation for the historical decomposition. The estimation sample is 1971:M1–2023:M8, the coverage of the index. Of 1,000,000 candidate rotations, 18,197 satisfy the restrictions (1.8 percent, close to the baseline acceptance rate).

Figure 11 reports the resulting full-sample counterfactual historical decomposition. The broad historical interpretation is unchanged. Aggregate-supply shocks remain the dominant contributor to the 2000s super-cycle, with a cumulative contribution of about 28 points at the June 2008 peak (about 25 in the baseline), and the collapse during the global financial crisis again reflects a joint reversal of the aggregate-demand and aggregate-supply contributions. Relative to the baseline, the late-stage contribution of commodity-market shocks during 2006–2008 is smaller and turns negative under this variant.

The two activity measures capture different aspects of the global economy. World industrial production measures quantities produced, whereas the Kilian index uses freight rates to measure activity in industrial commodity markets and may also reflect shipping capacity, inventories, or expectations (Kilian & Zhou, 2018; Hamilton, 2021). These differences may help explain the reallocation between activity and commodity-market shocks, but the exercise does not establish its cause. Aggregate-supply shocks remain the main contributor to the super-cycle under both measures; the finer allocation of the 2006–2008 increase between activity and commodity-market shocks is sensitive to the activity measure.

⁴⁵Kilian’s index is constructed from shipping costs in commodity markets; see Kilian (2009); Kilian & Zhou (2018).

D4 Relaxing the Zero Cross-Group Loading Restrictions

The DFM in Section III.2 restricts each group factor to load only within its own commodity group. These zero restrictions do not rule out cross-group spillovers; they assign broad-based spillovers to the global factor. We assess whether this assignment materially affects the estimated common component by allowing the energy factor to load on all metals and grains prices as well as energy prices, so those ten observation equations add $b_{E,i}f_{E,t}$. The new loadings receive the same $N(0, 1)$ prior as the other loadings. Using the exact baseline sample, 1970:M4–2024:M7, and otherwise identical data transformations, lag structure, factor dynamics, and priors, we estimate four independently initialized Markov chain Monte Carlo (MCMC) chains, each with 2,500 burn-in and 2,500 retained draws, for 10,000 retained draws in total.

The global component is virtually unchanged (Figure 12). The pooled posterior-median global factor from the relaxed model has a monthly correlation of 0.997 with the baseline factor and a correlation of 0.999 in 12-month moving sums (Panel A). Across the four chains, the corresponding correlations range from 0.979 to 0.996 and from 0.990 to 0.998. The time-series standard deviation of the pooled posterior median is 98.7 percent of the baseline value. Commodity-level global variance shares have a correlation of 0.993 across specifications and change by only 0.83 percentage points on average (Panel B). Average global shares rise from 23.9 to 24.9 percent for metals and from 10.1 to 11.6 percent for grains.

Because the structural FAVAR uses the global factor, this near identity leaves its empirical input effectively unchanged. We do not re-estimate the FAVAR on the relaxed factor, so the result here concerns the global factor and its variance contribution rather than the structural estimates. Relaxing them introduces rotational flexibility in how residual within-group variation is allocated among the group factors, so we do not interpret the individual added loadings as separately identified transmission parameters. The relevant result is that the estimated global commodity cycle and its variance contribution are stable when the restrictions are relaxed. Split $\hat{R}$ statistics for the global factor have a median of 1.03 across dates and a maximum of 1.51, so mixing is incomplete at some dates: the high chain correlations support the stability of the estimated factor path but not accurate pooled posterior inference at every date.

D5 Historical-Decomposition Uncertainty

Figure 5 uses the Fry & Pagan (2011) median-target model. We assess uncertainty across all 94,344 accepted rotations at the baseline $H = 3$ ordinary least squares estimates, from 5,000,000 candidates drawn independently of Table 3 (94,500 accepted). We then use 600 posterior parameter draws and

10,000 candidate rotations per draw, computing contributions with each model's own coefficients, covariance matrix, intercept, and residuals.

The posterior exercise accepts 110,342 parameter-rotation models (1.84 percent); 98.8 percent have stable companion matrices. Table 5 conditions on stability; retaining unstable models barely changes the episode probabilities. Uncertainty covers reduced-form FAVAR parameters and admissible rotations, conditional on the estimated DFM factor, maintained DFM specification, and identifying restrictions.

Figure 13 reports 68-percent and 90-percent credible bands for within-episode contribution changes. Using ten accepted rotations per posterior draw gives each parameter draw equal weight. Table 5's probabilities are nearly identical under equal-draw weighting and pooling all accepted rotations.

The principal historical conclusions are stable. From January 2003 to June 2008, aggregate-supply shocks are the largest contributor in 83.1 percent of posterior models and have a median contribution of 23.3 cumulated-factor points, with a 68-percent credible interval of [14.8, 31.0]. The median-target contribution for the same period is 25.7 points, inside that interval. During the June 2008–February 2009 collapse, aggregate-demand and aggregate-supply shocks are the largest negative contributor in 48.0 and 50.8 percent of models, respectively; commodity-market shocks are largest in only 1.2 percent. From April 2020 to March 2022, aggregate-demand shocks are the largest contributor in 87.9 percent of models, with a median contribution of 16.2 points and a 68-percent interval of [9.8, 22.4].

In the early 1970s, aggregate-supply shocks are most often largest near January 1974; commodity-market shocks are largest by January 1975 in 62.0 percent of posterior models.

E Additional Figures and Tables



FIGURE 6: GROUP COMMODITY PRICE FACTORS: 1970:M4–2024:M7

Note: Each panel shows the contribution of the global commodity factor and the group-specific factor to a representative commodity in each group: crude oil (Energy), gold (Precious Metals), aluminum (Metals), and maize (Grains). All series are 12-month moving sums. Shaded vertical intervals mark periods of oil price spikes as identified in Hamilton (2009) (amber shading) and U.S. recession periods dated by the National Bureau of Economic Research (NBER, blue-grey shading).




FIGURE 7: Cumulative Structural Shock Series, 1972–2023

Note: Cumulative structural shocks $\sum_{s=1}^t \hat{e}_{i,s}$ recovered using the impact matrix of the Fry-Pagan median-target rotation among the accepted sign-restriction draws. Shocks are normalized so that a positive value corresponds to an increase in the global commodity factor. Vertical dashed lines indicate key macroeconomic events.




FIGURE 8: Median Impulse Responses Under Multi-Period Sign Restrictions

Note: Median impulse responses for restriction horizons $H = 0$ (impact-only), $H = 3$ (baseline), $H = 6$, and $H = 12$. The darker line corresponds to the baseline $H = 3$ path restriction; the lighter lines report the alternative horizons. All shocks are normalized to increase the global commodity factor. Industrial production and global commodity factor responses are cumulated. Based on 5,000,000 rotation draws per horizon with cumulative restrictions on industrial production and the global commodity factor, and level restrictions on inflation.




FIGURE 9: HISTORICAL DECOMPOSITION: GLOBAL COMMODITY FACTOR (FULL SAMPLE)

Note: Solid line: cumulated global commodity factor. Dashed line: counterfactual path showing the evolution of the factor without each structural shock. Shaded area: cumulative contribution of the indicated shock. Counterfactuals are constructed following Kilian & Lewis (2011); see Appendix C3 for details.




FIGURE 10: Bayesian Impulse Response Functions with 68% and 90% Credible Bands (Explicit Near-Zero Negative Impact Bound)

Note: Bayesian impulse responses under the explicit near-zero negative bound $-0.10 \leq \Theta_{IP,3}(0) \leq 0$ for the commodity-market shock's contemporaneous effect on industrial production, while maintaining the $H = 3$ path restrictions. The outer band is the 90% Bayesian credible band (5th–95th percentile). The inner shaded region is the 68% Bayesian credible band (16th–84th percentile). The solid line is the Bayesian median response. The 445 rotations represent 74 posterior parameter draws; the sparse set limits the precision of the bands. Responses condition on the estimated DFM factor and maintained specifications. All shocks are signed to increase the global commodity factor. Industrial production and global commodity factor responses are cumulated.




FIGURE 11: COUNTERFACTUAL HISTORICAL DECOMPOSITION: GLOBAL COMMODITY FACTOR WITH KILIAN INDEX

Note: Solid line: cumulated global commodity factor under the Kilian-index specification, which retains the baseline 12-lag structure and $H = 3$ path restrictions, with the index entering in levels and the historical decomposition based on the Fry-Pagan median-target rotation (Appendix D3). Dashed line: counterfactual path without the indicated structural shock. Shaded area: cumulative contribution of the shock. Sample: 1972:M1–2023:M8 after lags.




FIGURE 12: RELAXING THE ZERO CROSS-GROUP LOADING RESTRICTIONS: ENERGY FACTOR LOADS ON METALS AND GRAINS

Note: Results from re-estimating the DFM over the exact baseline sample, 1970:M4–2024:M7, with the energy group factor allowed to load on all metals and grains prices (Appendix D4). Panel A: posterior-median global commodity factor from the baseline model (solid) and the relaxed-loadings model (dashed), shown as 12-month moving sums. The correlation is 0.997 for the monthly factors and 0.999 for the plotted 12-month sums. Across the four independently initialized MCMC chains, monthly correlations with the baseline range from 0.979 to 0.996. Panel B: commodity-level global-factor variance shares under the baseline and relaxed specifications; the diagonal denotes equality. The across-commodity correlation is 0.993, and the mean absolute change is 0.83 percentage points. Metals are shown in blue and grains in orange.




FIGURE 13: HISTORICAL-DECOMPOSITION UNCERTAINTY: EPISODE-REBASED CONTRIBUTIONS

Note: Each row rebases the contribution of every shock to zero at the beginning of the indicated episode. The solid colored line is the posterior median, the darker region is the 68% credible band, and the lighter region is the 90% credible band. The black line is the contribution from the Fry–Pagan median-target rotation used in Figure 5. Posterior bands use a balanced sample of ten accepted rotations from each reduced-form parameter draw with a stable companion matrix, under the baseline $H = 3$ sign-and-path restrictions. Of the 600 parameter draws, 593 are stable, giving 5,930 paths. Results are conditional on the estimated DFM factor.

TABLE 2: BASELINE CALIBRATION

Parameter	Description	Value	Source or Rationale
<i>Household</i>			
β	Discount factor	0.99	Standard (quarterly)
σ	Risk aversion	1.5	Standard
φ	Inverse Frisch elasticity	1.0	Chetty, Guren, Manoli, & Weber (2011); macro Frisch elasticity near 1
<i>Production</i>			
α	Commodity cost share	0.10	Energy plus materials share
μ	Labor-commodity substitution	1.0	Cobb-Douglas limit (numerically set to 1.001)
ε	Substitution between varieties	6.0	Steady-state markup of 20 percent
θ	Calvo parameter	0.60	Nakamura & Steinsson (2008); about 2.5 quarters
<i>Commodity markets</i>			
ζ	Commodity supply elasticity	0.20	Low short-run elasticity
η	Across-group substitution	0.30	Low cross-group substitutability
η_g	Within-group substitution	1.50	Disaggregated interpretation only
$\omega_E, \omega_L, \omega_F$	Commodity composite shares	0.5, 0.3, 0.2	Energy, metals, food weights
η_S	Desired-inventory price elasticity	0.15	Price-inelastic short-run inventory demand
ρ_S	Inventory stock persistence	0.85	Gradual adjustment of inventories
δ_S	Inventory depreciation	0.10	Replacement demand for stored commodities
s_S	Inventory stock share	0.80	Inventory stock relative to quarterly input use
<i>Monetary policy</i>			
ϕ_R	Interest rate smoothing	0.50	Moderate inertia
ϕ_π	Inflation response	1.50	Taylor principle
ϕ_y	Output-gap response	0.125	Standard Taylor-rule coefficient
<i>Shock persistence</i>			
ρ_A	TFP persistence	0.95	Highly persistent
ρ_ξ	Demand-shock persistence	0.70	Moderately persistent
ρ_v	Inventory-demand persistence	0.97	Persistent precautionary and storage demand
ρ_z	Group supply-shock persistence	0.70	Moderately persistent
$\rho_{v,g}$	Group inventory-demand persistence	0.50	Less persistent than the global motive

Note: The model is calibrated at quarterly frequency. The commodity share is set to 0.10 so that commodity-market disturbances primarily move commodity prices rather than aggregate output. The inventory parameters imply gradual stock adjustment and positive replacement demand, so inventory-demand shocks affect prices through physical stockbuilding rather than a direct shift in production demand. μ is implemented in Dynare as 1.001 to avoid the CES singularity at exactly one; the results are numerically indistinguishable from the Cobb–Douglas limit. The within-group elasticity η_g appears only in the disaggregated interpretation of the factor structure (Appendix B1.3) and does not enter the aggregate impulse responses, which treat each commodity group as a single composite.

TABLE 3: Multi-Period Sign Restrictions: Acceptance Rates

Restriction Horizon (H)	Accepted Draws	Acceptance Rate (%)	Total Draws
$H = 0$ (impact-only)	184,433	3.69	5,000,000
$H = 3$ (baseline)	94,500	1.89	5,000,000
$H = 6$	64,807	1.30	5,000,000
$H = 12$	37,751	0.76	5,000,000

Note: Number of rotation draws satisfying sign restrictions at all horizons $h = 0, \dots, H$. Restrictions on industrial production and global commodity factor responses are imposed on cumulative impulse responses; restrictions on inflation are imposed on the pointwise response. Total draws: 5,000,000 per horizon.

TABLE 4: Bayesian Restriction on Near-Zero Negative Output Impact Under $H = 3$

Specification	Accepted Draws	Acceptance Rate (%)	Median On-Impact Output	5th–95th On-Impact Output
Bayesian baseline: $H = 3$ path restrictions	3,277,292	1.821	−0.552	[−0.829, −0.302]
Bayesian restricted: $H = 3$ path restrictions plus $-0.10 \leq \Theta_{IP,3}(0) \leq 0$	445	0.00025	−0.087	[−0.099, −0.052]

Note: Each row uses an independent simulation with 600 posterior draws and 300,000 rotations per posterior draw. Output is measured in standard deviations of monthly industrial-production growth per unit-variance structural innovation. The first row imposes the maintained $H = 3$ path restrictions without any additional impact-magnitude bound on the commodity-market shock. The second row conditions the Bayesian posterior on the explicit near-zero negative band $-0.10 \leq \Theta_{IP,3}(0) \leq 0$. The restricted set is non-empty under parameter uncertainty, but its size is 0.0136% of the unrestricted accepted set. The 445 restricted rotations represent 74 parameter draws (1–63 acceptances each; median 3).

TABLE 5: HISTORICAL-DECOMPOSITION UNCERTAINTY: KEY EPISODE CHANGES

Shock	Jan. 2003–Jun. 2008 increase		Jun. 2008–Feb. 2009 decline		Apr. 2020–Mar. 2022 increase	
	Median [68% interval]	Largest (%)	Median [68% interval]	Largest decline (%)	Median [68% interval]	Largest (%)
Aggregate demand	3.5 [−1.1, 9.3]	3.5	−9.6 [−17.2, −3.6]	48.0	16.2 [9.8, 22.4]	87.9
Aggregate supply	23.3 [14.8, 31.0]	83.1	−10.1 [−15.9, −2.4]	50.8	0.4 [−4.2, 6.4]	6.4
Commodity market	7.5 [2.2, 15.3]	13.3	−0.3 [−4.4, 2.1]	1.2	2.5 [−1.3, 7.1]	5.8

Note: Contributions are changes in points on the cumulated global-factor scale. Intervals are the 16th–84th percentiles across accepted posterior parameter-rotation models with stable companion matrices. “Largest” reports the percentage of models in which the indicated shock makes the largest positive contribution to an increase; “Largest decline” reports the percentage in which it makes the largest negative contribution. The exercise uses 600 posterior reduced-form parameter draws and 10,000 candidate rotations per draw under the baseline $H = 3$ restrictions, yielding 110,342 accepted models, of which 109,047 are stable. Results condition on the estimated DFM factor.