

BEYOND OIL: THE ORIGINS OF COMMODITY PRICE FLUCTUATIONS

ALVIN LUMBANRAJA, SARAH MOUABBI, EVGENIA PASSARI AND ADRIEN ROUSSET PLANAT

ABSTRACT. Commodity supply shocks are a plausible but empirically elusive source of business-cycle fluctuations. We develop a comprehensive framework to measure them, constructing daily supply and demand proxies for 20 commodities—spanning energy, metals, agriculture, and livestock—from textual analysis of over one million news articles (2001–2023). These measures allow us to separate supply from demand across the full commodity market, not just oil. A striking finding emerges: non-oil supply disruptions affect inflation and industrial production at least as strongly as oil disturbances, a result previously undocumented in the literature. Transmission varies sharply with countries’ commodity trade positions: net importers experience more persistent output contractions and stronger inflation pass-through, while net exporters are partially insulated.

JEL classification codes: C19, F30, G15, Q02.

Keywords: commodities, textual analysis, news, demand, supply, business cycles, geopolitical risk, natural disasters, climate change.

1. INTRODUCTION

Ever since the fundamental work of [Kydland and Prescott \(1982\)](#), macroeconomists have searched for empirical measures of supply-side disturbances that drive business-cycle fluctuations. Commodities—as essential inputs across most production sectors—are a natural place to look: shocks to their supply propagate through input-output linkages and can generate economy-wide fluctuations in quantities

November 2025. A. Lumbanraja: Kellogg School of Management, Northwestern University, 2211 Campus Drive Evanston, IL 60208 (e-mail: alvin.lumbanraja@kellogg.northwestern.edu) . S. Mouabbi: Banque de France, 31 rue Croix des Petits Champs, 75049 Paris, France (e-mail: sarah.mouabbi@banque-france.fr). E. Passari: Université Paris-Dauphine, PSL and CEPR. DRM Finance, Université Paris Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris, France (e-mail: evgenia.passari@dauphine.psl.eu). A. Rousset Planat: Economics Department, London Business School, Regent’s Park, NW1 4SA London, UK (e-mail: adrienr@london.edu). The authors thank Amedeo Andriollo, Marina Azzimonti, Daniele Bianchi, Deepa Dhume Datta, Simona Delle Chiaie, Laurent Ferrara, Tarek Hassan, Matteo Iacoviello, Şebnem Kalemli-Özcan, Loukas Karabarbounis, Andreas Konstantinidis, Delphine Lautier, Michele Lenza, Gernot Müller, Anton Nakov, Luca Onorante, Elias Papaioannou, Dimitris Papanikolaou, Gert Peersman, Ricardo Reis, Hélène Rey, Philip Sauré, Martin Schneider, Daniele Siena, Vania Stavrakeva and seminar and conference participants at numerous venues for useful comments. Evgenia Passari gratefully acknowledges financial support from the Bank of Greece, the Europlace Institute of Finance (EIF) & the Labex Louis Bachelier and the FinTech & Digital Finance Chair at Université Paris Dauphine - PSL, as well as hospitality from the Kellogg School of Management, Northwestern University, the Federal Reserve Bank of Chicago and the Kiel Institute of the World Economy where this research was partly conducted. The views expressed in this paper are those of the authors and should under no circumstances be interpreted as reflecting those of the Banque de France or the Eurosystem.

and prices. Yet despite their theoretical importance, empirical progress in measuring commodity supply shocks has been largely confined to oil (Hamilton, 1983, 2003).¹ Modern economies, however, depend on a far broader set of commodity inputs—industrial metals, agricultural products, and diverse energy carriers—whose supply disruptions are varied and potentially consequential. The COVID-19 pandemic, Russia’s invasion of Ukraine, and ongoing geopolitical tensions have made this gap in the literature increasingly difficult to ignore.

We fill this gap by developing a comprehensive framework to measure supply and demand shocks across the full commodity market. Using a novel automated narrative approach applied to more than 1.1 million news articles, we build daily proxies that span 20 commodities—including 19 previously undocumented markets—covering energy, metals, agriculture, and livestock. The daily frequency allows us to control for autocorrelation and cross-market spillovers, pushing toward causal identification in a pseudo-high-frequency setup. Because the framework covers the entire commodity spectrum, it enables a question the existing literature cannot address: how do supply shocks originating in non-oil commodity markets compare to oil shocks in their macroeconomic impact?

The answer is striking: non-oil supply disturbances—originating in metals, grains, softs, and livestock markets—affect inflation and industrial production at least as strongly as oil disturbances. This result, previously undocumented, suggests that the literature’s narrow focus on petroleum has understated the macroeconomic importance of commodity supply shocks as a class.

A second result concerns the heterogeneous transmission of these shocks across countries. Interacting our supply disturbances with countries’ normalized trade balance—in a specification with both country and time fixed effects—we show that net importers suffer more persistent output contractions and stronger inflation pass-through, while net exporters are partially insulated. The contrasting patterns of industrial production and CPI responses reveal that international adjustment occurs primarily through quantities rather than prices.

The key identification challenge is separating supply-driven from demand-driven price fluctuations. We address it in three steps. First, we construct audited dictionaries that capture the words journalists use to describe commodity market forces, count their co-occurrence with directional modifiers, and normalize by total used article volume—yielding comparable daily measures of supply and demand pressure. Second, we decompose these pressures into underlying drivers (business cycles, geopolitical risk, natural disasters, climate change) by building dictionaries from authoritative reference texts. Third, we conduct an extensive human audit of randomly sampled articles; correlations between human-coded and machine-generated indices reach 85% in quarterly data.

We validate the methodology along three dimensions. First, supply contractions and demand expansions are both associated with higher commodity spot prices—not only at the index level but for

¹Additional work on oil and the business cycle includes Kilian (2009); Antolín-Díaz and Rubio-Ramírez (2018); Baumeister and Hamilton (2019); Caldara, Cavallo, and Iacoviello (2019) and Känzig (2021).

individual commodities—and this relationship is driven by the unanticipated component of our indicators. Second, for oil, the one market where structural benchmarks exist, our supply measure produces impulse responses closely aligned with those of [Baumeister and Hamilton \(2019\)](#) and [Känzig \(2021\)](#). Third, the cross-country interaction of our supply disturbances with trade-balance positions, estimated with country and time fixed effects, functions as a difference-in-differences test that pins the effects to commodity exposure.

The indices capture major market events: demand contracts sharply during the 2008 Great Recession and the COVID-19 crisis, while supply identifies natural disasters and logistical bottlenecks. At the commodity level, indicators track market-specific dynamics—oil supply responds to OPEC decisions, while wheat supply is driven by weather—and the theoretical distinction between supply and demand shocks emerges clearly: supply contractions raise prices while reducing output, whereas demand expansions increase both.

A further contribution is the decomposition of commodity supply and demand into their underlying drivers. Building on taxonomies from the IMF's *Commodity Special Feature* and the World Bank's *Commodity Markets Outlook*, we classify the forces behind commodity fluctuations into four categories: business-cycle developments, geopolitical risk, natural disasters, and climate change. Natural disasters emerge as the dominant driver of supply disruptions—concentrated in grains, softs, and energy—while business-cycle dynamics primarily shape demand, which co-moves strongly across commodity groups. The cross-category patterns are systematic: business-cycle effects are strongest in energy and industrial metals; geopolitical risks disproportionately affect energy and precious metals; natural disasters primarily disrupt agriculture; and climate-change concerns are most pronounced in energy and increasingly in industrial metals. Beyond supply, our methodology also extracts demand signals, enabling tests of New Keynesian predictions about demand propagation ([Christiano, Eichenbaum, and Evans, 2005](#); [Galí, 2015](#); [Kaplan, Moll, and Violante, 2018](#)).²

Despite long-standing theoretical interest in supply-side propagation—from the multi-sector models of [Long and Plosser \(1983\)](#) and the production network models of [Acemoglu et al. \(2012\)](#) to the multi-sector DSGE frameworks of [Horvath \(2000\)](#), [Atalay \(2017\)](#), and [Baqaee and Farhi \(2022\)](#)—direct empirical evidence has remained limited. The recent experience of COVID-19 has revived interest in supply-side explanations: [Bernanke and Blanchard \(2025\)](#) emphasize that supply chain disruptions were a key driver of pandemic-era inflation. But was COVID an exception or part of a broader pattern? Our analysis shows that a wide range of commodity supply disruptions—including those stemming from geopolitical conflicts, natural disasters, and climate change—systematically shape macroeconomic outcomes, providing direct empirical evidence for these theoretical mechanisms.

Our work connects to three strands of the literature. First, it builds on textual analysis of economic outcomes ([Tetlock, 2007](#); [Gentzkow and Shapiro, 2010](#); [Hoberg and Phillips, 2010](#); [Boudoukh et al.,](#)

²See also [Nakamura and Steinsson \(2014\)](#) on the importance of identifying demand shocks.

2013; Alexopoulos and Cohen, 2015; Baker, Bloom, and Davis, 2016; Allcott and Gentzkow, 2017; Hasan et al., 2019), bringing these methods to the commodity space where the breadth of price drivers makes narrative approaches especially powerful. Second, it complements the literature on supply and demand identification in energy markets (Kilian, 2008, 2009; Kilian, Rebucci, and Spatafora, 2009; Kilian and Vigfusson, 2017; Känzig, 2021), extending it beyond energy to the full commodity spectrum. Third, it relates to narrative shock identification in the spirit of Romer and Romer (2010), including Wu and Cavallo (2012), Caldara, Cavallo, and Iacoviello (2019), Datta and Dias (2019), and Loughran, McDonald, and Pragidis (2019). To our knowledge, this is the first automated narrative approach to build a real-time global commodity index spanning energy, metals, agriculture, and livestock.

The fundamental identification challenge is that commodity prices respond endogenously to macroeconomic conditions. In oil markets, SVAR frameworks address this using exclusion restrictions (Kilian, 2009), sign restrictions (Kilian and Murphy, 2012; Lippi and Nobili, 2012; Baumeister and Peersman, 2013; Baumeister and Hamilton, 2019), narrative information (Antolín-Díaz and Rubio-Ramírez, 2018; Caldara, Cavallo, and Iacoviello, 2019; Zhou, 2019), and high-frequency variation around OPEC announcements (Känzig, 2021). Extending these approaches beyond oil is difficult due to data limitations and heterogeneous production structures; only a handful of papers derive structural supply measures outside energy (Winne and Peersman, 2016; Castelnovo, Mori, and Peersman, 2024; Adjemian and Jo, 2025). Rather than estimating fully identified structural shocks across numerous markets, our narrative approach develops consistent, comparable proxies that span the commodity universe.

The remainder of the paper is organized as follows. Section 2 introduces our methodology for developing supply- and demand-side indicators in commodity markets. In Section 3 we perform a series of validation exercises for our indicators. Section 4 compares the macroeconomic responses to oil and non-oil disturbances and documents their comparable impact on inflation and output, as well as the heterogeneous transmission across countries' commodity trade positions. Section 5 decomposes the indices into factors driving commodity price fluctuations and documents the time-varying importance of these factors. Section 6 concludes.

2. BUILDING SUPPLY AND DEMAND COMMODITY MEASURES

In this section, we introduce our news-based narrative strategy for identifying supply- and demand-side fluctuations in commodity prices. The approach is flexible and applies across all major commodity groups, including energy, metals, agriculture, and livestock. Our indices clearly differentiate between supply and demand drivers, and their peaks coincide with major episodes such as the post-crisis collapse in commodity markets, market-specific disruptions, and the COVID-19 shock. The richness of the news data also allows us to pinpoint the fundamental forces shaping commodity dynamics, including business-cycle conditions, geopolitical risks, natural disasters, and climate change.

2.1. General Methodological Overview. Our methodological approach rests on the premise that financial news provides timely and economically meaningful information about the shocks that drive commodity markets. Two assumptions guide our approach: news outlets report economically significant developments with minimal delay, and the magnitude of a shock is positively associated with the breadth and intensity of coverage. Building on these assumptions, we construct text-based indices through a two-stage procedure. In the first stage, we identify words and phrases that signal supply- and demand-side developments at the sentence level, using audited dictionaries combined with standard directional terminology (see Appendix C.4). In the second stage, we refine these signals with algorithms that remove ambiguous contexts, adjust for negations, and incorporate commodity-specific exceptions. The net supply indicator equals supply contractions minus expansions, and the net demand indicator equals demand expansions minus contractions. These indicators can be aggregated at any time frequency and are comparable across all major commodity categories—energy, industrial and precious metals, agricultural commodities, and livestock. Additional details on the text-processing pipeline and contextual refinement algorithms are provided in Appendix C.

2.2. News Outlets and Article Selection. Our empirical analysis relies on more than 1,100,000 Reuters news articles published between May 2000 and May 2023, sourced through the Factiva database. We apply a multi-stage filtering and deduplication protocol—including the removal of near-duplicate articles, verification of textual fidelity, proper-noun disambiguation, and systematic negation handling—to ensure that the corpus includes only relevant and non-redundant content (see Appendix C for full details). When constructing commodity-specific indicators, we apply a relevance filter that excludes sentences whose primary focus is on substitute commodities, thereby reducing measurement error. In contrast, the aggregate commodity index follows an inclusive approach, incorporating all articles containing commodity-related terminology.

2.3. Disentangling Supply and Demand. We construct supply and demand dictionaries by identifying high-frequency terms in the corpus and classifying them according to their economic meaning, with human validation for accuracy. The core supply vocabulary includes terms such as *supply*, *production*, and *output*; the core demand vocabulary includes *demand*, *consumption*, *buy*, and *purchase*. For individual commodities, we augment these core terms with market-specific vocabulary (see Appendix C.4). To incorporate directional information, we combine these vocabularies with standard dictionaries of increase- and decrease-related terminology. Our methodology counts instances in which supply- or demand-related words appear alongside directional indicators within a sentence-level proximity window, yielding continuous measures of directional market pressure.

To ensure accurate classification, we supplement the baseline approach with contextual refinement algorithms that handle semantic negation and commodity-specific linguistic idiosyncrasies—for example, “high temperature” in agriculture signals supply disruption rather than expansion, and “long

maintenance” in energy markets indicates reduced output. These corrections, along with the full text pre-processing pipeline and dictionary construction details, are described in Appendix C.

A natural question is whether more sophisticated natural language processing methods—such as transformer-based models (e.g., BERT or GPT)—could improve upon our dictionary approach. We deliberately choose a parsimonious, rule-based methodology for several reasons. First, transparency and auditability are essential for credibility in applied macroeconomic research: every classification decision in our framework can be traced to specific dictionary entries and contextual rules, making it straightforward for researchers to inspect, replicate, and challenge our results. Second, our dictionaries are informed by domain-specific economic knowledge rather than statistical patterns learned from general-purpose training corpora, which may not reliably distinguish economically meaningful supply and demand signals in commodity news. Third, the simplicity of our approach minimizes the risk of overfitting to particular subsamples or time periods, a concern that is especially relevant given the evolving nature of commodity market reporting. Finally, our human audit exercise (Section 3.2) demonstrates that the dictionary-based classification achieves high accuracy, with quarterly correlations between human- and machine-generated indices reaching 85%, suggesting limited scope for improvement from more complex methods.

2.4. Narrative Drivers of Supply and Demand. In addition, we construct a separate classification of the key events that drive commodity markets. Drawing on the IMF’s *Commodity Special Feature* and the World Bank’s *Commodity Markets Outlook*, we organize the potential sources of commodity-market shocks into four primary drivers: business-cycle developments, geopolitical risk, natural disasters, and climate change. To identify language associated with each driver, we develop dictionaries derived from authoritative sources, following the methodology of [Engle et al. \(2020\)](#) for contexts in which pre-defined dictionaries are unavailable. For business-cycle terminology, we use the complete historical record of announcements issued by the NBER Business Cycle Dating Committee (1979-2020). To capture geopolitical risk, we rely on the full set of press statements from the President of the United Nations Security Council. For natural disasters, we draw on the [EM-DAT International Disaster Database](#). For climate-change terminology, we adopt the same authoritative texts used by [Engle et al. \(2020\)](#).

2.5. Supply and Demand Indicators. We next create indicators of supply-side and demand-side disturbances. We construct both aggregate indicators but also indicators for specific commodities. For the aggregate commodity market indicator, we employ an inclusive sampling strategy that captures all articles referencing any commodity included in the Standard & Poor’s Goldman Sachs Commodity Index (GSCI) spot price series excluding lead. This criterion ensures broad coverage of economically significant commodity markets, as the GSCI comprises major physical commodities with active

and liquid futures markets, providing a representative measure of broad commodity market developments. Table A1 lists the individual commodity constituents considered in our analysis, which encompasses the full range of major tradable commodities across energy, metals, and agricultural markets.

For each commodity i , we begin by constructing distinct measures of supply- and demand-driven directional market pressure. Specifically, we create four indicators (capturing increases and decreases in supply and demand) by counting, for each day t , the number of instances that mention supply or demand in market i together with a directional qualifier (+/-). Accordingly, the indicator for a supply increase is defined as:

$$S_{i,t}^+ \equiv \frac{\text{\# of Instances Indicating Commodity Supply Increase}_{i,t}}{\text{\# of Articles Read}_t}. \quad (1)$$

The remaining indicators—supply decreases $S_{i,t}^-$ and demand increases $D_{i,t}^+$ and decreases $D_{i,t}^-$ —are defined analogously. To account for shifts in publication volume, all indicators are normalized by the number of retained articles on day t . We also develop aggregate indices that pool information across markets. Thus, the aggregate supply-increase indicator is given by:

$$S_t^+ \equiv \frac{\text{\# of Instances Indicating Commodity Supply Increase}_t}{\text{\# of Articles Read}_t}. \quad (2)$$

Next, we derive net supply and net demand measures by taking the difference between opposing directional indicators: net supply as supply contractions minus supply expansions, and net demand as demand expansions minus demand contractions. Thus, for the aggregate indices, we define

$$S_t \equiv S_t^- - S_t^+ \quad \text{and} \quad D_t \equiv D_t^+ - D_t^-. \quad (3)$$

We construct analogous indices for each commodity market i .

Figure 1 displays the standardized monthly net supply and net demand indicators. To mitigate distortions arising from changes in publication volume over time, these indicators are normalized by the daily count of retained articles. The resulting time series align closely with well-documented commodity market dynamics. Major turning points in the demand indicator coincide with systemic economic events, including the 2001 US Recession, the global financial crisis (2007-2009), the COVID-19 pandemic (2020-2021). The supply-side indicator exhibits pronounced volatility during episodes of significant production adjustments, such as OPEC interventions following the 2001 recession and the September 2001 terrorist attacks, as well as during subsequent quota revisions. The supply measure additionally shows marked contractions during severe drought episodes that significantly impaired global grain production.

Figure 2 presents the decomposition of the net supply and net demand indicators into their underlying drivers across major commodity categories. The cross-sectional patterns highlight clear differences in both the incidence and synchronicity of market developments. Supply disturbances arise primarily in grains, softs, and energy markets, and exhibit limited temporal alignment across categories, indicating that supply-side pressures tend to be commodity-specific rather than systemic. In contrast, demand fluctuations show a high degree of cross-sectional co-movement, with pronounced cyclical patterns evident across most commodity groups. This synchronicity aligns with the influence of shared macroeconomic forces shaping global demand conditions. Grain markets represent a notable exception: they demonstrate comparatively muted demand elasticity outside episodes of trade conflict, reflecting the inherently inelastic consumption characteristics of staple agricultural goods.

Overall, the commodity-level indicators for crude oil, wheat, corn, sugar, and copper reveal marked heterogeneity in the drivers of price dynamics. Energy and metal markets are highly responsive to the business cycle, with demand movements closely aligned with broader macroeconomic conditions. Agricultural commodities, by contrast, are more heavily affected by natural-disaster shocks and show more inelastic demand due to their specific production processes and consumption profiles.

2.5.1. Energy Commodities: Crude Oil. Our crude oil indicators capture well-documented market dynamics, particularly production adjustments following OPEC ministerial meetings. OPEC's supply management decisions remain the dominant source of supply-side variation in global oil markets. Additional supply disruptions arise from weather-related production impairments (e.g., hurricanes), the structural expansion of US production capacity during 2014-2016, and the acute oversupply conditions of April 2020, when pandemic-induced demand collapses triggered both a price breakdown and severe physical storage constraints. The demand for petroleum products displays pronounced cyclicity, with major contractions aligning with key macroeconomic downturns: the 2001 US recession, the 2007-2009 global financial crisis, and the early phase of the COVID-19 pandemic. Figure 3 further illustrates the rapid demand recovery observed in 2021, after the global roll-out of vaccines, and in early 2023, after China's economic reopening. These episodes underscore the sensitivity of petroleum consumption to mobility restrictions and broader economic activity.

2.5.2. Agricultural Commodities: Wheat, Corn and Sugar. Appendix Figure A1 presents the standardized net supply and demand indicators for wheat markets from 2001 through 2023. The supply indicator identifies major production disruptions coincident with extreme meteorological events affecting global wheat cultivation. Adverse growing conditions—including drought episodes, extreme temperature events, and other detrimental weather patterns—register as significant negative supply developments in our index. Demand indicators for wheat exhibit distinct patterns relative to energy commodities, with pronounced positive demand disturbances corresponding to bilateral trade

agreements between major market participants (notably US-China agricultural procurement commitments) and precautionary consumption behaviors, particularly evident in the stockpiling observed during the initial pandemic phase. This pattern highlights an important cross-commodity distinction: while energy demand indicators primarily reflect cyclical variations in economic activity, agricultural commodity demand indicators are more responsive to precautionary consumption motives and trade policy developments.

Appendix Figure A2 presents the standardized net supply and demand indicators for corn markets throughout the sample period. Corn demand exhibits distinctively different patterns compared to wheat, displaying greater elasticity and industrial utilization sensitivity. Pronounced demand increases frequently coincide with biofuel sector expansion, reflecting corn's significant role as an ethanol feedstock. This industrial demand component introduces cyclical sensitivity absent in more consumption-oriented agricultural commodities. The COVID-19 pandemic demonstrates this relationship, as corn demand contracted significantly following the collapse in both ethanol production and livestock markets—key downstream sectors dependent on corn inputs. Additional demand fluctuations stem from trade policy adjustments and tariff implementations that affect international market access. Supply-side developments in corn markets, similar to wheat, are predominantly influenced by meteorological conditions affecting cultivation outcomes. Differences in supply disruption timing typically reflect the geographical distribution of production regions rather than fundamentally different causal mechanisms.

Appendix Figure A3 displays the standardized net supply and demand indicators for sugar markets from 2001 through 2023. Sugar demand exhibits distinctive temporal patterns, with pronounced seasonality driven by international holidays and festivals that generate systematic consumption peaks. Additionally, discrete demand surges coincide with geopolitical tensions as market participants engage in precautionary inventory accumulation during periods of elevated uncertainty. Supply-side fluctuations primarily reflect weather-induced production variability and policy-driven adjustments to production quotas, particularly changes in regulatory constraints on sugar beet cultivation over the sample period.

2.5.3. *Industrial Metals: Copper.* Appendix Figure A4 presents the standardized net supply and demand indicators for copper markets spanning 2001 through 2023. Supply-side developments predominantly reflect production conditions in the major copper-producing nations—Chile, Peru, China, Russia, and Indonesia. Negative supply shocks correspond to several distinct categories of production disruptions: natural disasters (predominantly earthquakes) affecting mining operations, technical disruptions to extraction and refining processes, labor disputes manifesting as production stoppages, and international sanctions restricting market access. Demand for copper exhibits pronounced cyclicity, with significant sensitivity to macroeconomic conditions in major consuming regions, particularly the

United States and China. The indicator captures substantial demand contractions during systematic economic downturns, most notably the 2007-2009 global financial crisis and the COVID-19 pandemic. This cyclical sensitivity reflects copper’s extensive utilization in construction, electrical infrastructure, and durable goods production-sectors characterized by high income elasticity and pronounced business cycle exposure.

Taken together, the commodity-level indicators—constructed for crude oil, wheat, corn, sugar, and copper—reveal substantial heterogeneity in the forces driving market developments across commodity groups. Energy and metal markets show pronounced sensitivity to the business cycle, with demand fluctuations closely tracking broader macroeconomic conditions. In contrast, agricultural commodities, including both grains and softs, are more strongly shaped by natural-disaster shocks and exhibit characteristically inelastic demand, reflecting their distinct production technologies and consumption patterns.

3. VALIDATION

In this section, we conduct several validation exercises to assess the reliability of our commodity indicators. First, we illustrate how major supply shocks are reflected in contemporaneous news coverage. Second, we implement human-based auditing to evaluate the accuracy of our article classification procedure. Third, we test whether the signed demand and supply disturbances produce price responses consistent with economic theory, and we employ a placebo experiment to confirm that these correlations are not spurious. Finally, we compare a subset of our indicators—specifically, supply-side disturbances in the oil market—with established measures of oil supply shocks.

3.1. Illustrative example. To assess the empirical validity of our key assumption—that shocks with substantial effects on commodity markets generate immediate news coverage—we examine a setting in which market disruptions stem from exogenous and unexpected events. We focus on seismic activity in Chile, a major copper producer, to evaluate the temporal correspondence between actual supply disturbances and their media footprint. These episodes offer a clean environment for testing whether market-relevant information is reflected in news coverage without delay.

Table 1 documents the correspondence between major seismic events in Chile’s copper-producing regions and our textual supply indicators. Using seismic data from EM-DAT merged with our daily indicators, we report standardized values of both the aggregate copper supply-decrease measure and its natural-disaster component on the trading day following each earthquake. The results reveal a clear and systematic relationship between large seismic events (magnitude above 7.5) and our indicators, with particularly strong responses for earthquakes occurring near mining sites where production capacity is directly affected. This high-frequency alignment provides compelling validation that our textual methodology captures exogenous supply disruptions with appropriate temporal precision.

Post-event dynamics show acute increases in the copper supply-decrease indicator, with amplitudes of 1-2 standard deviations above the mean on the first trading day after the shock. This magnitude is economically meaningful, especially given the wide range of other supply-side disturbances—such as labor disputes, logistical bottlenecks, regulatory changes, energy shortages, mechanical failures, and safety incidents—that routinely generate production losses, operational suspensions, or project delays in copper markets.

To isolate the effects of geological disturbances, we decompose the supply-decrease indicator and analyze the natural-disaster component separately. This more granular measure exhibits markedly stronger reactions, with peaks of 3.5-5.4 standard deviations above baseline levels on the day following major earthquakes. The pronounced responsiveness of this sub-indicator provides strong evidence that our textual methodology not only detects commodity market disruptions but also classifies them with appropriate specificity and temporal accuracy.

Taken together, these results show that financial news coverage incorporates major supply shocks rapidly and with precise characterization. This validation supports the construct validity of our textual indicators as timely and reliable measures of supply-side disruptions in commodity markets.

3.2. Human Auditing. We validate our algorithmic classifications using a structured human-audit protocol designed to minimize bias. Each month, we draw a fixed sample of articles and randomize their evaluation order to reduce exposure to temporal coverage patterns and learning effects. Two independent auditors review each article, assessing commodity relevance, identifying market developments, and evaluating the completeness of dictionary-based classifications.

The correspondence between human- and algorithm-generated indicators is strong. Quarterly correlations reach 0.85 across market dimensions, indicating substantial agreement between the two approaches. Time-series comparisons show similarly tight alignment between human-coded and algorithmic measures. The audit procedure is described in detail in Section C.7 of the Appendix. Figures A5 through A8 display the evolution of human- and algorithm-derived indicators for supply, demand, and composite indices, visually confirming the robustness of our classification methodology.

3.3. Relation to Commodity Price Fluctuations. We next examine the relationship between our supply- and demand-side disturbances and high-frequency movements in commodity prices. Conceptually, commodity price increases should coincide with supply contractions or demand expansions. We provide supporting evidence for this prediction using a panel specification that exploits cross-commodity variation.

3.3.1. Panel Evidence. We focus on the analysis of individual commodities. Specifically, we estimate:

$$R_{i,t} = \beta S_{i,t} + \gamma D_{i,t} + c\mathbf{Z}_{i,t} + u_{i,t}. \quad (4)$$

Our main coefficients of interest are β and γ that capture the impact of a one-standard deviation reduction in supply (an increase in $S_{i,t}$) and increase in demand $D_{i,t}$, respectively, on the daily (log) returns of commodity i . We focus on twenty distinct commodity markets, including energy products (crude oil, gasoil, gasoline, heating oil, natural gas), industrial metals (copper, aluminum, zinc, nickel), precious metals (gold, silver), grains (corn, wheat, soybean), softs (sugar, cocoa, cotton, coffee), and livestock (hogs, live cattle). The vector of controls \mathbf{Z} includes fixed effects for individual commodities and calendar time (day). The dependent variable is the return on spot GSCI component commodity prices between day $t - 1$ and t . We cluster the errors by day and commodity.

Columns (1) to (6) of Table 2 report the estimated coefficients β and γ . These panel estimates reveal a statistically and economically significant relation between our narrative-based supply and demand indicators and individual commodity (log) returns. In columns (1) to (3) the main dependent variables are the net supply and demand indicators while in columns (4) to (6) we replace them with the residualized net supply and demand indicators. Specifications in columns (1) and (4) control for global risk aversion (VIX), monetary conditions (the Fed Funds Rate) and exchange rate movements (the trade-weighted USD index) and include daily fixed effects. Specifications in columns (2) and (5) include both commodity and daily fixed effects and in columns (3) and (6) we also add one lag of the dependent variable.

Overall, a unit standard deviation decline in the residualized supply-side indicator leads to a 7-8 basis point increase in the price of the commodity on the same day; for the demand-side indicator, a one-standard deviation increase in demand leads to a 12-15 basis point increase depending on the specification. The estimated effects on the supply side are approximately half the size of the demand-side coefficients—suggesting asymmetric price responsiveness to supply versus demand innovations. These relationships remain robust across alternative specifications.

Table A2 in the Appendix reports the estimation results for the eighteen commodity markets based on individual time-series regressions. The estimated coefficients on the net supply indicator are uniformly positive and statistically significant at conventional levels for fourteen of the eighteen commodities; only natural gas, nickel, silver, and gold show insignificant supply effects. The net demand indicator exhibits consistently positive and statistically significant effects across all commodities (except aluminum for which the coefficient is statistically insignificant), in line with theoretical predictions.

The robustness of these relationships across heterogeneous commodity markets provides strong evidence for the broad applicability of our text-based methodology. The estimation results in Table 2 and Table A2 show that our commodity-specific supply and demand indicators have significant explanatory power for S&P GSCI single-commodity index returns, even after controlling for global risk aversion (VIX), monetary conditions (the Fed Funds Rate) and exchange rate movements (the trade-weighted USD index). This consistent pattern across diverse commodity classes suggests that

our textual indicators capture economically meaningful variation in the fundamental supply-demand dynamics that drive individual commodity price formation.

3.3.2. *Placebo Test.* Thus far, we have shown that our supply- and demand-side disturbances are significantly related to commodity returns—and, crucially, that the signs of these relationships align with our text-based classification of supply and demand. These relationships remain robust even after controlling for standard financial market variables. The consistent statistical association between our indicators and market-wide investable commodity index returns suggests that our approach captures economically meaningful variation in global commodity market conditions. We next implement a series of falsification tests to verify that these relationships are not spurious. Specifically, we re-estimate equation (4) for the daily GSCI component returns of commodity i , but replace the commodity-specific supply and demand indicators with those constructed for commodity $j \neq i$. This cross-commodity specification serves as a placebo test, allowing us to assess whether the relationships documented in the columns of Table 2 reflect genuine commodity-specific information rather than spurious correlations or omitted variables that might produce artificial statistical significance.

Examining Table 3, we find that none of the cross-commodity supply indicators exhibit statistically significant coefficients at conventional levels. The absence of significant effects provides strong evidence against the concern that our baseline results are driven by common shocks, measurement error, or statistical artifacts. Instead, these null findings support the interpretation that our commodity-specific textual indicators capture genuinely idiosyncratic information relevant to their corresponding markets, thereby reinforcing the credibility of the baseline estimates. In specifications that omit daily temporal fixed effects, demand indicators remain statistically significant, but with markedly attenuated coefficient magnitudes. This systematic reduction in estimated elasticities is consistent with the presence of substantial common components in demand innovations across commodity markets (see Chiaie, Ferrara, and Giannone, 2022). The attenuation pattern provides empirical support for the theoretical expectation that demand shocks are more correlated across commodities than supply shocks, which tend to be more idiosyncratic and commodity-specific. Once temporal fixed effects absorb the shared variation in commodity markets, the remaining identifying variation in demand indicators reflects primarily idiosyncratic components, resulting in diminished coefficients. This empirical regularity aligns with macroeconomic theories that view aggregate demand as a key common factor across commodity classes, typically linked to global business cycle fluctuations, monetary conditions, and broad-based shifts in risk appetite that simultaneously influence multiple commodity markets through parallel channels.

3.4. **Do our news-based indicators contain actual news?** A critical question for interpreting our news-based indicators as genuine shocks is whether they represent information that markets did not

already incorporate into prices. If our indicators merely capture anticipated developments, their correlation with contemporaneous price movements would understate their true impact and limit their usefulness for causal analysis.

Our empirical investigation provides compelling evidence that these indicators indeed represent unexpected news—particularly when analyzed at monthly frequencies. We employ a two-stage approach to rigorously assess this claim.

First, we purge our indicators of any predictable components by estimating AR(p) processes,

$$S_{i,t}^r \equiv S_{i,t} - B(L)S_{i,t} \quad \text{and} \quad D_{i,t}^r \equiv D_{i,t} - B(L)D_{i,t}. \quad (5)$$

We estimate (5) separately for each commodity. Using a conservative lag structure of 20 days ensures we remove even subtle forms of serial dependence.³ We proceed to standardize these residuals to unit standard deviation. These residualized indicators represent the truly unpredictable components of our supply and demand measures.

Second, we estimate the dynamic price response around these news events,

$$R_{i,t+k} = \beta(k) S_{i,t}^r + \gamma(k) D_{i,t}^r + \alpha_i + \mu_t + u_{i,t}. \quad (6)$$

The dependent variable is the daily (log) return on commodity i in day $t + k$, while the key independent variables are our residualized net supply and demand indicators from (5). We include commodity and calendar date fixed effects, and cluster the errors by day and commodity.

Figure 4 plots the estimated coefficients $\beta(k)$ and $\gamma(k)$ for horizons of $k = -20$ to $k = 5$. Examining the figure, two things become apparent. Most striking is the sharp price reaction precisely on the day of news release ($k = 0$). A one standard deviation supply contraction triggers an immediate 7 basis point price increase, while a comparable demand expansion generates a 13 basis point jump. This pronounced contemporaneous response provides definitive evidence that our indicators capture genuine news that market participants immediately incorporate into prices.

Equally revealing is what we don't see in the data. Examining the pre-event window ($k < 0$), we observe essentially zero anticipatory price movements beyond a horizon of roughly ten days before the news release. There is modest evidence of a small drift—approximately 1–2 basis points—confined to the few days immediately preceding publication, primarily for supply indicators. Importantly, this narrow-window drift is economically small (less than 30% of the contemporaneous effect) and is statistically indistinguishable from zero when tested jointly.⁴ Beyond this narrow window, the pre-event coefficients are flat and centered on zero, decisively ruling out the possibility that our indicators merely capture long-anticipated developments.

³We obtain quantitatively similar results with longer lag lengths.

⁴An F -test of the null hypothesis that all pre-event coefficients $\{\beta(k)\}_{k=-20}^{-1}$ are jointly zero fails to reject at conventional significance levels.

Crucially, this pattern poses no concern for our macroeconomic analysis, which is conducted at the monthly frequency. Any minor information leakage occurring a few days before the news release falls *within the same monthly observation window* as the main effect. Once aggregated to monthly or quarterly frequency, the pre-event and contemporaneous effects are subsumed into a single period, so the slight daily-level drift does not attenuate or bias our estimates. Put differently, at the monthly horizon at which we estimate the local projections in Section 4, our indicators behave as clean contemporaneous shocks.

These findings substantially strengthen the case for using our news-based indicators in studying causal relationships between commodity market developments and macroeconomic outcomes.

3.5. Comparison to Existing Indicators. As a final validation exercise, we compare our supply- and demand-side disturbances with established indicators. Because the existing literature has predominantly focused on the oil market, this comparison is feasible only for our oil-specific indicators. The standard approach to disentangling supply- and demand-driven movements in oil prices relies on structural vector autoregressions (SVARs) (see, e.g., Kilian, 2009; Baumeister and Hamilton, 2019). More recently, Känzig (2021) uses high-frequency variation in oil prices around OPEC announcements to identify supply-side shocks. To validate our narrative-based indicators, we focus on two benchmark measures: the oil supply shock series of Baumeister and Hamilton (2019) and the news shock about future oil supply identified by Känzig (2021). We compare the impulse responses of key quantities and prices to these two shocks with those generated by the “surprise” component of our Net Supply Oil Index, defined as the residuals from an AR(1) regression. We implement this comparison using local projections following Jordà (2023):

$$x_{t+k} = \beta(k)S_{oil,t} + \sum_{l=1}^L c_l(k) x_{t-l} + c(k)\mathbf{Z}_t + u_t \quad (7)$$

Figure 5 compares the estimated coefficients $\beta(k)$ from the local projections for the two existing measures (Columns one and two) with those obtained from our oil-specific supply disturbance. The top three rows of the figure report the responses of oil-market prices and quantities. The impulse responses associated with our supply-side indicator closely resemble those of the benchmark measures: all three series show that a negative supply shock leads to higher prices and lower quantities. The fourth row presents the response of OECD inventories, where a negative supply shock initially triggers an accumulation of inventories, followed by a gradual depletion over time—except in the case of the middle shock. The final three rows display the responses of US inflation as well as world and US industrial production. Consistent with theoretical predictions, a negative supply shock is associated with higher US prices and lower industrial production both globally and in the US.

In sum, a negative surprise in our net supply index accurately predicts a future increase in the real WTI price, a significant short- and medium-term decline in global and OPEC oil production, a

reduction in future OECD inventories over time, an increase in US CPI, and a medium-term decline in world industrial production.

To quantify the degree of alignment, we compute correlations between our oil supply surprise series and the benchmark measures. Our Net Supply Oil Index (computed as the residual of an AR(1) process) correlates 29% with the oil supply shock of [Baumeister and Hamilton \(2019\)](#) and 34% with the news shock of [Känzig \(2021\)](#)—for reference, the correlation between these two established measures is 37%. On the demand side, our Net Demand Oil Index correlates 47% with the sum of the three demand components of [Baumeister and Hamilton \(2019\)](#). These moderate correlations are expected given the fundamentally different identification approaches (textual narrative versus structural restrictions versus high-frequency), and the qualitative concordance of impulse responses documented above is arguably more informative than correlation magnitudes.

Notably, despite excluding price data from our methodology, our approach successfully captures market dynamics comparable to those identified by methods that explicitly incorporate oil prices. This demonstrates that textual sources contain sufficient information to recover the underlying structural factors driving commodity markets without relying on endogenous price signals. This concordance serves a crucial methodological purpose: it validates our textual approach in the one commodity market where credible structural benchmarks exist. With this benchmark in place, we can extend our narrative-based methodology beyond oil, providing new measures of supply and demand forces in markets such as natural gas and agricultural commodities, where traditional identification strategies face substantial practical limitations.

4. COMMODITY SUPPLY SHOCKS AND THE BUSINESS CYCLE

We now turn to the main objective of the paper: assessing whether commodity supply shocks drive business-cycle fluctuations. Because our indices span the entire commodity market, we can go beyond the existing literature’s narrow focus on petroleum and examine the macroeconomic contribution of both oil and non-oil supply disturbances—and compare them directly.

4.1. Aggregation. The first step is to aggregate our daily indicators to a monthly or quarterly frequency. Specifically, we define the supply-increase indicator for commodity market i and month (or quarter) τ as:

$$S_{i,\tau}^+ \equiv \sum_{d \in D(\tau)} S_{i,d}^+$$

where the set $D(\tau)$ includes all calendar days in month or quarter τ for which we have articles. The second step is to aggregate the commodity-level indicators into overall indices that reflect the relative importance of each commodity for the aggregate economy. Unlike our earlier composite index, which implicitly weighted commodities based on news coverage, we now use production-based weights.

The existing literature has largely focused on supply shocks in the oil market. To underscore the role of non-oil supply shocks for the business cycle, we construct two indicators: one specific to the oil market and another covering all commodities *excluding* oil. Accordingly, our aggregate index capturing supply increases in non-oil commodities is defined as:

$$S_{non-oil,\tau}^+ \equiv \sum_{i \in \mathcal{T} \setminus Oil} w_{i,\tau} S_{i,\tau}^+,$$

where w_i denotes the production weight for commodity i , obtained from the GSCI Composite construction methodology. These weights, calculated as five-year moving averages, provide a market-based measure of each commodity's relative importance in global production.⁵ The indicators for supply decreases and for demand increases or decreases are defined analogously. Finally, we construct the net supply and net demand indicators following equation (3).

As a preliminary validation of the aggregated indices, Figure 6 displays the medium-run response of real commodity spot prices to a one-standard-deviation supply contraction. For the oil index, Brent prices rise by approximately 3% over the following year; for the non-oil index, the production-weighted commodity basket increases by 5–7%. These magnitudes confirm that the aggregated supply indices capture economically meaningful variation in commodity markets and provide a natural benchmark for interpreting the macroeconomic effects that follow.

4.2. Macroeconomic effects. To quantify the macroeconomic impact of commodity supply disturbances, we estimate local projections following Jordà (2023) on a panel of 60 countries encompassing both advanced and emerging market economies. Macroeconomic variables—consumer price indices and industrial production—are obtained from the International Monetary Fund's International Financial Statistics (IFS) database. To measure each country's exposure to commodity markets, we construct a normalized trade balance for each commodity category using disaggregated trade flow data from the United Nations Comtrade database:

$$TB_{i,t} = \frac{X_{i,t} - M_{i,t}}{X_{i,t} + M_{i,t}}, \quad (8)$$

where $X_{i,t}$ and $M_{i,t}$ denote the exports and imports of country i in period t , respectively. The index $TB_{i,t}$ lies in $[-1, 1]$: a country is a net exporter when $TB_{i,t} > 0$ and a net importer when $TB_{i,t} < 0$. This measure captures both the direction and the intensity of commodity dependence.

⁵We implement a methodological adjustment to account for the absence of lead from our composite indices. This reweighting ensures that the relative importance of commodities is properly represented in both the aggregate and non-petroleum composite measures, preserving the economic interpretation of the transmission estimates despite incomplete coverage relative to the standard GSCI production-weighted benchmark. For each year after 2010, we use the reference production dollar weight (RPDW) for the GSCI index in that year, obtained from documents published by S&P. As no publicly available document exists for 2009 and prior, we use the RPDW for 2009 listed in the 2010 publication; for earlier years, we use the 2009 weight.

The construction of the supply disturbance $S_{oil/non-oil,t}$ proceeds through a multi-step residualization designed to isolate genuinely unexpected supply innovations. We first remove autocorrelation from the daily supply index by estimating an AR(20) process and retaining the residuals. Next, for the oil (or non-oil) market, we control for own-market demand by including the one-day lagged AR(20) residuals of the oil (or non-oil) demand indicator. To address cross-market influences, we further include the one-day lagged AR(20) residuals of the non-oil (or oil) supply and demand. These daily residuals are then aggregated to monthly frequency.

The panel local projection specification is:

$$x_{i,t+k} = \beta_i(k) S_{oil/non-oil,t} + \sum_{l=1}^L c_{i,l}(k) x_{i,t-l} + c_i(k) \mathbf{Z}_{i,t} + \alpha_i(k) + u_{i,t}. \quad (9)$$

where i indexes countries, $\alpha_i(k)$ represents country fixed effects, and $\mathbf{Z}_{i,t}$ is a vector of macroeconomic and financial controls. We include $L = 12$ lags of the dependent variable and adjust the standard errors to account for heteroskedasticity and cross-sectional dependence arising from shared global shocks following [Driscoll and Kraay \(1998\)](#).

Average effects. Figure 7 presents the estimated average impulse responses across all 60 countries. Commodity supply shocks generate economically meaningful macroeconomic effects: a one-standard-deviation non-oil supply contraction raises CPI by approximately 5 percentage points and contracts industrial production by roughly 2 percent; the output decline is frontloaded and statistically significant over the first months. Oil supply disturbances generate effects of comparable order—CPI responses are somewhat weaker and become marginally significant only at longer horizons, while industrial production falls by about 1–2 percent. A striking result thus emerges: non-oil supply disturbances matter at least as much as oil disturbances for global inflation and output, a finding previously undocumented in the cross-country literature.

The finding that non-oil supply disturbances generate effects at least as large as oil disturbances merits discussion. Several factors may contribute to this result. First, non-oil commodities collectively represent a broad set of inputs across manufacturing, construction, and agriculture, so a one-standard-deviation movement in the production-weighted non-oil index may reflect correlated disruptions across multiple sectors simultaneously, amplifying the aggregate impact relative to a single-commodity oil shock. Second, substitutability differs across commodity classes: oil has well-developed substitutes (natural gas, renewables, energy efficiency), whereas many non-oil inputs—such as specific industrial metals for electronics or wheat for food production—face tighter substitution constraints, making their supply shortfalls more binding on downstream production. Third, the 2001–2023 sample period encompasses the US shale revolution, which fundamentally expanded domestic

oil supply capacity and reduced the economy’s vulnerability to oil-specific disruptions. The declining macroeconomic sensitivity to oil supply shocks may thus partly reflect a secular structural change rather than a universal ranking of commodity importance.

Heterogeneous transmission. Whether commodity supply shocks matter for the global business cycle depends not only on average effects but on how transmission varies with countries’ exposure. Net importers face higher input costs and supply bottlenecks when world supply contracts, whereas net exporters may benefit from price gains. To test this heterogeneity, we interact the supply disturbance with the normalized trade balance:

$$x_{i,t+k} = \beta_i(k) S_{oil/non-oil,t} + \gamma_i(k) (S_{oil/non-oil,t} \times TB_{i,t}) + \sum_{l=1}^L c_{i,l}(k) x_{i,t-l} + c_i(k) \mathbf{Z}_{i,t} + \alpha_i(k) + \delta_t(k) + u_{i,t}. \quad (10)$$

where $\alpha_i(k)$ are country fixed effects and $\delta_t(k)$ are time fixed effects. Time fixed effects absorb all common global variation, so that $\gamma_i(k)$ is identified solely from *cross-sectional* variation in trade balance within each period. Country fixed effects remove time-invariant differences in macroeconomic sensitivity (e.g., sectoral structure, average openness).

Figure 8 presents the estimated interaction coefficients for industrial production. The coefficient $\gamma_i(k)$ is negative: net importers ($TB_{i,t} < 0$) exhibit significantly larger and more persistent output declines following supply contractions than net exporters ($TB_{i,t} > 0$). This pattern holds for both oil and non-oil disturbances and is robust to the inclusion of time fixed effects (Panel B).

A notable feature of the results is the difference in timing between the average and the differential response. The average output contraction (Figure 7, Panel B) is frontloaded, with the bulk of the decline materializing within the first few months. By contrast, the importer–exporter gap (Figure 8) is small on impact and widens in the medium run, peaking roughly one year after the shock. Several mechanisms can account for this pattern. First, commodity inventories buffer the initial disruption symmetrically: firms in all countries draw down stocks to sustain production. As buffers deplete, importers—who depend on world markets for restocking—face increasingly binding constraints, while exporters replenish from domestic production. Second, terms-of-trade income effects build gradually: higher commodity prices generate export revenues, fiscal windfalls, and currency appreciation that support demand in exporting countries with a lag, while the corresponding income drain in importing countries accumulates as trade-balance deterioration feeds through to spending. Third, supply disruptions propagate through input-output networks over quarters, not weeks; in importing countries, where no domestic production can substitute for lost imports, this cascade is more severe. Finally, divergent policy responses may contribute: monetary tightening in response to commodity-driven inflation imposes a contractionary impulse in importing countries with the usual

six-to-eighteen-month lag, while exporting countries may ease fiscal policy from revenue gains. Together, the initial shock is largely common across countries; the structural divergence—determined by trade position, inventory capacity, and policy space—takes time to materialize.

Descriptive evidence corroborates this finding. Appendix Figure A10 stratifies the sample into net importers and net exporters: industrial production contracts significantly more in importing countries, while exporter responses are smaller and less precisely estimated, consistent with domestic production partially insulating these economies.

Turning to prices, one might expect CPI responses to also differ by trade position—for instance, through exchange rate appreciation in exporting countries. However, the primary channel is the direct effect of world commodity prices, which rise identically for all countries regardless of trade position. Consistent with this reasoning, the CPI interaction coefficients show weaker and less robust importer–exporter differences than the output effects (Appendix Figure A11), indicating that international adjustment to commodity supply shocks occurs primarily through quantities rather than prices.

These international results are consistent with US-only evidence. Appendix Figure A9 reports local projections for the United States, confirming that supply contractions raise CPI and lower industrial production, and that non-oil supply shocks generate effects at least as large as oil shocks over the sample period.

In sum, commodity supply shocks are global in reach but heterogeneous in impact. Net commodity importers face systematically larger and more persistent output losses, while net exporters are partially insulated—a gap that is concentrated in quantities rather than prices, and that the interaction specification with country and time fixed effects pins to the causal channel linking supply disruptions to countries’ commodity exposure.

5. DECOMPOSING THE DRIVERS OF SUPPLY AND DEMAND DEVELOPMENTS

Different types of supply and demand shocks can have distinct implications for macroeconomic outcomes, financial markets, and price stability (e.g., Caballero, Farhi, and Gourinchas, 2008). For example, natural disasters typically affect inflation over the short to medium term, whereas environmental regulations can have a more persistent impact by permanently altering relative prices through expectations. Decomposing commodity price fluctuations into their key drivers allows us to distinguish between short- and long-term effects, and to assess the severity and transmission mechanisms of each shock. In the following, we present the main drivers of commodity supply and demand developments. Drawing on the commodity risk factors highlighted in the *Commodity Special Feature* of the IMF’s World Economic Outlook and the World Bank’s *Commodity Markets Outlook*, we group these into four broad categories that represent the main drivers of commodity market dynamics: business cycle developments, geopolitical risk, natural disasters, and climate change.

Figure 9 displays the business cycle driver for the period 2001-2023. The series captures major global macroeconomic developments, with the largest spikes corresponding to the 2001 recession, the Global Financial Crisis, the August 2011 stock market decline driven by fears of European sovereign debt contagion, the COVID-19 crisis, and the US debt-ceiling episode. Figure 10 presents the geopolitical risk driver over the same period. The indicator reflects several major geopolitical events, including the 9/11 attacks, the 2003 Iraq invasion, the Arab Spring, the escalation of the US-China trade war, and the Russian invasion of Ukraine. Figure 11 shows the natural disasters driver, highlighting events such as the 2001 Gujarat earthquake, the intense Atlantic hurricane seasons (e.g., Katrina, Emily), the 2012 US summer drought, and the COVID-19 crisis. Finally, Figure 12 plots the climate change driver for 2005-2023. The series captures key climate-related developments, warnings, and policy initiatives, including European Commission proposals and US Environmental Protection Agency regulations. Decomposing the drivers across commodity categories reveals that: (i) the business cycle is most relevant for energy, industrial metals, and soft commodities that are also used as biofuels, (ii) geopolitical risk events tend to affect energy and soft commodities, with grains becoming important primarily during the trade war period, (iii) natural disasters mainly impact grains and soft commodities, and to a lesser extent, energy, and (iv) climate change effects are concentrated in energy commodities and, more recently, industrial metals.

A key methodological contribution of our analysis is the decomposition of commodity supply and demand components into their underlying fundamental drivers. As shown in Figure 13, net supply fluctuations across commodity categories are primarily driven by exogenous events, such as natural disasters. By contrast, net demand variations are largely shaped by endogenous business cycle dynamics, with the notable exception of the COVID-19 pandemic, which manifests as an unusually large exogenous negative natural disaster shock in 2020. This contrasting pattern of driver importance aligns with theoretical expectations regarding the relative exogeneity of supply disruptions versus the endogenous nature of demand fluctuations in commodity markets.

6. CONCLUDING REMARKS

Identifying supply-side disturbances that drive business-cycle fluctuations has been a central challenge since [Kydland and Prescott \(1982\)](#). This paper develops a comprehensive textual methodology to measure supply and demand shocks across the full commodity spectrum—energy, metals, agriculture, and livestock—providing empirical traction on a question that has been largely confined to the oil market. A striking result emerges: supply disruptions in non-oil commodity markets generate macroeconomic effects at least as large as those originating in petroleum markets. This finding suggests that the literature’s narrow focus on oil has substantially understated the role of commodity supply shocks in driving aggregate fluctuations.

Our decomposition of fundamental drivers across commodity groups uncovers clear systematic patterns. Business cycle fluctuations predominantly influence energy and industrial metals; geopolitical risks disproportionately affect energy and precious metals; natural disasters primarily disrupt agricultural commodities with spillovers into energy markets; and climate-related risks increasingly shape both energy and industrial metals markets.

The heterogeneous transmission documented in our cross-country analysis carries direct policy implications. For import-dependent economies, the vulnerability to non-oil commodity supply disruptions suggests that strategic diversification of import sources and the development of buffer stocks—currently focused almost exclusively on petroleum—should be extended to critical non-oil commodities such as industrial metals and key agricultural inputs. For monetary policymakers, the finding that commodity supply shocks generate both inflationary and contractionary effects poses the classical dilemma: our results indicate that the inflation pass-through is strongest and most persistent for net importers, suggesting that central banks in commodity-importing economies may need to give greater weight to non-oil commodity developments in their reaction functions. More broadly, the quantity-versus-price adjustment pattern we document—where the output response dominates the price response for net importers—implies that the real costs of commodity supply disruptions are borne disproportionately through reduced economic activity rather than through price adjustments alone, underscoring the importance of supply-side resilience in commodity-dependent economies.

Taken together, these findings make the case for a comprehensive framework for understanding commodity market dynamics—one that extends well beyond oil. By mapping the multidimensional sources of commodity price movements, our approach uncovers both enduring regularities and emerging trends in global commodity markets, with direct implications for macroeconomic policy.

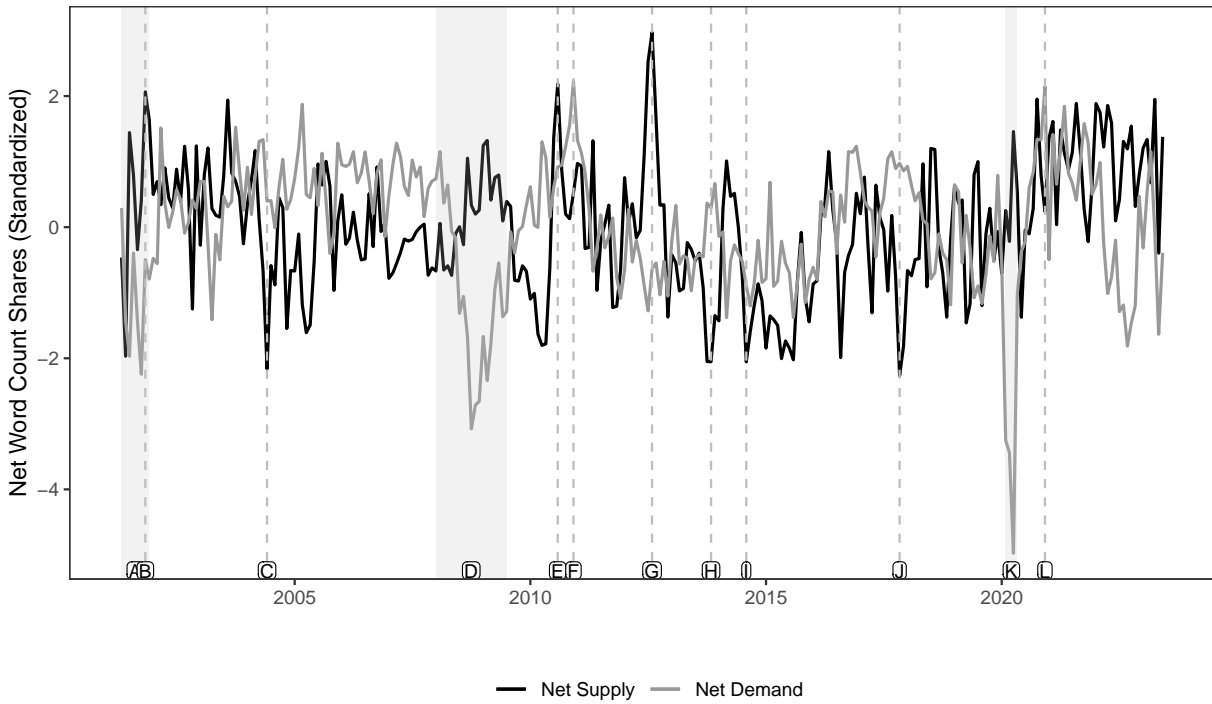
REFERENCES

- Acemoglu, Daron, Vasco Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80 (5):1977–2016.
- Adjemian, Michael and Jungkeon Jo. 2025. "The Macroeconomic Effects of Agricultural Supply News." Working Paper.
- Alexopoulos, Michelle and Jon Cohen. 2015. "The power of print: Uncertainty shocks, markets, and the economy." *International Review of Economics & Finance* 40 (C):8–28.
- Allcott, Hunt and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31 (2):211–236.
- Antolín-Díaz, Juan and Juan F. Rubio-Ramírez. 2018. "Narrative Sign Restrictions for SVARs." *American Economic Review* 108 (10):2802–2829.
- Atalay, Englin. 2017. "How Important Are Sectoral Shocks?" *American Economic Journal: Macroeconomics* 9 (4):254–280.
- Baker, Scott, Nicholas Bloom, and Steven Davis. 2016. "Measuring Economic Policy Uncertainty." *The Quarterly Journal of Economics* 131 (4):1593–1636.
- Baqae, David and Emmanuel Farhi. 2022. "Supply and Demand in Disaggregated Keynesian Economies with an Application to the Covid-19 Crisis." *American Economic Review* 112 (5):1397–1436.
- Baumeister, Christiane and James D. Hamilton. 2019. "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks." *American Economic Review* 109 (5):1873–1910.
- Baumeister, Christiane and Gert Peersman. 2013. "Time-Varying Effects of Oil Supply Shocks on the US Economy." *American Economic Journal: Macroeconomics* 5 (4):1–28.
- Bernanke, Ben and Olivier Blanchard. 2025. "What Caused the US Pandemic-Era Inflation?" *American Economic Journal: Macroeconomics* 17 (3):1–35.
- Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson. 2013. "Which News Moves Stock Prices? A Textual Analysis." NBER Working Papers 18725, National Bureau of Economic Research, Inc.
- Caballero, Ricardo J., Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2008. "Financial Crash, Commodity Prices and Global Imbalances." NBER Working Papers 14521, National Bureau of Economic Research, Inc.
- Caldara, Dario, Michele Cavallo, and Matteo Iacoviello. 2019. "Oil price elasticities and oil price fluctuations." *Journal of Monetary Economics* 103 (C):1–20.

- Castelnuovo, Efrem, Lorenzo Mori, and Gert Peersman. 2024. "Commodity Price Shocks and Global Cycles: Monetary Policy Matters." Working Papers of Faculty of Economics and Business Administration, Ghent University, Belgium 24/1087, Ghent University, Faculty of Economics and Business Administration.
- Chiaie, Simona Delle, Laurent Ferrara, and Domenico Giannone. 2022. "Common factors of commodity prices." *Journal of Applied Econometrics* 37 (3):461–476.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1):1–45.
- Datta, D. Deepa and A. Daniel Dias. 2019. "Oil Shocks: A Textual Analysis Approach." Tech. rep.
- Driscoll, John C. and Aart C. Kraay. 1998. "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data." *Review of Economics and Statistics* 80:549–560.
- Engle, Robert F, Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. 2020. "Hedging Climate Change News." *Review of Financial Studies* 33 (3):1184–1216.
- Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, NJ: Princeton University Press, 2 ed.
- Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78 (1):35–71.
- Hamilton, James D. 1983. "Oil and the Macroeconomy since World War II." *Journal of Political Economy* 91 (2):228–248.
- Hamilton, James D. 2003. "What is an oil shock?" *Journal of Econometrics* 113 (2):363–398.
- Hassan, Tarek A, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. "Firm-Level Political Risk: Measurement and Effects." *The Quarterly Journal of Economics* 134 (4):2135–2202.
- Hoberg, Gerard and Gordon Phillips. 2010. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *Review of Financial Studies* 23 (10):3773–3811.
- Horvath, Michael. 2000. "Sectoral Shocks and Aggregate Fluctuations." *Journal of Monetary Economics* 45 (1):69–106.
- Jordà, Òscar. 2023. "Local Projections for Applied Economics." *Annual Review of Economics* 15 (1):607–631.
- Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review* 111 (4):1092–1125.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review* 108 (3):697–743.
- Kilian, Lutz. 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *The Review of Economics and Statistics* 90 (2):216–240.
- . 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3):1053–1069.

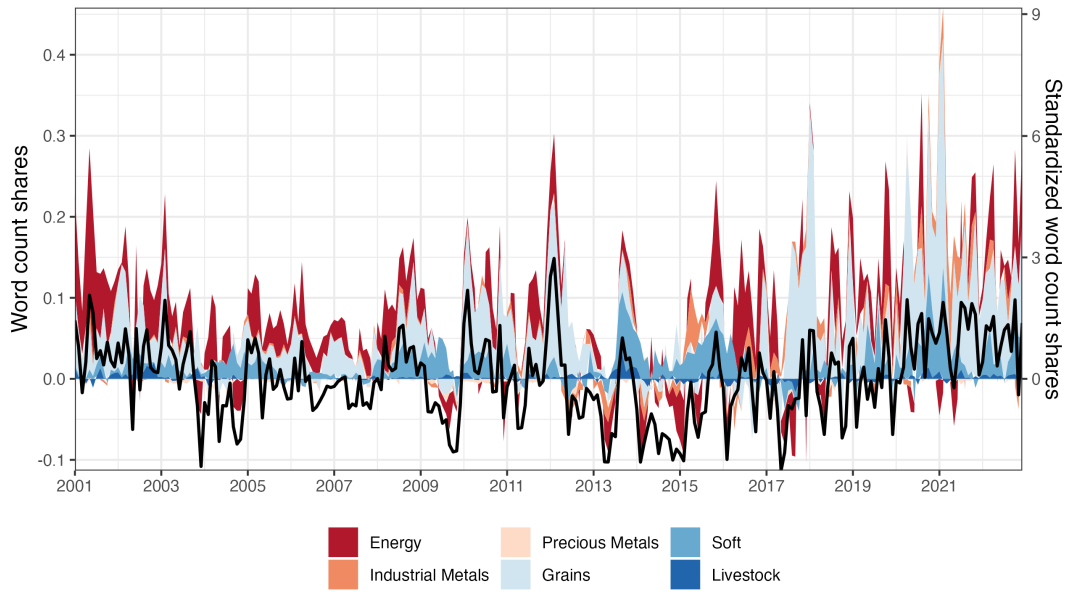
- Kilian, Lutz and Daniel P. Murphy. 2012. "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics Of Oil Market Var Models." *Journal of the European Economic Association* 10 (5):1166–1188.
- . 2014. "The Role Of Inventories And Speculative Trading In The Global Market For Crude Oil." *Journal of Applied Econometrics* 29 (3):454–478.
- Kilian, Lutz, Alessandro Rebucci, and Nikola Spatafora. 2009. "Oil shocks and external balances." *Journal of International Economics* 77 (2):181–194.
- Kilian, Lutz and Robert J. Vigfusson. 2017. "The Role of Oil Price Shocks in Causing U.S. Recessions." *Journal of Money, Credit and Banking* 49 (8):1747–1776.
- Kydland, Finn E and Edward C Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50 (6):1345–1370.
- Lippi, Francesco and Andrea Nobili. 2012. "Oil And The Macroeconomy: A Quantitative Structural Analysis." *Journal of the European Economic Association* 10 (5):1059–1083.
- Long, John and Charles Plosser. 1983. "Real Business Cycles." *Journal of Political Economy* 91 (1):39–69.
- Loughran, Tim, Bill McDonald, and Ioannis Pravidis. 2019. "Assimilation of oil news into prices." *International Review of Financial Analysis* 63 (C):105–118.
- Nakamura, Emi and Jón Steinsson. 2014. "Fiscal Stimulus in a Monetary Union: Evidence from US Regions." *American Economic Review* 104 (3):753–92.
- Romer, Christina D. and David H. Romer. 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review* 100 (3):763–801.
- Tetlock, Paul C. 2007. "Giving Content to Investor Sentiment: The Role of Media in the Stock Market." *Journal of Finance* 62 (3):1139–1168.
- Winne, Jasmien De and Gert Peersman. 2016. "Macroeconomic Effects of Disruptions in Global Food Commodity Markets: Evidence for the United States." *Brookings Papers on Economic Activity* 47 (2 (Fall)):183–286.
- Wu, Tao and Michele Cavallo. 2012. "Measuring Oil-Price Shocks Using Market-Based Information." IMF Working Papers 12/19, International Monetary Fund.
- Zhou, Xiaoqing. 2019. "Refining the Workhorse Oil Market Model." Working Papers 1910, Federal Reserve Bank of Dallas.

FIGURE 1. Net supply and net demand indices across all commodity-related news articles




Note: This figure plots the standardized net supply and net demand indicators from 2001 to 2023, constructed using all commodity-related news articles in the Reuters corpus. The bars and lines map a number of well-known commodity-wide developments that exceeded ± 2 S.D. These events are: [A] US Recession and 9/11, [B] OPEC 118th Extraordinary Meeting (production cut by 1mbpd following US Recession and 9/11), [C] OPEC 131st Extraordinary Meeting (production increase by 2mbpd), [D] 2008 Global Financial Crisis, [E] Multiple natural disasters in Pakistan, Russia, China, and India, [F] Positive outlook on global recovery increased global demand, [G] Worst drought in more than 50 years in the U.S. severely cut corn and soybeans productions, [H] Record US corn harvest, third-largest US soybean harvest, [I] Record grain harvest in US, Brazil, and South Africa, [J] Record-high US crude production, 5-year high crude production in Libya, [K] Onset of Covid-19 Pandemic, US Recession, [L] Covid vaccine approval increases broad-based demand for commodities

FIGURE 2. Decomposition of composite indices across categories




Panel A: Net supply



Panel B: Net demand

FIGURE 3. Crude oil index: Net supply and net demand



Note: This figure plots the standardized net supply and demand indicators for the period between 2021 and 2023. The bars map a number of well-known oil sector developments that are in the top 1% (two-tailed) of events. These events are: [A] US Recession, [B] OPEC 118th Extraordinary Meeting (production cut by 1mbpd following US Recession and 9/11), [C] OPEC 131st Extraordinary Meeting (production increase by 2mbpd), [D] Niger Delta conflict affecting Nigeria oil output, [E] OPEC 135th Meeting (production increase by 0.5mbpd, production at record high), [F] Global Financial Crisis, [G] U.S. record oil production (above 10 million bpd for the first time since the 1970's) fed worries about oversupply, [H] OPEC+ extended production cut, US imposed sanction on Venezuela, Iran, [I] Onset of Covid-19 Pandemic, US Recession, [J] Bad US unemployment data and strong dollar depressed oil demand, [K] Saudi announced 1mbpd voluntary cut, winter storm in Texas cut U.S. production by more than 10%, [L] Very positive global demand forecast due to vaccine rollout and expected recovery, [M] Global oil demand jumped by 1.4mbpd (mtm) on continued recovery, [N] Major economies started to reopen borders and China lifted strict COVID-19 restrictions, boosting demand growth expectations

FIGURE 4. Are our indicators anticipated?



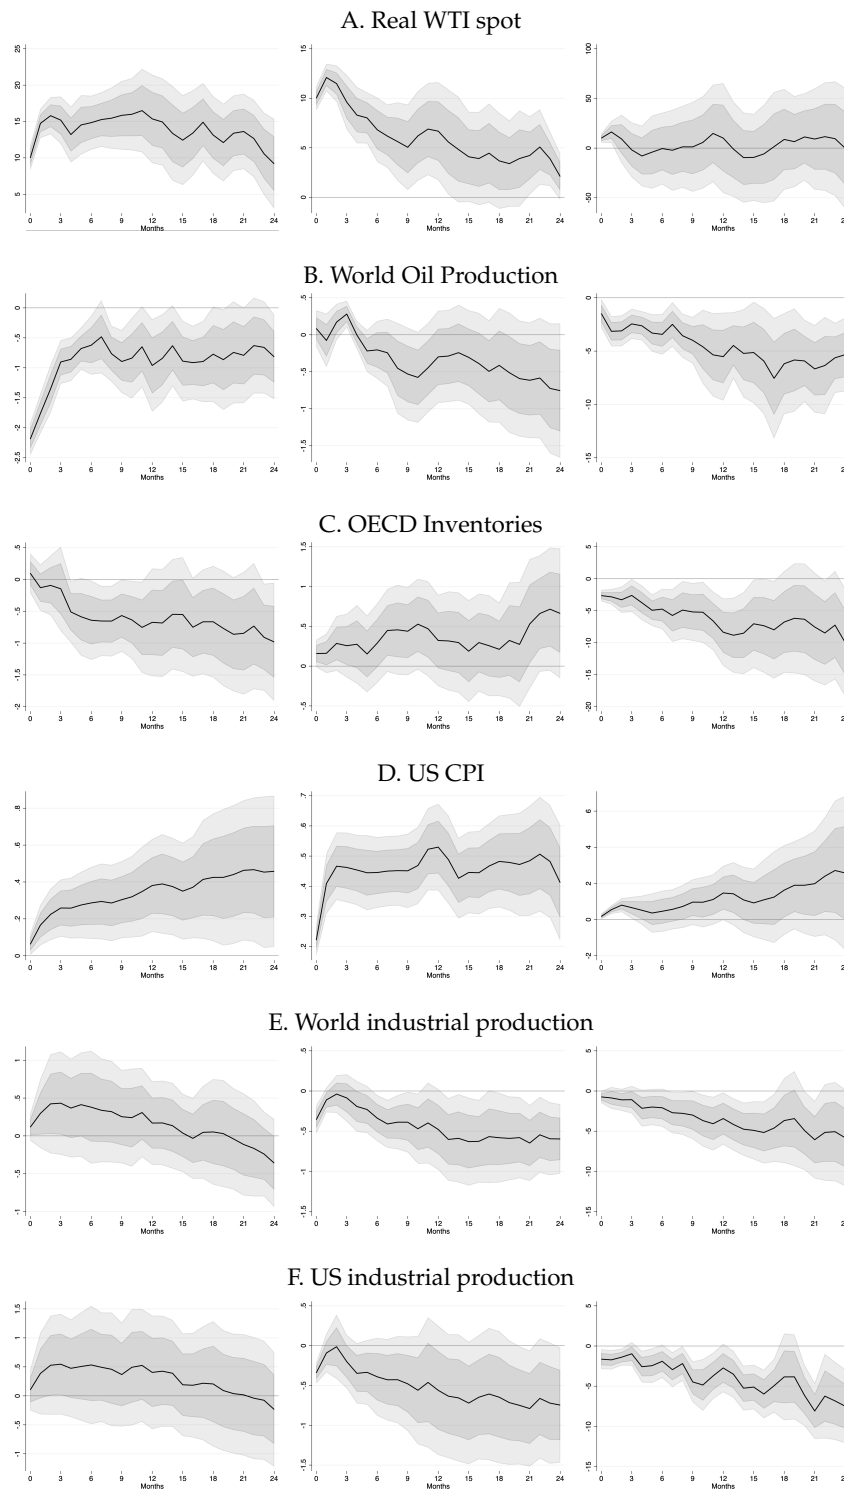


FIGURE 5. Local projections impulse responses




(I) Baumeister & Hamilton, 2019

(II) Känzig, 2021

(III) LMPRP, 2025.

Note: This figure shows the impulse responses to a one standard deviation oil supply shock for the Real Spot Price of Oil, World Production, OPEC Production, OECD Inventories, US CPI and World Industrial Production, employing the oil supply shock series of [Baumeister and Hamilton \(2019\)](#), [Känzig \(2021\)](#) and a “surprise” series, calculated from our Net Supply Crude Oil Index as the residual of an AR(1), reported in Columns, I, II, and III, respectively.

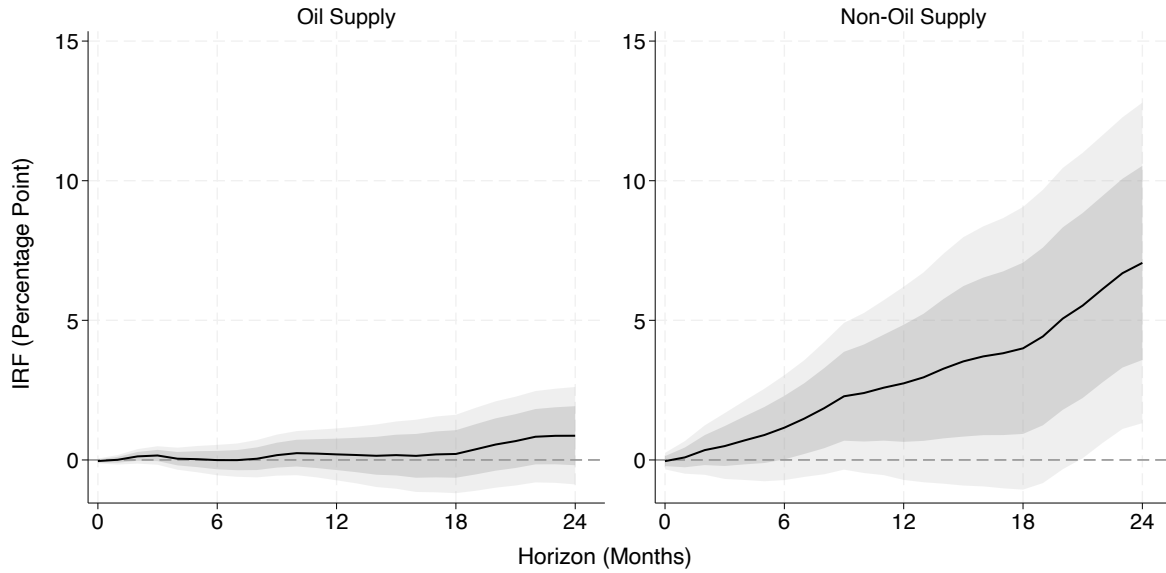
FIGURE 6. Commodity price response to supply indices



Note: This figure shows the impulse responses of real commodity spot prices to a one-standard-deviation supply contraction. The left panel uses the oil supply index with Brent as the dependent variable; the right panel uses the non-oil production-weighted supply index with a weighted non-oil commodity basket as the dependent variable. Supply indices are constructed as AR(20) residuals aggregated to monthly frequency.

FIGURE 7. Local projections: Oil and non-oil impulse responses

A. World CPI (left: Brent; right: weighted non-oil commodity basket)



B. World Industrial Production (left: Brent; right: weighted non-oil commodity basket)

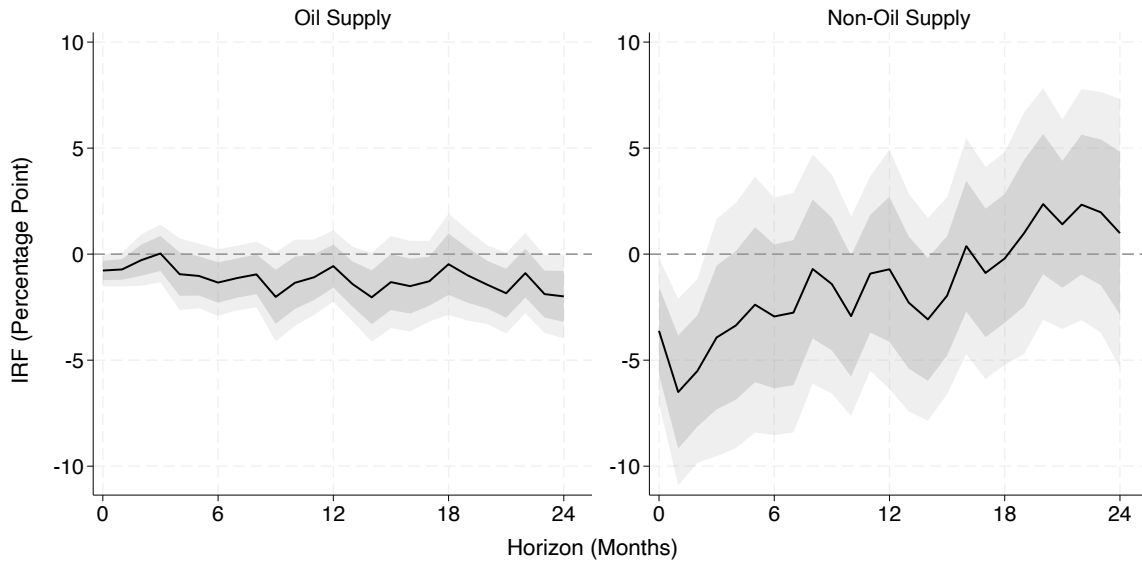
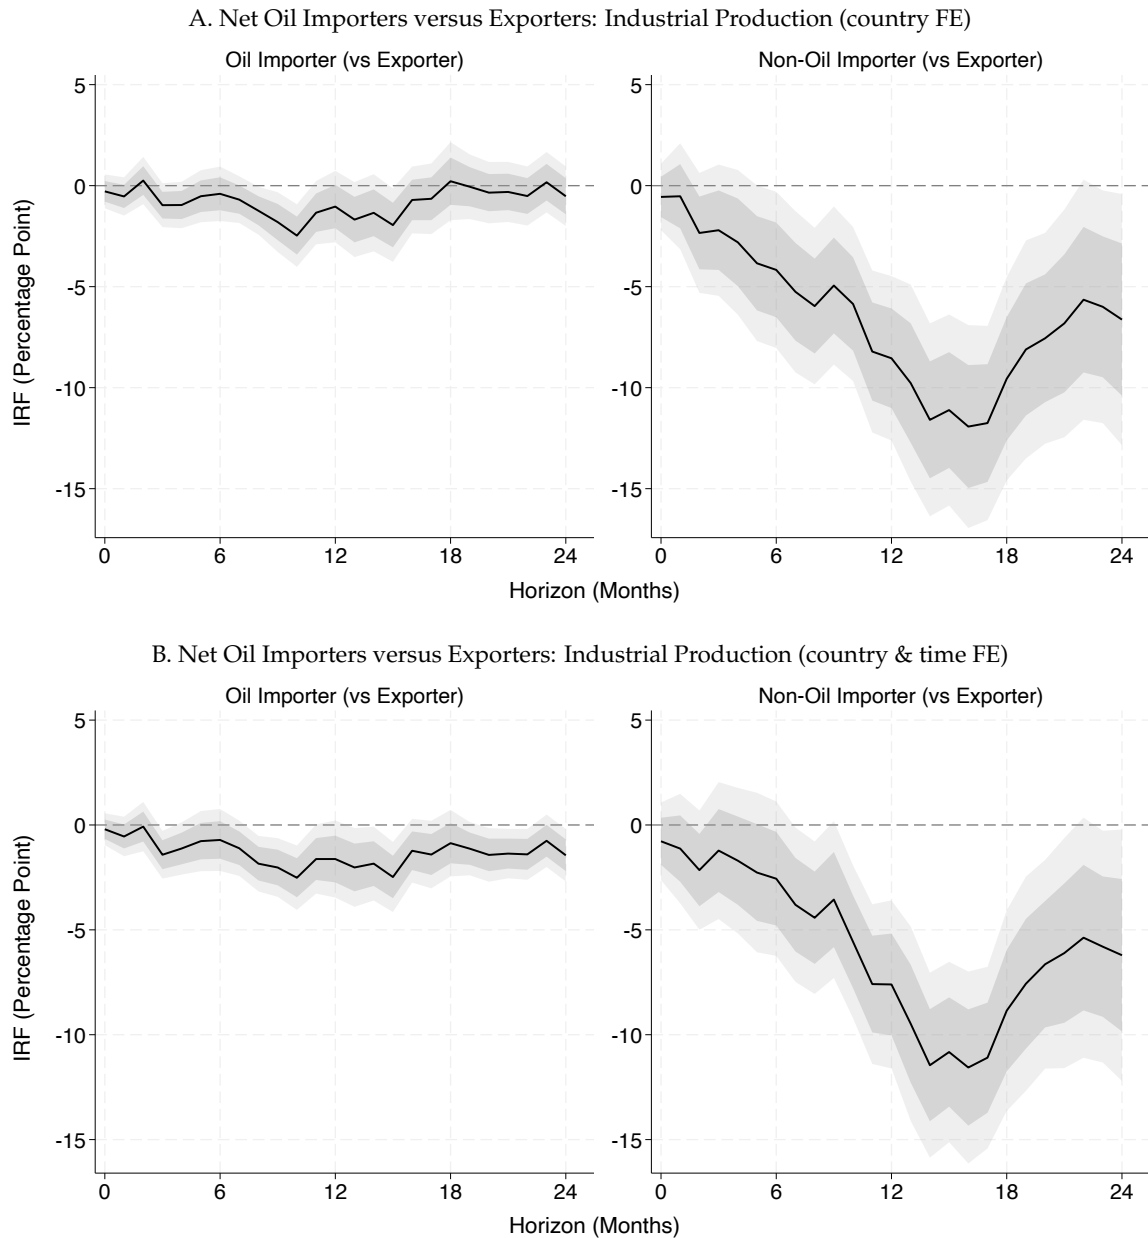



FIGURE 8. Local Projections: Importers versus Exporters Industrial Production Responses




Note: This figure shows the differential impulse response of industrial production to a one-standard-deviation commodity supply contraction, interacted with the normalized trade balance. Panel A includes country fixed effects; Panel B adds time fixed effects. The left column uses the oil supply index; the right column uses the non-oil production-weighted supply index. Time fixed effects absorb common global shocks, so identification comes from cross-sectional variation in the trade balance within each period. Driscoll-Kraay standard errors are used.

FIGURE 9. Business cycle

A. Occurrence in news articles




B. Decomposition across commodity categories



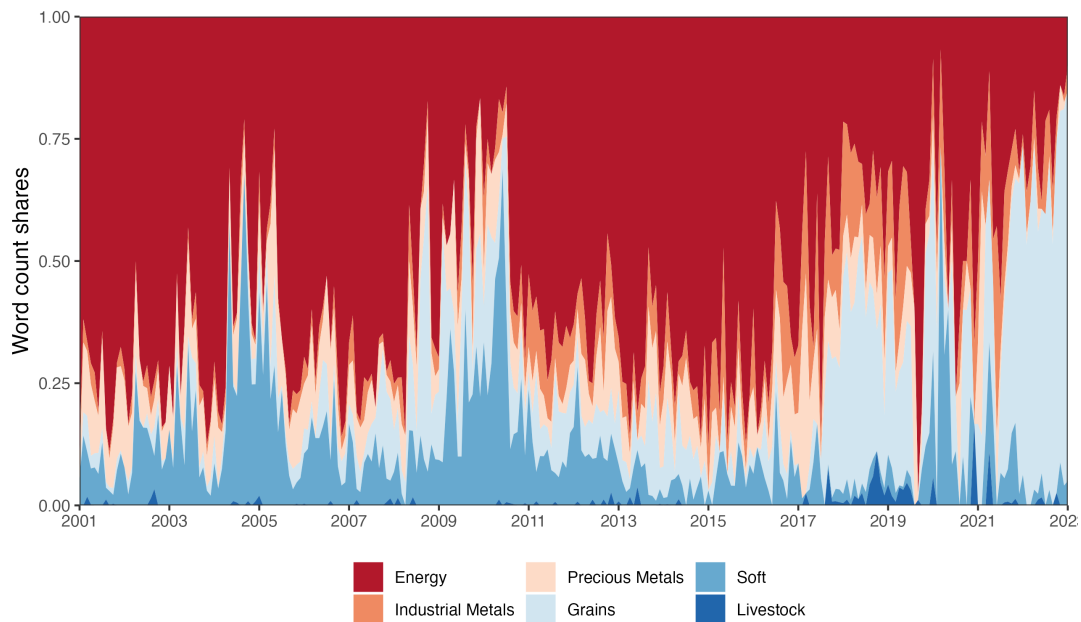
Note: This figure plots the Business Cycle Theme for the period between 2001 and 2023. The bars map a number of important economic developments (i.e. in the top 1% (two-tailed) of events) reflecting the state of the business cycle. These events respectively are: [A] US Recession and 9/11, [B] Global Financial Crisis, [C] Onset of Covid-19 Pandemic, US Recession, [D] Bullish sentiment following signs of recovery and fiscal stimulus, [E] Global central banks' steep monetary tightening induced global market sell-off

FIGURE 10. Geopolitical risk

A. Occurrence in news articles




B. Decomposition across commodity categories



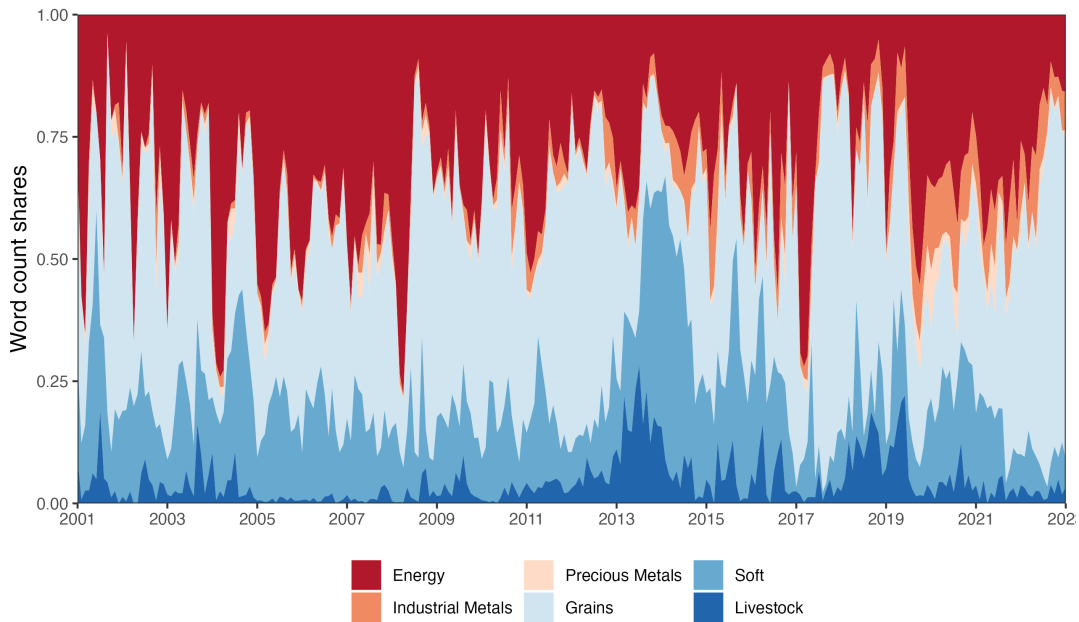
Note: This figure plots the Geopolitical Risk Theme for the period between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) geopolitical events. These events respectively are: [A] 9/11 attacks, [B] 2003 Iraq invasion, [C] China-United States trade war, [D] Russian invasion of Ukraine.

FIGURE 11. Natural disaster

A. Occurrence in news articles




B. Decomposition across commodity categories




Note: This figure plots the Natural Disasters Theme for the period between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) natural disasters. These events respectively are: [A] 2005 Atlantic hurricane season (Katrina, Emily, Rita, Wilma), [B] US Summer 2012 Drought, [C] Onset of Covid-19 Pandemic

FIGURE 12. Climate change

A. Occurrence in news articles



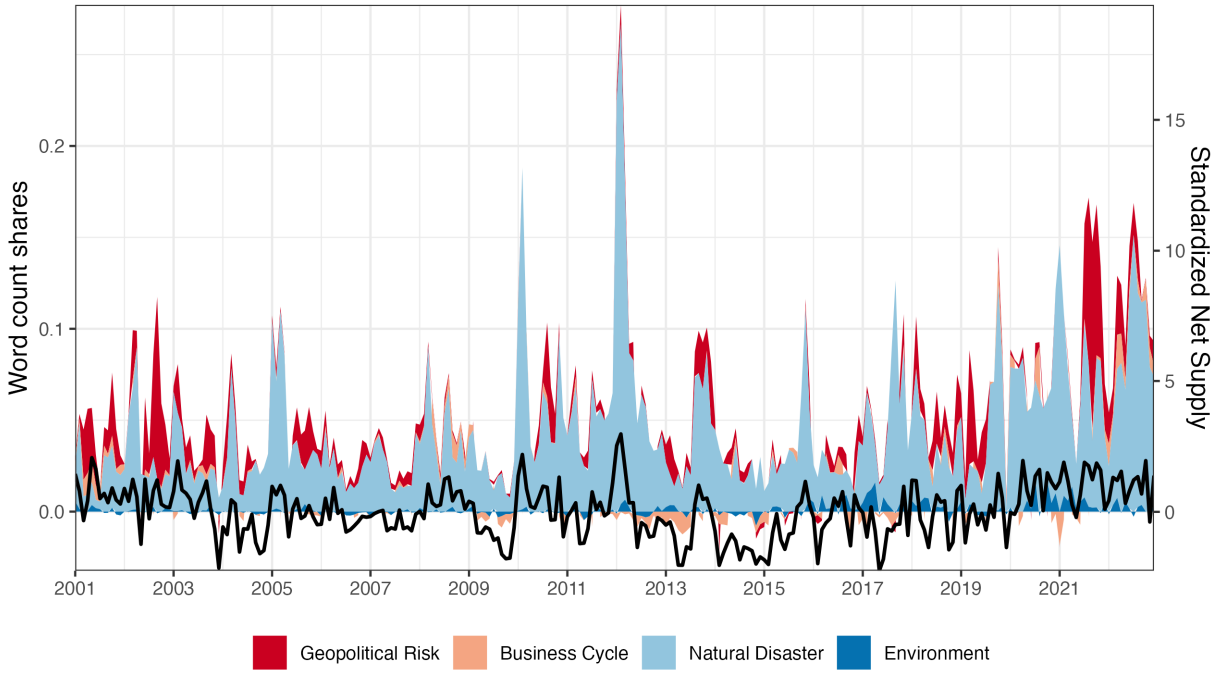
B. Decomposition across commodity categories



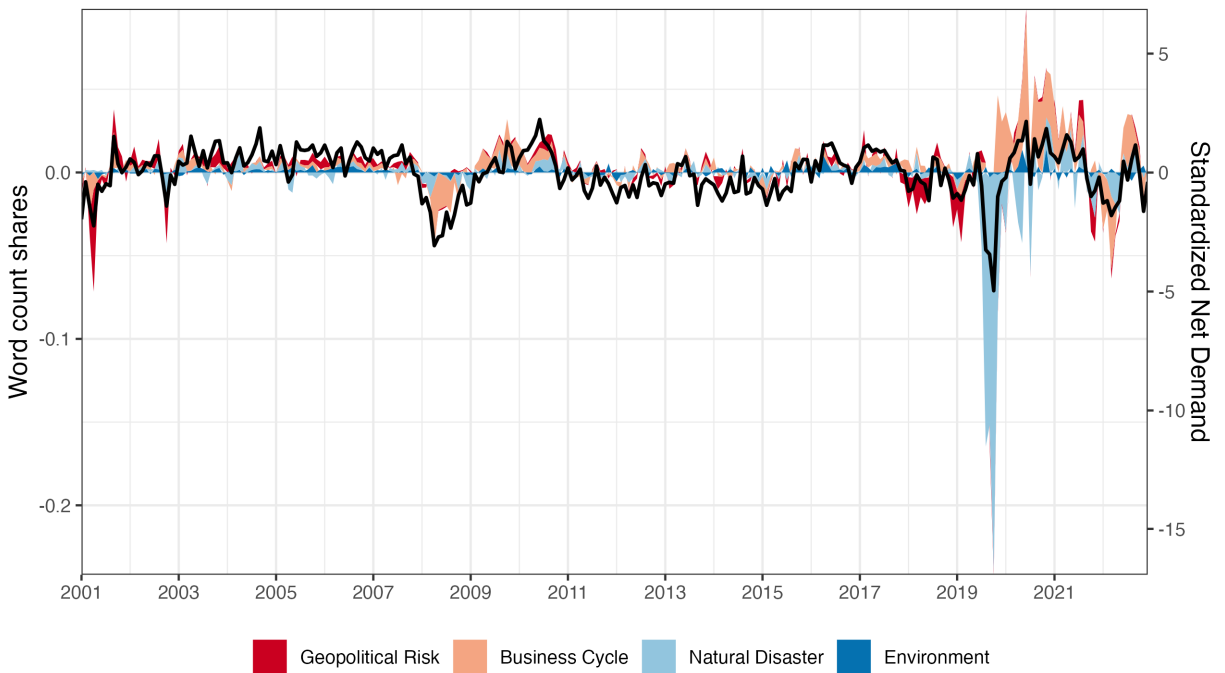
Note: This figure plots the Climate Change Theme between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) climate change events. [A] The U.S. Environmental Protection Agency requires fuel companies to blend slightly more renewable fuels into the nation's fuel supply; major oil spill on the Keystone pipeline in South Dakota sparked environmental debates; China's crackdown on genetically modified (GMO) soybean imports and its intensifying war on smog this winter will boost soymeal prices

FIGURE 13. Factor decomposition: Net supply and net demand

A. Net supply



B. Net demand



Note: This figure plots the factor decomposition in commodity articles narrating supply and demand events between 2001 and 2023. The factors are: [A] Geopolitical Risk, [B] Business Cycle, [C] Natural Disaster, [D] Environment.

TABLE 1. Major earthquakes in Latin America and change in copper net supply indices

Country	Date	Magnitude	Total Affected	Supply Decrease (Copper, All)	Supply Decrease (Copper, Nat. Dis.)
Chile	2007-11-14	7.7	25,155	1.53	5.40
Chile	2010-02-27	8.8	2,671,556	0.72	4.58
Chile	2014-04-01	8.2	513,387	1.15	4.13
Chile	2015-09-16	8.3	681,499	1.90	3.50

Note: This table reports four large seismic episodes in Chile. The second and third columns report the date and magnitude of the event. The fourth column reports the EM-DAT human impact, as the sum of the number of injured, number affected and number homeless. The fifth and sixth columns report daily the values of the copper supply decrease index and the natural disaster component of the copper supply decrease index, respectively, on the day following the event.

TABLE 2. Panel Regression of commodity futures daily returns on commodity-level supply & demand indicators

Commodity Returns (%)	Narrative Indices (Raw Series)			Narrative Indices (AR20 residuals)		
	(1)	(2)	(3)	(4)	(5)	(6)
	Net Supply	0.087*** (5.71)	0.079*** (6.19)	0.080*** (6.21)	0.088*** (5.21)	0.082*** (5.74)
Net Demand	0.153*** (8.53)	0.128*** (8.13)	0.130*** (8.39)	0.154*** (8.24)	0.129*** (7.94)	0.130*** (8.15)
Observations	106880	111940	111920	106480	111540	111540
Controls:						
Commodity Fixed Effects	Y	Y	Y	Y	Y	Y
VIX, FF, USD	Y			Y		
Time Fixed Effects		Y	Y		Y	Y
Lagged Return			Y			Y

TABLE 3. Placebo panel regression of commodity futures daily returns on commodity-level supply & demand indicators

	(1)	(2)	(3)	(4)	(5)	(6)
Net Demand (Standardized)	0.0180*** (0.0029)		0.0180*** (0.0029)	0.0180*** (0.0029)	0.0066* (0.0030)	-0.0063 (0.0033)
Net Supply (Standardized)		0.0047 (0.0032)	0.0046 (0.0032)	0.0046 (0.0032)	-0.0010 (0.0030)	-0.0046 (0.0027)
Within R^2	0.000	0.000	0.000	0.000	0.000	0.000
N	100,746	100,746	100,746	100,746	100,746	100,746
Time FE	No	No	No	No	Weekly	Daily
Commodity FE	No	No	No	Yes	Yes	Yes

Note: This table reproduces columns (4) to (6) of Table 2 but this time we estimate Equation (4) for the daily GSCI component returns of commodity i , but replace the commodity-specific supply and demand indicators with those constructed for commodity $j \neq i$.

APPENDIX A. COMMODITY INDICATORS: ADDITIONAL SENSITIVITY CHECKS

APPENDIX B. DATA SOURCES AND VARIABLE DEFINITIONS

APPENDIX C. DATA CONSTRUCTION, DICTIONARIES, AND AUDIT PROCESS

C.1. **Sourcing.** We obtain news articles from Reuters via the Factiva Snapshots API over the period May 2000 to May 2023. We apply the following filters and restrictions to construct our news sample.

We restrict attention to Reuters news outlets by limiting the sample to articles with Factiva `restrictor_codes` set to `lba` or `eline`. We further restrict to English-language content by imposing `language_code = 'en'`. We focus on commodity-specific news. We identify such articles using Factiva `subject_codes` that correspond to individual commodities, and retain only those articles whose subject codes include at least one of the following:

`mcroi, mhtgol, mnatgs, mcrntg, malum, mcopp, mnickl, mlead, mzinc, m1421, msilv, mwheat, mcorn, mcoeff, msugar, mcocoa, mcott, mhogs, mgrfds`

We exclude articles whose primary Factiva `subject_codes` correspond to the following generic or non-article content types:

- `nnam`: republished news;
- `nrgn`: obituaries, sports, and calendar-type content;
- `ncal`: calendar of events;
- `nlist`: headline listings.

These exclusions remove material that is not original news reporting or that does not contain substantive textual content.

C.2. **Cleaning.** We apply a sequence of standard text-cleaning steps. Each article is first split into sentences and then into individual tokens. Tokens are lowercased, stripped of non-alphanumeric characters (with digits preserved), and standard stopwords are removed. We exclude articles that correspond to “Take a Look” headline-summary pieces. Within the cleaned token lists, we additionally remove placeholder codes such as `[NUC]`, `[ENV]`, `[RNW]`, and `[CO2]`, which would otherwise distort dictionary-based word counts.

To address near-duplicate reporting, we group articles into rolling five-day windows and compare their sentence-level content. For each pair of articles within a window, we compute the proportion of identical sentences. Articles sharing more than 70% identical sentences are flagged as near-duplicates and removed from the final sample. This procedure yields a final dataset of 1,059,669 articles after removing 128,973 near-duplicate observations.

C.3. **Text Pre-processing.** Before implementing the text-reading algorithms, we standardize the structure of the news articles through a set of pre-processing steps. We segment each article into sentences, which serve as the basic units for detecting economically relevant patterns. Following standard practice in textual analysis, we remove stop words and punctuation from our corpus. Stop-word lists

typically contain high-frequency functional words (e.g., “the”, “is”, “at”, “which”, “and”, “on”) that carry little informational content for economic analysis. This normalization reduces stylistic variation across journalists and allows the algorithms to focus on content-bearing words that convey substantive economic information.

C.4. Dictionary Construction. Our dictionary development follows a data-driven approach. We first identify high-frequency terms in commodity news articles. Each term is evaluated individually and classified as supply-related or demand-related according to its economic significance. This computational classification is supplemented with extensive human validation checks conducted on a representative sample of news articles to ensure semantic consistency and contextual accuracy. These verification procedures establish the foundation for our supply and demand dictionaries.

When building the dictionaries for the aggregate commodity index, we include a parsimonious set of words that unambiguously refer to supply and demand. For individual commodities, we augment the core dictionaries with commodity-specific terms that capture the distinct supply and demand conditions of each market. These lexicons are designed to remain structurally simple to limit classification errors. Table A4 presents the supply and demand word roots used for the global commodity index and for crude oil, natural gas, and wheat, together with the standard list of increase and decrease words common across commodities. Our results are robust to using both words and lemmas.

To incorporate directional information, we employ standardized dictionaries of increase- and decrease-related words validated in the economic textual-analysis literature. We combine supply and demand terms with these directional words to form combinations that indicate both the side of the market (supply or demand) and the direction of change (expansion or contraction).

C.5. Co-occurrence Counting. Our methodology counts how often supply- or demand-related words appear near directional words within the same sentence. Specifically, for each supply or demand term, we count instances in which a directional word appears within a specified word-distance threshold on either side. These co-occurrence counts are used to construct the supply and demand indicators described in Section 2. The algorithm retains sentences that explicitly reference the focal commodity and excludes content centered on competing commodities.

C.6. Contextual Refinement. We supplement the baseline approach with rules that handle semantic negation and commodity-specific linguistic quirks. This step is essential because supply- or demand-related words combined with directional modifiers do not always match their literal meaning. Several commodity groups exhibit systematic reversals that would otherwise introduce measurement error. For agricultural commodities, positive directional terms paired with supply-related words can signal supply contractions—for example, “high temperature” denotes adverse growing conditions and reduced supply rather than expansion. In energy markets, “long maintenance” similarly indicates

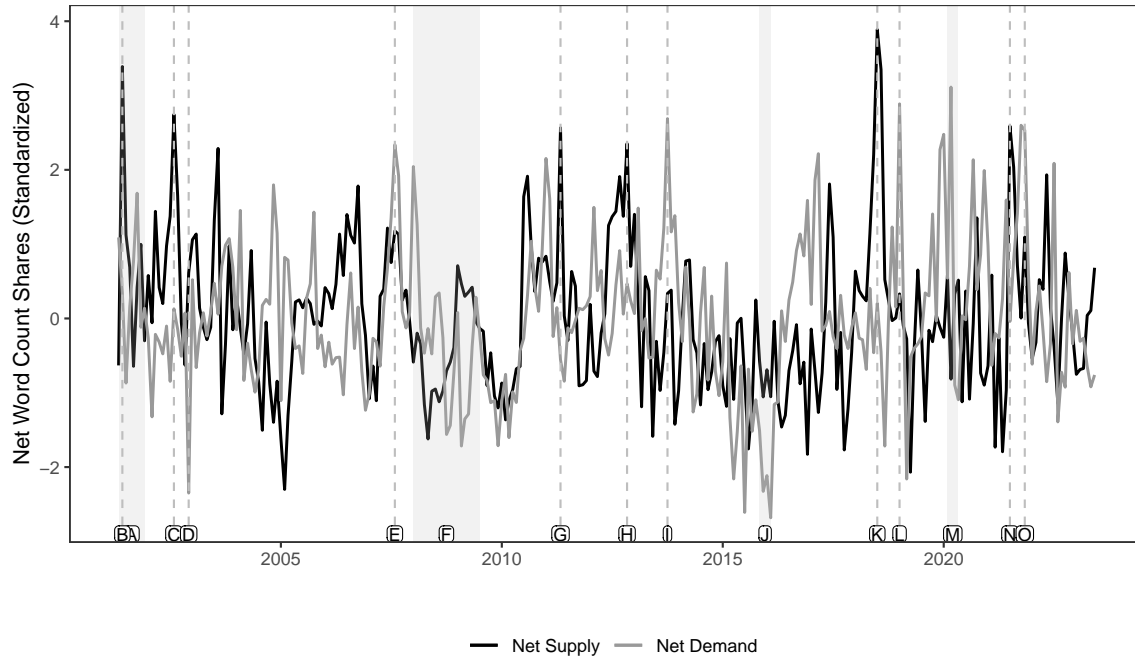
supply disruption rather than increase. Our algorithm incorporates commodity-specific exception rules to correct these reversals and maintain consistent interpretation across commodity classes.

C.7. Audit Process Based on Human Readings. We first randomly sample (without replacement) a fixed number of articles per month. The idea is to insulate the sample from changes in the number of articles that result from increasing trends in coverage, as well as from the differential effect of supply and demand developments. We randomize the order in which auditors review and code assigned articles to ensure that learning effects do not bias the assessment of differences over time. We assign the same sample to at least two auditors, accepting a smaller coverage in exchange for greater accuracy and confidence in our human-coded benchmarks.

When it comes to the information we retain, we first assess whether the article refers to and provides details regarding the commodity in question; we code the different outcomes accordingly [if yes, 2; if only to a limited extent, 1; else, 0]. Next, we verify whether the article is about supply and demand market developments [if yes, `supply_demand = 1`, else `supply_demand = 0`], and if so, we subsequently identify these developments as supply increase, supply decrease, demand increase, and or demand decrease. We also count the number of supply, demand, increase, decrease words as well as combinations and examine whether the supply and demand words actually imply the right market developments (e.g. increased number of planted crops (supply increase) vs pasta plants boosted orders (demand increase)). At this stage we additionally examine if any relevant demand or supply words are missing from the current dictionaries.


The auditing results also serve as a performance benchmark for our computer-based indices. The correlation between our human- and computer-built indices is high; it ranges between 0.64 (demand decrease) and 0.85 (demand increase) in quarterly data. In particular, to dynamically assess the time-series performance implied by our machine-generated process, we examine the co-evolution of the human- and computer-generated indices. Figure A5 to Figure A8 plot the quarterly comparison for each of the supply and demand components along with the net index performance of the two methods.

FIGURE A1. Wheat index: Net supply and net demand



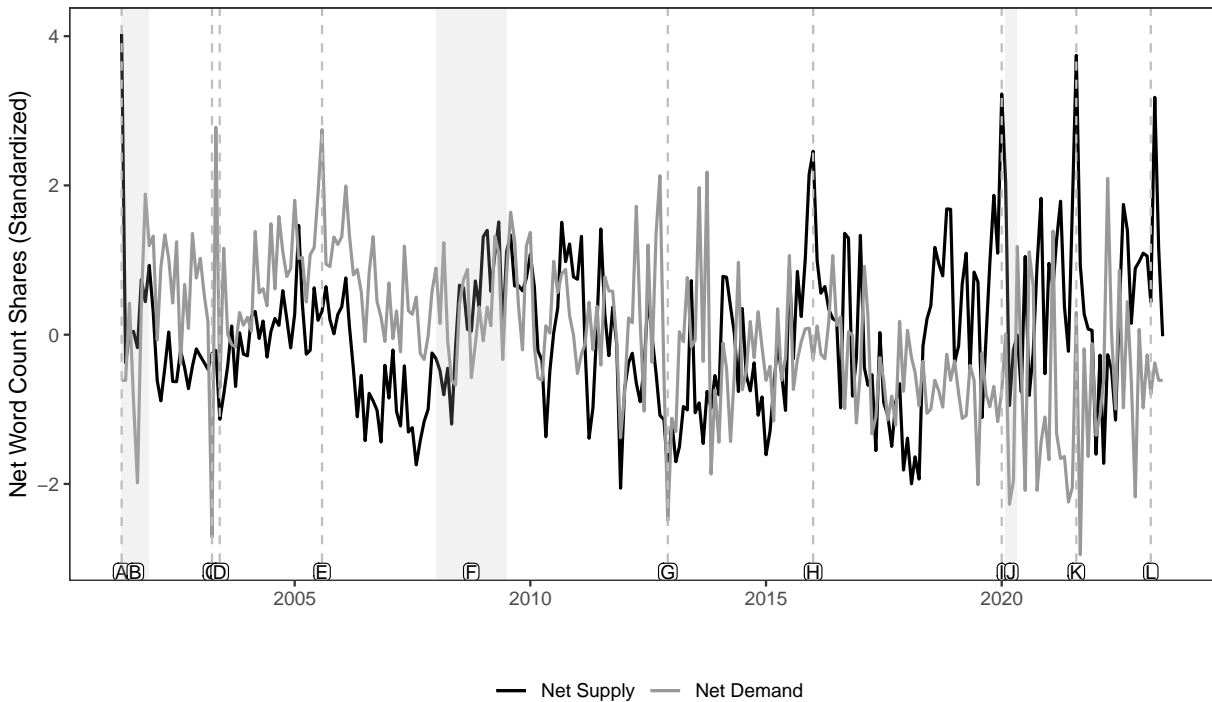
Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the wheat sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] US wheat production declined by 10% from previous year, [C] Severe drought in Canada cut spring wheat production by more than 25%. Worrisome dry weather in Australia and U.S. Northern Plains area threatened yields' prospects, [D] Decline in global wheat import putting pressure on price, [E] Strong global growth drives prices to record high, Ukraine imposes export restriction, [F] Global Financial Crisis, [G] Dry weather stressed the hard red winter wheat crop in the Plains, while wet weather slowed spring wheat seeding in the northern Plains. Dry weather in France led to concerns about harm to the wheat crop, [H] Expectations of smaller wheat crops in Argentina and Australia, dry conditions in the US Midwest, and low supply in the Black Sea region sparked worries about world supply, [I] Prospect of strong demand from China, [J] Sluggish demand for U.S. supplies, partly due to appreciating USD, China's policy shift that favoured other grains, tensions with traders over Egyptian policy on grain fungus ergot, [K] Widespread severe drought developed over Northern and Central Europe in summer 2018 led to a significant drop in harvests, [L] Resuming trade talks between China and the U.S. boosted hopes for higher wheat demand, [M] Covid-19 Recession, global shutdown induced panic buying of essential goods, [N] Adverse hot weather in North America, torrential rain in western Europe and reduced expectations for Russia's harvest created concern about global supply, [O] Ongoing supply chain disruption and robust demand after global reopening drives wheat to 9-year high.

FIGURE A2. Corn index: Net supply and net demand




Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the corn sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] Weak US export demand and abundant supply put pressure on corn price, [C] Persistent weakness in demand prompted reduction in US corn export forecast, [D] Concern over bird flu in East Asia depress demand for corn feed, [E] Global Financial Crisis, [F] Worst drought in more than 50 years in the US severely cut corn and soybeans productions, [G] Record US corn harvest, third-largest US soybean harvest, [H] Record US grain harvest in US, Brazil, and South Africa, [I] Strong global demand for U.S. corn from traditional buyers, including Mexico, Japan and South Korea, [J] El Niño induced drought and hit corn production in Southern Africa, [K] Strong export demand for U.S. corn, [L] Covid-19 Recession; global shutdown induced panic buying of essential goods, [M] China's recovery from swine fever boosted demand for corn, [N] Lower U.S. and EU corn crops amid dry weather

FIGURE A3. Sugar index: Net supply and net demand



Note: This figure plots the standardized net supply and demand indicators for sugar, for the period between 2001 and 2023. The bars map a number of important developments in the sugar sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] EU permanent production quota reduction and adverse weather in key producers (e.g. Brazil, China) reduced sugar production, [B] US Recession and 9/11, [C] Shrinking demand partly triggered by the U.S.-led war in Iraq, as Middle East countries had already stockpiled sugar in recent months over war fears, [D] Notable uptick in sugar futures driven by trade (Libya, Taiwan, Russia) and speculative buying, buoyed the market from a recent low, [E] Robust consumer buying from Russia, Pakistan and the Middle East, [F] Global Financial Crisis, [G] Sluggish demand due to winter season and rising supplies amid ongoing cane crushing in India [H] Lower-than-expected sugar output worldwide, with drought affecting India, slow start in Central America, disappointing yields in Thailand, and lowest EU beet crop in 40 years, [I] Poor growing conditions in US and India resulted in poor sugar harvest, [J] Onset of Covid-19 Pandemic, US Recession, [K] Lower production in Brazil due to the earlier drought and recent waves of frost, [L] Concerns on lower output from droughts in key producing regions (Thailand, India)

FIGURE A4. Copper index: Net supply and net demand



Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the copper sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] producers in North and South America and Asia announced production cuts to battle steep mining and processing costs and historically low metal prices, [C] Strong demand from China and good U.S. manufacturing-related news indicated demand for copper would remain sustained, [D] Global Financial Crisis, [E] Well-supplied market amid higher output from mining groups (such as BHP Billiton) and China’s domestic production, [F] Higher copper output in China (jumped by 21 percent y.o.y) and Chile (copper mine’s output surged back after tough 2012), [G] Concerns about rising supply as Newmont Mining Corp raised its full-year copper production forecast, [H] The global copper market is expected to move into surplus as boom time-driven new mine supply ramps up mostly from Peru, [I] Concerns about supplies after disruptions in top producer Chile (strikes), Indonesia (dispute with the government over mining rights) and Peru (strikes), [J] Onset of Covid-19 Pandemic, US Recession, [K] Protests hitting top mines in Peru, halting 20% of national copper output, coupled with a drop in output from top producer Chile and the potential for more sanctions on Russia raised supply shortages, [L] Signs of improving copper demand from China amid the easing of COVID-19 restrictions

TABLE A1. Keywords for the global commodity supply and demand index

Energy	Industrial metals	Precious metals	Agricultural Grains	Softs	Livestock
Crude Oil	Aluminum	Gold	Corn	Cocoa	Cattle
Heating Oil	Copper	Silver	Soybeans	Sugar	Lean Hogs
Natural Gas	Nickel		Wheat	Coffee	
Gasoil	Zinc			Cotton	
Gasoline					

Note: This table lists the constituent commodities of the S&P GSCI composite index. These items serve as the keyword set used to construct the Global Commodity Supply and Demand Index series.

TABLE A2. Regression of S&P GSCI commodity-specific daily returns on net supply and demand indices

Panel A: Energy

	Gas	Gasoline	Oil
Net Demand (Standardized)	0.0712*** (0.0134)	0.0533*** (0.0134)	0.0846*** (0.0136)
Net Supply (Standardized)	0.0129 (0.0134)	0.0281* (0.0133)	0.1215*** (0.0134)
Within R^2	0.006	0.010	0.026
N	5,597	5,597	5,597
Controls	Yes	Yes	Yes

Panel B: Industrial metals

	Aluminium	Copper	Nickel	Zinc
Net Demand (Standardized)	0.0269* (0.0134)	0.1509*** (0.0132)	0.0575*** (0.0134)	0.0321* (0.0134)
Net Supply (Standardized)	0.0511*** (0.0134)	0.0442*** (0.0132)	0.0250 (0.0134)	0.0377** (0.0134)
Within R^2	0.009	0.030	0.005	0.007
N	5,597	5,597	5,597	5,597
Controls	Yes	Yes	Yes	Yes

Panel C: Precious metals

	Gold	Silver
Net Demand (Standardized)	0.0687*** (0.0133)	0.0583*** (0.0134)
Net Supply (Standardized)	0.0029 (0.0134)	-0.0007 (0.0134)
Within R^2	0.005	0.004
N	5,597	5,597
Controls	Yes	Yes

Panel D: Grains

	Corn	Soybean	Wheat
Net Demand (Standardized)	0.0984*** (0.0133)	0.1283*** (0.0133)	0.1016*** (0.0132)
Net Supply (Standardized)	0.0684*** (0.0133)	0.0602*** (0.0133)	0.1326*** (0.0132)
Within R^2	0.014	0.021	0.029
N	5,597	5,597	5,597
Controls	Yes	Yes	Yes

Panel E: Softs

	Cocoa	Coffee	Cotton	Sugar
Net Demand (Standardized)	0.0938*** (0.0134)	0.0923*** (0.0134)	0.1302*** (0.0132)	0.0922*** (0.0133)
Net Supply (Standardized)	0.0534*** (0.0133)	0.0513*** (0.0134)	0.0726*** (0.0132)	0.0781*** (0.0133)
Within R^2	0.012	0.011	0.024	0.014
N	5,597	5,597	5,597	5,597
Controls	Yes	Yes	Yes	Yes

Panel F: Livestock

	Cattle	Hog
Net Demand (Standardized)	0.2151*** (0.0131)	0.1622*** (0.0132)
Net Supply (Standardized)	0.0469*** (0.0130)	0.0504*** (0.0132)
Within R^2	0.051	0.030
N	5,597	5,597
Controls	Yes	Yes

TABLE A3. Data Sources and Variable Definitions

Variable	Description	Source
<i>Energy</i>		
WTI Crude Oil	S&P GSCI Crude Oil Spot (GSCLSPT)	Datastream
Brent Crude Oil	S&P GSCI Brent Crude Spot (GSBRSPT)	Datastream
Heating Oil	S&P GSCI Heating Oil Spot (GSHOSPT)	Datastream
Natural Gas	S&P GSCI Natural Gas Spot (GSNGSPT)	Datastream
<i>Metals</i>		
Aluminum	S&P GSCI Aluminum Spot (GSIASPT)	Datastream
Copper	S&P GSCI Copper Spot (GSICSPT)	Datastream
Nickel	S&P GSCI Nickel Spot (GSIKSPT)	Datastream
Zinc	S&P GSCI Zinc Spot (GSIZSPT)	Datastream
Gold	S&P GSCI Gold Spot (GSGCSPT)	Datastream
Silver	S&P GSCI Silver Spot (GSSISPT)	Datastream
<i>Agriculture</i>		
Corn	S&P GSCI Corn Spot (GSCNSPT)	Datastream
Soybeans	S&P GSCI Soybeans Spot (GSSOSPT)	Datastream
Wheat (CBOT)	S&P GSCI Wheat (CBOT) Spot (GSWHSPT)	Datastream
Cocoa	S&P GSCI Cocoa Spot (GSCCSPT)	Datastream
Coffee	S&P GSCI Coffee Spot (GSKCSPT)	Datastream
Sugar	S&P GSCI Sugar Spot (GSSBSPT)	Datastream
Cotton	S&P GSCI Cotton Spot (GSCTSPT)	Datastream
<i>Livestock</i>		
Cattle	S&P GSCI Live Cattle Spot (GSLCSPT)	Datastream
Lean Hogs	S&P GSCI Lean Hogs Spot (GSLHSPT)	Datastream
<i>Control variables</i>		
USD Index (broad)	Nominal Broad U.S. Dollar Index (Goods Only) (DTWEXB)	FRED
USD Index (broad)	Nominal Broad U.S. Dollar Index (DTWEXBGS)	FRED
Fed funds rate	Federal Funds Effective Rate (DFF)	FRED
VIX Index	CBOE Volatility Index: VIX (VIXCLS)	FRED

Variable	Description	Source
<i>Macro variables</i>		
World oil production	World oil production (EIA1955)	Datastream
World industrial production	Industrial production of OECD+6 (Brazil, China, India, Indonesia, Russia, South Africa), Baumeister and Hamilton (2019) (OECD+6IP)	Baumeister's web-page
OECD crude oil inventories	OECD crude oil inventories from petroleum stock data (EIA1976, EIA1533, EIA1541), as in Kilian and Murphy (2014)	Datastream
US industrial production	US industrial production index (INDPRO)	FRED
US CPI	US CPI for all urban consumers: all items (CPIAUCSL)	FRED
US PPI	Producer Price Index: all commodities (PPI-ACO)	FRED
US unemployment rate	Civilian unemployment rate (UNRATE)	FRED
US real GDP	Real Gross Domestic Product (GDPC1)	FRED
US real Investment	Real Gross Private Domestic Investment (GPDIC1)	FRED

TABLE A4. Dictionary words: Supply and demand by commodity, and increase/decrease word stems

Panel A: Global Commodity Index	
<i>Supply</i>	suppl*, produc*, output
<i>Demand</i>	demand*, consum*, buy*, purchas*
Panel B: Crude Oil	
<i>Supply</i>	suppl*, produc*, output, discovery, glut*, reserv*, surplus*, rig*
<i>Demand</i>	demand*, consum*, buy*, util*, drain*, deplet*, refin*, purchas*
Panel C: Natural Gas	
<i>Supply</i>	suppl*, produc*, output, rig*, drill*
<i>Demand</i>	demand*, consum*, buy*, purchas*
Panel D: Wheat	
<i>Supply</i>	suppl*, produc*, output, crop*, planting*, farm*, harvest*
<i>Demand</i>	demand*, consum*, buy*, purchas*
Panel E: Increase words	
	accru*, climb*, improve*, rais*, soar*, accumulat*, elevat*, increas*, rall*, spik*, add*, enlarg*, inflat*, reach*, spring*, advanc*, escalat*, jackup*, rebuil*, spurt*, augment*, expand*, jump*, recoup*, strengthen*, bolster*, firm*, leap*, recover*, surg, boom*, gain*, lift*, regain*, surpass*, boost*, grow*, perk*, resurgenc*, swell*, buil*, height*, pickup, reviv*, up*, bullish*, high*, pop*, ris*, buoyant*, hik*, propell*, *rocket*
Panel F: Decrease words	
	abat*, dent*, evaporat*, pullback*, slump*, bearish*, depress*, fad*, reced*, small*, below, deteriorat*, fall*, reduc*, squeez*, collaps*, diminish*, falter*, restrict*, stumbli*, compress*, dip*, halt*, retre*, sub*, contract*, disappear*, landslid*, sink*, suppress*, crash*, disappoint*, less*, short*, suspend*, crimp*, disrupt*, los*, shrink*, tank*, crush*, div*, low*, shut*, tap*, curtail*, down*, meltdown*, slack*, tight*, cut*, drawdown*, nosediv*, slash*, dampe*, drop*, outage*, slid*, declin*, dwindle*, plummet*, slip*, decreas*, ebb*, plung*, slow*

FIGURE A5. Human- and computer-based increase indices




FIGURE A6. Human- and computer-based decrease indices




FIGURE A7. Human- and computer-based net indices




FIGURE A8. Human- and computer-based standardized net indices

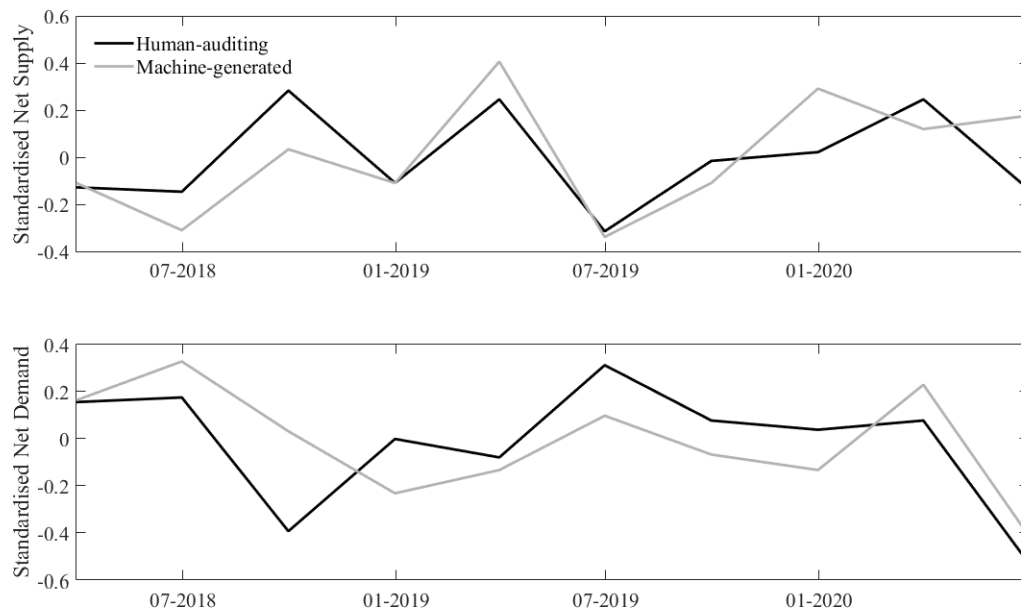
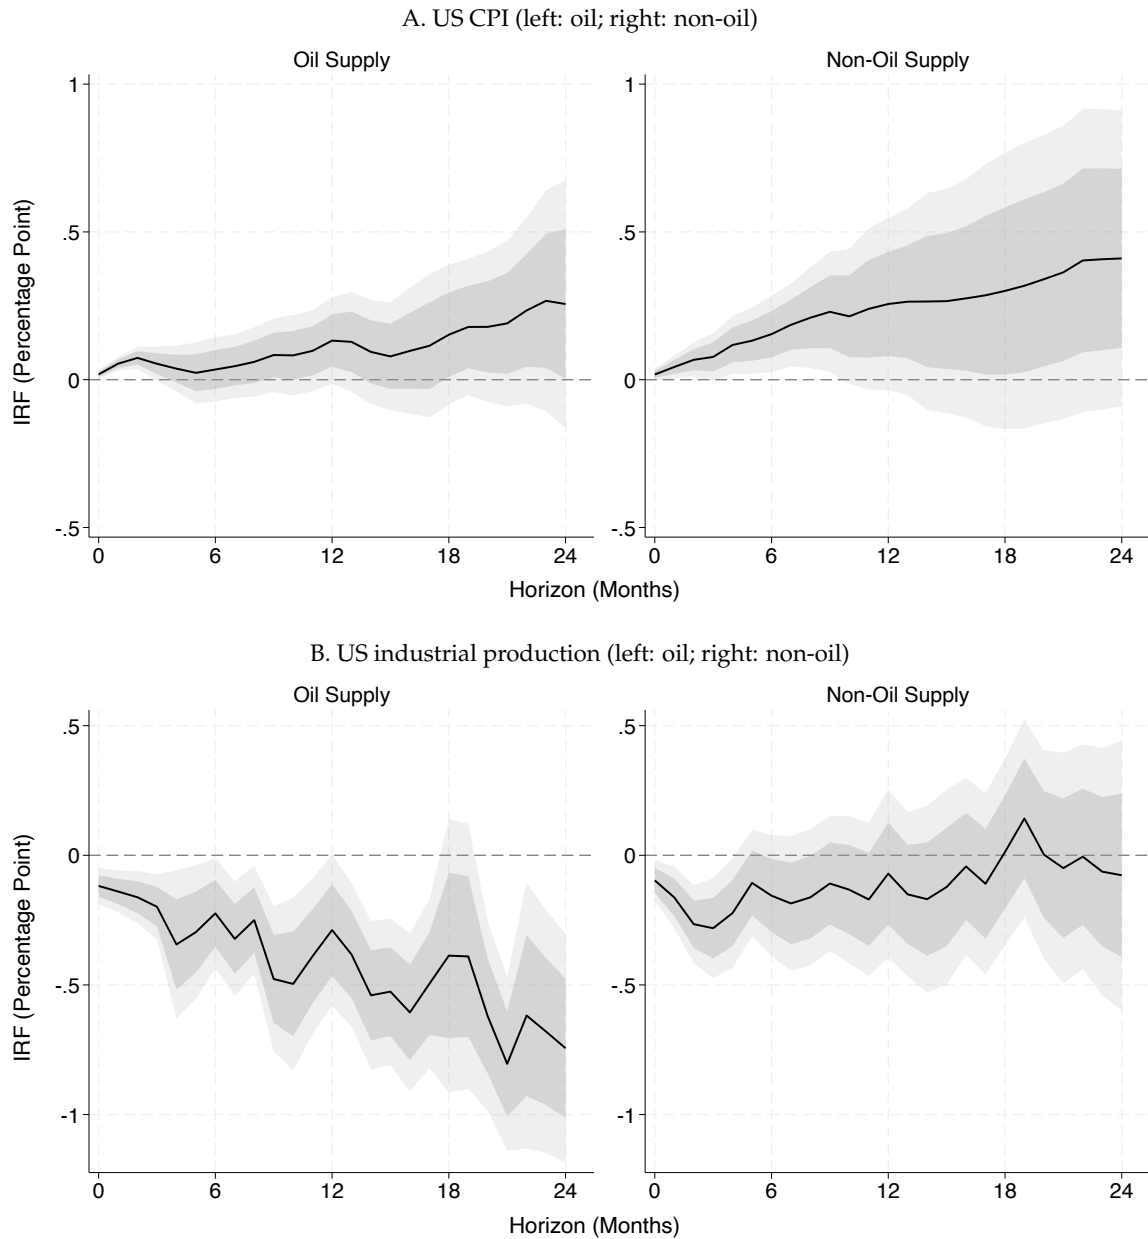
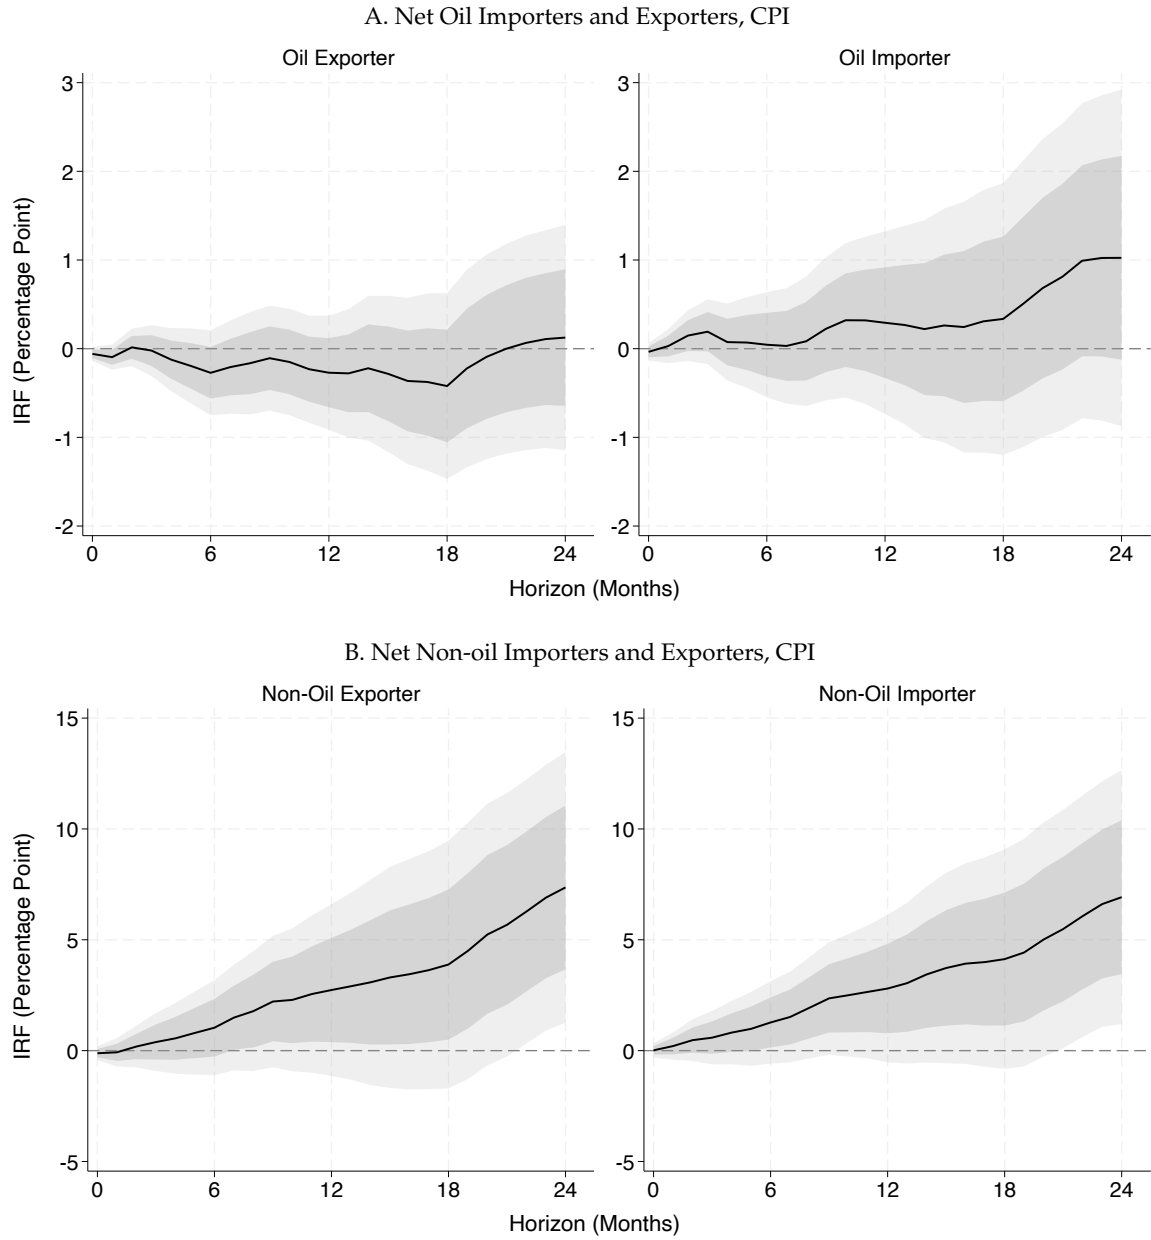


FIGURE A9. Local projections: US impulse responses to oil and non-oil supply shocks

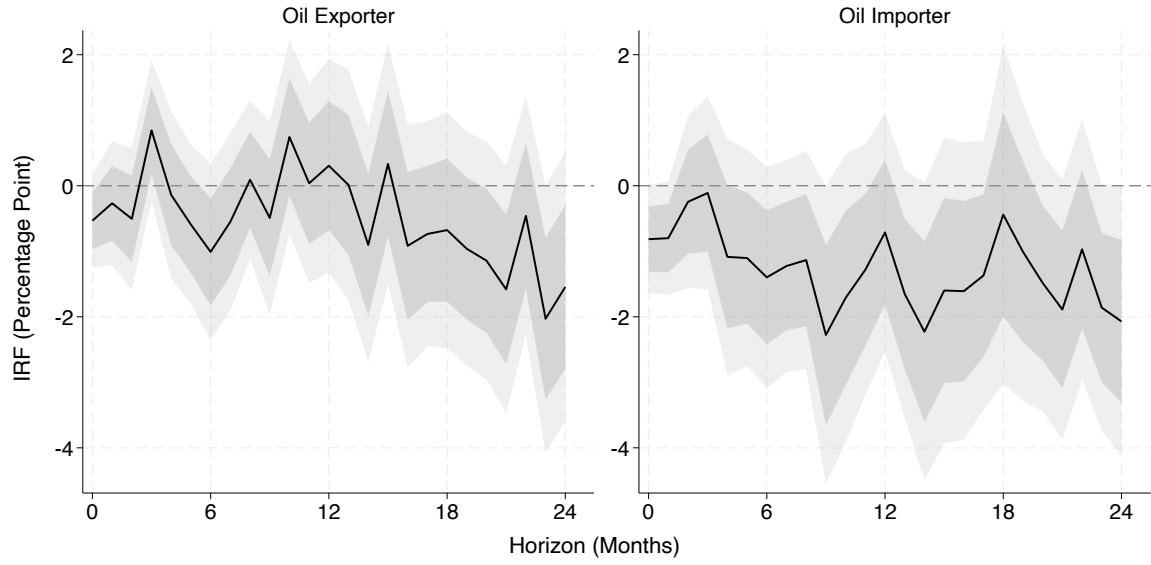


Note: This figure shows the impulse responses of US CPI and US industrial production to a one-standard-deviation supply contraction, employing the oil supply index (left) and the non-oil production-weighted supply index (right), constructed as AR(20) residuals aggregated to monthly frequency.

FIGURE A10. Local projections: Commodity Net Exporters and Commodity Net Importers impulse responses



C. Net Oil Importers and Exporters, Industrial Production



D. Net Non-oil Importers and Exporters, Industrial Production

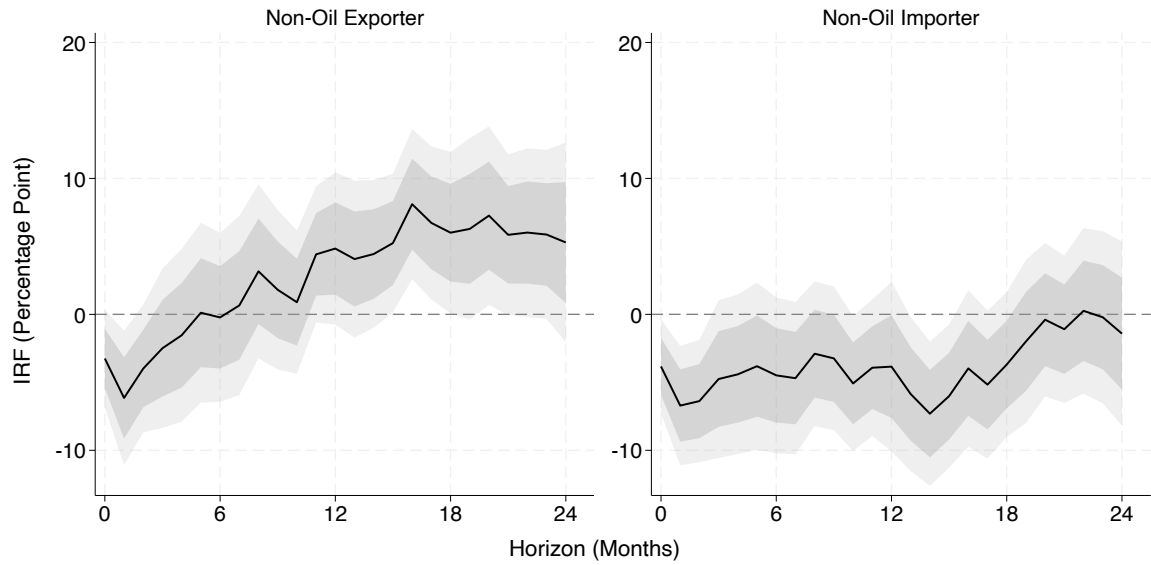
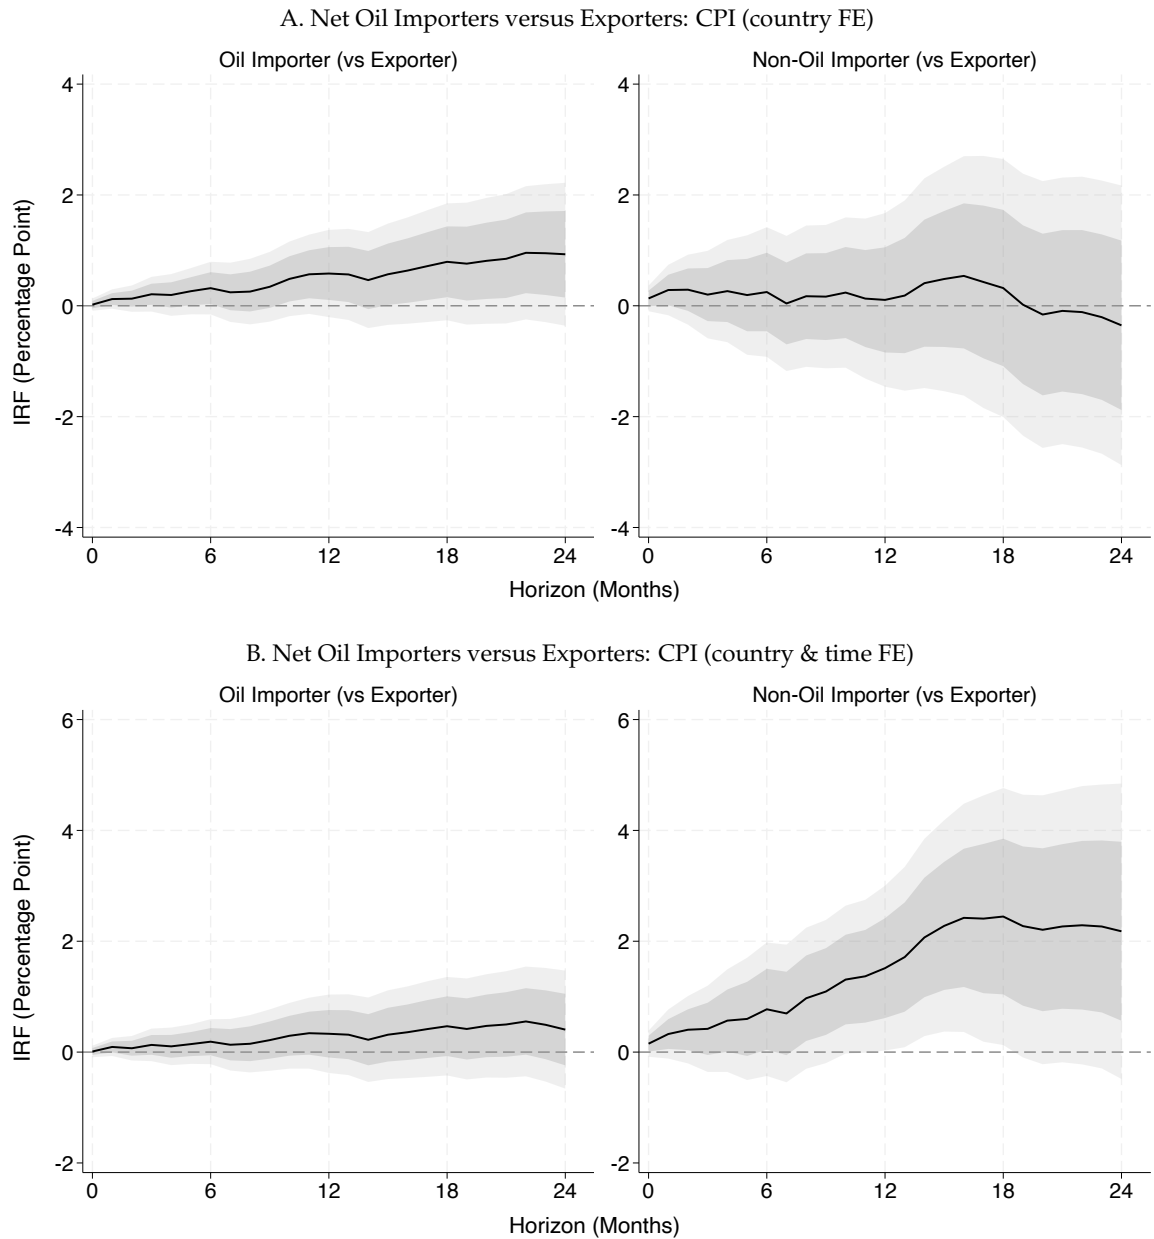


FIGURE A11. Local Projections: Importers versus Exporters CPI Responses



Note: This figure shows the differential impulse response of CPI to a one-standard-deviation commodity supply contraction, interacted with the normalized trade balance. Panel A includes country fixed effects; Panel B adds time fixed effects. The left column uses the oil supply index; the right column uses the non-oil production-weighted supply index. Driscoll-Kraay standard errors are used.