Beyond Oil: The Origins of Commodity Price Fluctuations*

Alvin Lumbanraja, Sarah Mouabbi, Evgenia Passari and Adrien Rousset Planat

September 13, 2026

Abstract

Commodity supply shocks are a plausible but elusive source of business-cycle fluctuations, and evidence has been largely confined to oil. Using textual analysis of more than 1.1 million news articles (2001–2023), we construct daily supply and demand proxies for 20 commodities spanning energy, metals, agriculture, and livestock. Our central finding is that non-oil supply disruptions move inflation and industrial production at least as strongly as oil shocks. Aggregate local projections across 60 countries establish this, and a country-commodity-month panel confirms it from within-commodity variation in exposure alone, with the effects concentrated among net importers. Commodity supply shocks matter well beyond oil.

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

Keywords: commodities, textual analysis, news, demand, supply, business cycles,

*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, Pol Antràs, 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, Juan Pablo Nicolini, Luca Onorante, Elias Papaioannou, Dimitris Papanikolaou, Gert Peersman, Ricardo Reis, Hélène Rey, Philip Sauré, Martin Schneider, Daniele Siena, Vania Stavrakeva, Alan Taylor and seminar and conference participants at numerous venues for useful comments including the conference participants of the SED 2026, ASSA 2024, ESSIM CEPR 2022, and BSE Summer Forum 2022 for comments and discussions. 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.

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 and prices. However, despite their theoretical importance, empirical progress has been largely confined to oil (Hamilton, 1983, 2003).[1] A central reason is that commodity prices respond to both supply and demand, and separating the two has proved tractable only in that market. Modern economies, however, depend on a far broader set of commodity inputs, such as industrial metals, agricultural products, and various 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 methodology to measure both supply and demand across the entire commodity market. Using an automated narrative approach applied to more than 1.1 million news articles, we construct daily proxies for supply and demand pressure that span 20 commodities—including 19 previously undocumented markets across energy, metals, agriculture, and livestock. We further classify the forces behind these movements into four drivers—business-cycle developments, geopolitical risk, natural disasters, and climate change—and find that natural disasters and geopolitical risk dominate supply disruptions while the business cycle shapes demand. The daily frequency is central to identification: it lets us purge predictable and cross-market components so the series isolate unanticipated, commodity-specific disturbances. This breadth lets us take up two questions the literature has struggled to answer: whether commodity supply shocks beyond oil matter for the macroeconomy, and how their transmission depends on a country's trade position.

We disentangle supply from demand in commodity prices using audited dictionaries that capture the language journalists use to describe market forces, counting the co-occurrence of supply- and demand-related terms with directional modifiers and normalizing by article volume; an extensive human audit of randomly sampled articles yields quarterly correlations with our machine-generated indices of 85 percent. Three further pieces of evidence indicate that the series capture genuine, unanticipated shocks. First, supply contractions and demand expansions both raise commodity spot prices—at the aggregate level and commodity by commodity—and the response is concentrated on the day the news breaks, with negligible anticipation, so the indicators behave as clean shocks at the monthly frequency of our macroeconomic analysis. Second, for oil—the one commodity for which existing measures of supply shocks are available—our supply measure reproduces the impulse responses of

[1] Additional 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).

Baumeister and Hamilton (2019) and Känzig (2021). Third, the indices track the major episodes one would expect—demand collapses in the 2008–2009 recession and the COVID-19 pandemic, and supply disruptions around OPEC decisions, hurricanes, and droughts—and the theory-consistent signatures emerge, with supply contractions raising prices while depressing output and demand expansions raising both.

Our first finding is that commodity supply shocks are a quantitatively important source of macroeconomic fluctuations—and that this is not driven by oil alone. Estimating the average response of inflation and industrial production to an aggregate supply shock, separately for oil and non-oil commodities, we find that non-oil disturbances—originating in metals, grains, softs, and livestock—affect inflation and industrial production at least as strongly as oil, and often more. The literature's near-exclusive focus on petroleum has thus understated the macroeconomic importance of commodity supply shocks as a class.

These aggregate effects are identified from time-series variation alone and need not reflect causal transmission. To isolate it, we exploit the full granularity of our data—a panel of 60 countries observed at the country-commodity-month level. Interacting each commodity's supply shock with countries' commodity-specific trade positions and absorbing country-by-commodity and commodity-by-month fixed effects, we identify the differential response to a given commodity shock from within-commodity variation in exposure alone—ruling out confounding from countries' broader macroeconomic characteristics. Transmission is sharply heterogeneous: net importers suffer larger, more persistent output contractions and stronger inflation than net exporters, who are partially insulated on both margins.[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 disruptions associated with COVID-19 and geopolitical tensions have revived interest in supply-side explanations: Bernanke and Blanchard (2025) emphasize that supply chain disruptions were a key driver of pandemic-era inflation. Whether COVID was an exception or part of a broader pattern has been unclear. 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.,

[2]The stronger inflation among importers is consistent with the commodity-currency channel—an adverse terms-of-trade shift depreciates the importer's currency, and the higher domestic-currency price of imports passes through to consumer prices (Campa and Goldberg, 2005)—dominating the demand channel that pushes the other way for exporters (Cashin, Cespedes, and Sahay, 2004; Drechsel, McLeay, and Tenreyro, 2019; Chen, Turnovsky, and Zivot, 2014).

2013; Alexopoulos and Cohen, 2015; Baker, Bloom, and Davis, 2016; Allcott and Gentzkow, 2017; Hassan 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; Castelnuovo, 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 1 introduces our methodology for developing supply- and demand-side indicators in commodity markets and decomposes them into the fundamental factors that drive commodity price fluctuations. In Section 2 we perform a series of validation exercises for our indicators. Section 3 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 4 concludes.

1 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 applies across all major commodity groups—energy, metals, agriculture, and livestock—and clearly separates supply from demand: the resulting indices peak during major episodes such as the post-crisis collapse in commodity markets, market-specific disruptions, and the COVID-19 shock, and are granular enough to pinpoint the fundamental forces

shaping commodity dynamics—business-cycle conditions, geopolitical risks, natural disasters, and climate change.

The strategy rests on two assumptions: news outlets report economically significant developments with minimal delay, and the magnitude of a shock scales with the breadth and intensity of its coverage. Under these assumptions, we construct the indices in two stages—first extracting directional supply- and demand-side signals at the sentence level, then refining them with contextual algorithms—and aggregate the resulting counts into net supply and net demand indicators that are comparable across markets and time frequencies. The subsections below describe each step in turn; the full text-processing pipeline is detailed in Appendix 5.

1.1 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 5 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.

1.2 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 5.4). 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 make this concrete, a sentence reporting that "copper *output slumped* after a strike halted mining at a major Chilean site" pairs a supply term (*output*) with a decrease qualifier (*slumped*) within the proximity window, and is therefore recorded as a single supply-decrease instance for copper.

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 5.

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 2.2) 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.

1.3 Supply and Demand Indicators

We construct indicators of supply- and demand-side disturbances, both at the aggregate level and for individual commodities. The aggregate index draws on all articles referencing any commodity in the Standard & Poor's Goldman Sachs Commodity Index (GSCI) spot price series (excluding lead)—a universe of major physical commodities with active, liquid futures markets. Table A.1 lists the 20 constituents, which span energy, metals, agriculture, and livestock.

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:

$$S^{+}_{i,t} \equiv \frac{\text{\# of Instances Indicating Commodity Supply Increase}_{i,t}}{\text{\# of Retained Articles}_{t}}. \tag{1}$$

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:

$$S^+_t \equiv \frac{\text{\# of Instances Indicating Commodity Supply Increase}_t}{\text{\# of Retained Articles}_t}. \tag{2}$$

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

$$S_t \equiv S^-_t - S^+_t \qquad \text{and} \quad D_t \equiv D^+_t - D^-_t. \tag{3}$$

We construct analogous indices for each commodity market i.

Figure 1 displays the standardized monthly net supply and net demand indicators. 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 net supply and net demand indicators broken down by major commodity category. 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.

To illustrate how the commodity-level indicators behave, we next examine five representative markets in turn—crude oil, wheat, corn, sugar, and copper.

1.3.1 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 cyclicality, 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 A.1 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.

1.3.2 Agricultural Commodities: Wheat, Corn and Sugar

Appendix Figure A.2 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.

Appendix Figure A.3 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.

As with wheat, supply-side developments in corn markets are predominantly weather-driven, with differences in timing reflecting the geography of production.

Appendix Figure A.4 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.

1.3.3 Industrial Metals: Copper

Appendix Figure A.5 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 cyclicality, 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.

1.4 Narrative Drivers of Supply and Demand

Different supply and demand shocks carry distinct implications for macroeconomic outcomes, financial markets, and price stability (e.g., Caballero, Farhi, and Gourinchas, 2008): natural disasters, for example, typically affect inflation over the short to medium term, whereas environmental regulation can alter relative prices more persistently through expectations. To distinguish these effects, we construct a separate classification of the events that drive commodity markets. Drawing on the risk factors highlighted in the IMF's *Commodity Special Feature* and the World Bank's *Commodity Markets Outlook*, we group them into four broad drivers: business-cycle developments, geopolitical risk, natural disasters, and climate change. To identify the language associated with each driver, we build dictionaries from authoritative sources, following Engle et al. (2020) for settings in which predefined dictionaries are unavailable: the complete record of NBER Business Cycle Dating Committee announcements (1979–2020) for the business cycle, the press statements of the President of the United Nations Security Council for geopolitical risk, the EM-DAT International Disaster Database for natural disasters, and the same climate-change texts used by Engle et al. (2020).

Figure A.6 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 A.7 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 A.8 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 A.9 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 3, 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.

2 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 show that our indicators identify the price impact of supply and demand shocks at high frequency: their signed comovement with commodity prices is theory-consistent and commodity-specific, and the price reaction is concentrated on the day the news breaks, with negligible anticipation—so the indicators behave as clean shocks at the monthly frequency used in our macroeconomic analysis. Finally, for the oil market—the one commodity for which alternative supply shocks exist—we show that our supply disturbance reproduces the impulse responses of established supply-shock measures.

2.1 Illustrative Examples

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.

2.2 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 between the two range from 0.64 for demand decreases to 0.85 for demand increases, 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 5.7 of the Appendix. Figures A.10 through A.13 display the evolution of human- and algorithm-derived indicators for supply, demand, and composite indices, visually confirming the robustness of our classification methodology.

2.3 Commodity Prices and High-Frequency Identification

We now relate our supply- and demand-side indicators to high-frequency movements in commodity prices. We first ask whether the price responses to our indicators carry the signs economic theory

predicts—a supply contraction and a demand expansion both raising prices. We then ask whether those responses represent *news*, in the sense that prices react when the indicators move but not before. The former speaks to whether the indices capture supply and demand shocks at all; the latter, to whether they capture *unanticipated* shocks suitable for causal analysis. Daily data and calendar-day fixed effects—which strip out whatever is common to all commodities on a given day—let us address both at high frequency and, ultimately, treat our indicators as clean shocks in the macroeconomic analysis of Section 3.

Throughout, we use both our indicators and their residualized "surprise" counterparts. The latter purge the predictable component of each series by removing an AR(p) fit,

$$S^r_{i,t} \equiv S_{i,t} - B(L)S_{i,t} \qquad \text{and} \quad D^r_{i,t} \equiv D_{i,t} - B(L)D_{i,t}, \tag{4}$$

estimated separately for each commodity. A conservative lag length of 20 days removes even subtle serial dependence—longer lags yield quantitatively similar results—and we standardize the residuals to unit standard deviation.

2.3.1 Price Responses

We turn to individual commodities, estimating:

$$R_{i,t} = \beta\, S_{i,t} + \gamma\, D_{i,t} + c\mathbf{Z}_{i,t} + u_{i,t}. \tag{5}$$

The coefficients of interest, β and γ, capture the effect of a one-standard-deviation reduction in supply (a rise in $S_{i,t}$) and a one-standard-deviation increase in demand ($D_{i,t}$) on commodity i's daily (log) return. 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 standard financial-market variables and calendar-day fixed effects, with fixed effects for individual commodities added in some specifications. 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. Both are significant and carry the signs theory predicts—a supply contraction and a demand expansion each raise prices—which is direct evidence that the indices are genuinely picking up supply and demand shocks. In columns (1) to (3) the key explanatory variables are the net supply and demand indicators, while in columns (4) to (6)

we replace them with their residualized counterparts. 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 one-standard-deviation decline in the residualized supply-side indicator is associated with a 7–8 basis point increase in the commodity's price on the same day; a one-standard-deviation increase in demand is associated with a 12–15 basis point increase, depending on the specification. The supply-side effects are approximately half the size of the demand-side coefficients—an asymmetry in price responsiveness to supply versus demand innovations—and these relationships hold across alternative specifications.

Because the calendar-day fixed effects absorb every shock common to all commodities on a given day, β and γ are identified from within-day, cross-commodity variation alone: the estimates cannot be driven by aggregate forces—global risk appetite, the dollar, or the business cycle—that move all commodities together, which is precisely what lets us read them as commodity-specific supply and demand responses.

The relationship holds market by market, not merely on average: re-estimating equation (5) separately for eighteen commodities (Table A.2)—spanning energy, industrial and precious metals, grains, softs, and livestock—we find that the net supply indicator enters positively in every market and is statistically significant in fourteen of the eighteen, with only natural gas, nickel, silver, and gold showing insignificant supply effects. The net demand indicator is positive and significant in every market except aluminum. That a single text-based procedure recovers theory-consistent supply and demand relationships across markets as different as crude oil and live cattle is exactly the external validity a framework spanning the full commodity spectrum requires.

2.3.2 Placebo Tests

The price responses documented above survive controls for standard financial-market variables, but we should still confirm that they reflect information specific to each individual market. We now verify that these associations are not spurious. Specifically, we re-estimate equation (5) for the daily GSCI return of commodity i but replace its own supply and demand indicators with those constructed for a different commodity $j \neq i$. This cross-commodity specification serves as a placebo: it asks whether the relationships in Table 2 reflect genuine commodity-specific information rather than common shocks or omitted variables that could generate artificial significance.

Examining Table 3, we find that none of the cross-commodity supply indicators carry statistically significant coefficients. This absence of spurious effects weighs against the concern that our baseline results reflect common shocks, measurement error, or statistical artifacts, and supports the interpretation that our commodity-specific supply indicators capture genuinely idiosyncratic information about their own markets. Demand indicators behave differently: when daily fixed effects are omitted, they remain significant but with markedly attenuated magnitudes. This attenuation is consistent with a substantial common component in demand innovations across commodities (see Chiaie, Ferrara, and Giannone, 2022)—demand shocks are more correlated across markets than supply shocks, which are largely idiosyncratic. Once the daily fixed effects absorb this shared variation, the identifying variation in the demand indicators is primarily commodity-specific and the coefficients shrink accordingly—consistent with the view of aggregate demand as a common factor linked to the global business cycle, monetary conditions, and shifts in risk appetite.

2.3.3 Unanticipated News

We now turn to the second question: whether the price response reflects news rather than developments the market had already anticipated. If our indicators merely captured expected developments, their comovement with contemporaneous returns would understate the true impact of the shock and limit their usefulness for causal analysis. To test for anticipation, we trace the dynamic price response around each news event using the residualized "surprise" indicators from (4),

$$R_{i,t+k} = \beta(k)\, S^r_{i,t} + \gamma(k)\, D^r_{i,t} + \alpha_i + \mu_t + u_{i,t}. \tag{6}$$

The dependent variable is the daily (log) return on commodity i at horizon $t+k$; we include commodity and calendar-date fixed effects and cluster by day and commodity.

Figure 4 plots $\beta(k)$ and $\gamma(k)$ for horizons $k = -20$ to $k = 5$. On impact ($k = 0$), a one-standard-deviation supply contraction raises prices by 7 basis points and a comparable demand expansion by 13 basis points—reproducing the contemporaneous panel estimates above, now traced as an event study around the news date. What is informative here is not this magnitude, which we have already established, but the *timing* of the response.

By contrast, prices barely move beforehand: in the pre-event window ($k < 0$), anticipatory movements are essentially zero beyond roughly ten days before 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. This narrow-window drift is economically small, less than 30 percent

of the contemporaneous effect, and an F-test of the null that all pre-event coefficients $\{\beta(k)\}_{k=-20}^{-1}$ are jointly zero fails to reject at conventional levels

Crucially, this pattern should pose little 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 is unlikely to attenuate or bias our estimates. Put differently, at the monthly horizon at which we estimate the local projections in Section 3, our indicators can plausibly be treated as clean contemporaneous shocks.

Taken together, the two results—theory-consistent, commodity-specific price responses and the lack of meaningful anticipation—point toward a high-frequency identification of the price impact of supply and demand shocks, and strongly support treating our indicators as shocks in the macroeconomic analysis of Section 3.

2.4 Comparison to Existing Indicators

As a final validation exercise, we benchmark our supply- and demand-side disturbances against 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) exploits high-frequency variation in oil prices around OPEC announcements to identify supply-side shocks. We ask whether our oil supply disturbance generates dynamic responses consistent with those of two benchmark measures: the oil supply shock series of Baumeister and Hamilton (2019) and the news shock about future oil supply of Känzig (2021). To keep the comparison clean, we evaluate all three shock series within a single local-projection specification, so that the shock is the only object that varies across them. We take the "surprise" component of our Net Supply Oil Index—the residual from an AR(1) regression—estimate its dynamic effects with the local projection of Jordà (2023), and run the two benchmark shocks through the same projection for comparison:

$$x_{t+k} = \beta(k) S_{oil,t} + \sum_{l=1}^{L} c_l(k)\, x_{t-l} + c(k)\mathbf{Z}_t + u_t, \qquad k = 0 \ldots 24. \tag{7}$$

This specification stays close to the benchmark oil-market studies rather than the cross-country panel of Section 3: the response variables x—the real price of oil, world oil production, OECD crude inventories, US CPI, and world and US industrial production—together with the controls $\mathbf{Z}$, the lag length L, and

the estimation sample, are consistent with those in Baumeister and Hamilton (2019) and Känzig (2021). Because each of the three shock series enters this identical design, any divergence in the resulting impulse responses reflects the shock itself rather than differences in estimation.

Appendix Figures A.15 and A.16 plot the estimated coefficients $\beta(k)$ for the two established measures (first two columns) alongside those of our oil-specific supply disturbance. Across both market and macroeconomic outcomes, our surprise tracks the benchmarks closely. For oil-market prices and quantities (Figure A.15), all three series agree that a negative supply shock raises prices and lowers quantities, with OECD inventories first accumulating and then gradually depleting. For macroeconomic outcomes (Figure A.16), our surprise reproduces the responses the established measures generate—higher US prices and lower industrial production, both globally and domestically. The emphasis here is on whether the responses agree, rather than on the size of the macroeconomic effects themselves, which we take up in Section 3.

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.

3 Commodity Supply Shocks and Macroeconomic Outcomes

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. We proceed in two steps. We first estimate the average response of inflation and industrial production to an aggregate supply shock, separately for oil and non-oil commodities, and find that non-oil disturbances matter at least as much as oil. Because these average effects are identified from time-series variation alone—and are therefore open to confounding from other global forces—we then exploit countries' cross-sectional differences in commodity exposure to identify the *differential* transmission of a given commodity shock across net importers and net exporters, holding the shock itself fixed—so that this second step, rather than the aggregate time-series evidence, is what supports a causal interpretation of the transmission.

3.1 Data Construction

Recall from Section 1 that our textual procedure yields, for each of the 20 commodities in our sample—spanning energy, metals, agriculture, and livestock—daily news-based indicators of supply and demand pressure. From these we construct the shocks that enter our projections in three stages. First, we aggregate the daily commodity-level indicators into three production-weighted indices—a *composite* index spanning all commodities, an *oil* index, and a *non-oil* index covering all commodities except oil. For market m, the daily supply-increase index is

$$S^{+}_{m,d} \equiv \sum_{i \in m} w_i \, S^{+}_{i,d},$$

where $S^{+}_{i,d}$ is the commodity-level supply-increase indicator on day d and w_i is the production weight of commodity i; the composite index sums over all commodities, the oil index over petroleum alone, and the non-oil index over $\mathcal{I} \setminus Oil$. Unlike our earlier composite index, which implicitly weighted commodities based on news coverage, these weights are production-based—obtained from the GSCI Composite construction methodology and calculated as five-year moving averages—and thus provide a market-based measure of each commodity's relative importance in global production.[3] The indicators

[3]We 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.

for supply decreases and for demand increases or decreases are defined analogously, and we construct the net supply and net demand indices following equation (3), which defines net supply as supply contractions net of expansions—so that a *positive* value denotes a net supply contraction, the adverse shock whose transmission we trace throughout. Separating oil from non-oil lets us isolate the role of non-oil supply disturbances, which the existing literature has largely neglected in favor of petroleum.

Second, we isolate genuinely unexpected, supply-specific innovations from each market's daily net supply index through a multi-step residualization carried out *at the daily frequency*, before any temporal aggregation. We first purge predictable components with the same AR(20) residualization used in Section 2 (equation (4)), now applied to the aggregate market index rather than to individual commodities, so that what remains is news to market participants. We then go further, controlling for own-market demand by including the one-day lagged AR(20) residuals of the corresponding demand indicator, removing variation predictable from lagged demand innovations that would otherwise blend demand pressure into the supply shock. For the oil and non-oil indices, we further purge the one-day lagged supply and demand residuals of the *other* market, so that the two shocks do not proxy for one another; the composite index, which already spans all markets, requires no such cross-market adjustment.

Only then do we aggregate the resulting daily residuals to the monthly frequency, summing over the calendar days in each period:

$$S_{m,t} \equiv \sum_{d \in D(t)} \tilde{S}_{m,d},$$

where $\tilde{S}_{m,d}$ is the daily supply residual for market m and $D(t)$ collects the days in period t for which we have articles. By construction, the resulting shock is orthogonal to its own past, to lagged own-market demand, and—for the oil and non-oil indices—to lagged movements in the other market, so the variation identifying our macroeconomic responses reflects unanticipated, supply-specific disturbances rather than demand-driven or cross-market forces. An analogous procedure, applied to each commodity individually rather than to the aggregate indices, delivers the commodity-specific shocks $S_{i,t}$ that enter the heterogeneous-exposure design of equation (10).

Before turning to macroeconomic outcomes, we verify that these aggregated supply shocks move commodity prices and compute the price change that a one-standard-deviation shock represents. This exercise is distinct from the high-frequency validation in Section 2: there we documented the *contemporaneous* daily price response of individual commodities, whereas here we ask whether the *production-weighted* oil and non-oil supply indices that enter our local projections move aggregate real prices at the monthly frequency. Appendix Figure A.14 displays the response of real commodity spot prices to a one-standard-deviation supply contraction. The magnitude we carry forward is the response

on impact: a one-standard-deviation oil supply shock raises Brent by approximately 1 percent, while a one-standard-deviation non-oil supply shock raises the production-weighted basket by approximately 4 percent. The two indices differ in persistence—the oil response is concentrated on impact and dissipates over the following year, whereas the non-oil response builds and is sustained—but it is this on-impact change that provides the price denomination we use throughout this section, mapping each one-standard-deviation supply shock into the commodity-price movement it represents.

3.2 Aggregate 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.

We begin by estimating the following specification:

$$x_{j,t+k} = \beta(k)\, S_t + \sum_{l=1}^{L} c_l(k)\, x_{j,t-l} + c(k)\mathbf{Z}_{j,t} + \alpha_j(k) + u_{j,t}, \qquad k = 0 \ldots 24. \tag{8}$$

We focus on two dependent variables x, the consumer price index and industrial production—both series in logs. Both sets of series are obtained from the International Monetary Fund's International Financial Statistics (IFS) database. In equation (8), j indexes countries and k indexes the horizon k. We include country fixed effects $\alpha_j(k)$ to account for country-specific trends in industrial production or the price level, and a vector $Z_{j,t}$ of macroeconomic and financial controls. Last, we control for $L = 12$ 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).

Figure 5 presents the estimated coefficients $\beta(k)$ from equation (8) across all 60 countries. The top panel of the figure focuses on inflation, while the bottom panel focuses on industrial production. The first column of the figure presents results using the composite supply index, while the second and third column present results separately for supply shocks in the oil market and the other commodities, respectively. Examining the figure, we see that commodity supply shocks generate economically meaningful macroeconomic effects: and contracts industrial production by roughly 2 percent; the decline in industrial production is frontloaded and statistically significant over the first months.

Comparing the second and third column of the figure, we see that our effects are not driven exclusively by oil. A one-standard deviation negative oil supply shock—which corresponds to approximately a 1 percent increase in oil prices on impact— has only a modest (and not statistically significant) impact on inflation, while industrial production falls by approximately [C: 0.2] percentage points over the next year. By contrast, a one-standard deviation decrease in the supply of non-oil commodities—which

corresponds to approximately a 4 percent increase in the price of non-oil commodities—is associated with approximately 1 to 1.5 percentage point increase in inflation over the next two years, and a (short-run) decrease of approximately 1 percent in industrial production.[4]

The fact that non-oil supply disturbances matter at least as much as oil disturbances for global inflation and output, is particularly important given the fact that the existing literature has almost exclusively focused on the role of oil supply shocks. 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. The first of these explanations, however, doubles as a caveat: because the non-oil index bundles many commodities, a one-standard-deviation non-oil shock is not directly comparable to a single-commodity oil shock, so the aggregate comparison partly reflects bundled versus single-market disturbances. The commodity-by-commodity design we develop next sidesteps this concern, comparing exposures to the *same* commodity shock across countries rather than aggregating heterogeneous markets into a single index.

The average, cross-country responses above mask potential heterogeneity. Estimating specification (8) separately for countries that are overall net commodity importers and net exporters (Appendix Figures A.18 and A.19)—classified by the sign of the *commodity* component of their trade balance (net commodity exports summed across all commodities), rather than their total trade balance—reveals that the output contraction is concentrated among importers: industrial production falls for oil importers and considerably more for non-oil importers, while exporters are insulated (oil) or expand outright (non-oil). Inflation rises among importers as well: oil importers see prices climb by roughly 1 percentage point, while oil exporters are insulated. For non-oil commodities, however, exporters' prices rise by as much as those of importers. Motivated by these facts, and building on economic theory, the next

[4]These international results are consistent with US-only evidence: Appendix Figure A.17 reports local projections for the United States, confirming that supply contractions raise CPI and lower industrial production, with non-oil supply shocks generating effects at least as large as oil shocks over the sample period.

section exploits cross-country differences in commodity exposure to identify how a country's trade position shapes its response to commodity supply shocks.

3.3 Exploiting Heterogeneous Exposures

The average results in the previous section are identified purely from time-series variation, and the importer–exporter split, though cross-sectional, rests on a coarse, country-level measure of trade position. Neither of these designs allows for a causal interpretation of the results. Other macroeconomic factors may move with our commodity supply shocks, and a country's overall trade status may proxy for confounding differences across economies. We therefore sharpen the design by exploiting countries' cross-sectional differences in exposure to a *particular* commodity. To the extent that these country–commodity differences are pre-determined, they provide an additional source of variation. This design identifies a different object than the aggregate specification above: rather than the absolute response of a country to a marketwide supply shock, it recovers the *differential* response—how much more a highly exposed country reacts than an unexposed one to the same commodity shock.

Existing models deliver predictions for how these responses should vary with a country's commodity-specific net-trade position. A negative supply shock—a contraction that raises the world price of the commodity—acts as an adverse terms-of-trade movement for net importers and a favorable one for net exporters. For industrial production, the two forces reinforce one another: net importers face higher input costs and a negative income effect that depress activity, whereas net exporters enjoy a terms-of-trade windfall and an expansion of the commodity-producing sector itself, both of which cushion output. We therefore expect the output contraction to be systematically larger for net importers than for net exporters—consistent with the evidence on how commodity terms-of-trade movements shape output and external balances in commodity-dependent economies (Deaton and Miller, 1996; Kilian, Rebucci, and Spatafora, 2009).

For inflation, by contrast, two channels pull in opposite directions. The demand channel raises inflation among exporters, whose windfall is partly spent on domestically produced goods; the commodity-currency channel works the other way, as the adverse terms of trade depreciate the importer's currency and appreciate the exporter's, raising the domestic-currency price of imports for the former and lowering it for the latter. That said, in economies whose exchange rates are free to adjust, we expect the commodity-currency channel to dominate on balance: under the law of one price, import prices move with the exchange rate immediately and nearly one-for-one, whereas the demand channel loads only

on the smaller home-goods share of the consumption basket and transmits slowly through sticky home-goods prices.[5]

In sum, we expect net importers to display a larger contraction of output than net exporters in response to a negative commodity supply shock, and—assuming a sufficiently flexible exchange rate regime—a larger inflationary response as well. To test these predictions, we construct a more granular, commodity-level exposure index, computing the trade balance of country j in commodity i as

$$TB_{j,i,t} = \frac{X_{j,i,t} - M_{j,i,t}}{Y_{j,t}}. \tag{9}$$

Here, $X_{j,i,t}$ and $M_{j,i,t}$ denote the exports and imports of country j of commodity i in period t, respectively. As before, $Y_{j,t}$ corresponds to the country's GDP. We sign the index so that a country is a net exporter when $TB_{j,i,t} > 0$ and a net importer when $TB_{j,i,t} < 0$. To aid with the interpretation, we standardize (9) to zero mean and unit standard deviation.

We next estimate the following specification at the country j, commodity i, and month t level,

$$x_{j,i,t+k} = \gamma(k)\, S_{i,t} \times \left(-TB_{j,i,t}\right) + \sum_{l=1}^{L} c_l(k)\, x_{j,i,t-l} + c(k)\mathbf{Z}_{j,i,t} + \alpha_{j,i}(k) + \delta_{i,t}(k) + u_{j,i,t}. \tag{10}$$

Equation (10) is a stacked version of equation (8) above: the dependent variable and the controls do not vary across commodities for a given country–month, hence it is necessary to cluster the standard errors by both country and month; we also cluster by commodity to account for any residual serial correlation in our commodity-specific indices. In addition, we now include both country–by–commodity $\alpha_{j,i}$ and commodity-by-month $\delta_{i,t}$ fixed effects. The former absorb time-invariant unobserved heterogeneity in a country's exposure to a specific commodity (e.g., persistent differences in production technology or storage capacity), while the latter absorb the commodity-specific shock itself, along with any other time-varying disturbance common to all countries in that commodity-month.

Equation (10) allows us to identify the impact of a commodity supply shock on the dependent variable by comparing the response of countries that are highly exposed (importers, negative trade balance) to the average. That is, $\gamma(k)$ is identified purely from *within commodity-month* variation: holding fixed the global shock to a given commodity in a given month, do countries more reliant on that commodity for their net trade contract by more than countries with no exposure to it? By comparing countries facing an identical commodity shock in the identical month, this design nets out the confounding from countries' broader macroeconomic characteristics that could not be ruled out in

[5]Consistent with this view, commodity exporters' real exchange rates comove strongly with commodity prices (Cashin, Cespedes, and Sahay, 2004), exchange-rate adjustment is the central margin through which commodity price shocks transmit in these economies (Drechsel, McLeay, and Tenreyro, 2019), and world commodity-price movements have strong predictive power for CPI and PPI inflation in commodity-sensitive economies (Chen, Turnovsky, and Zivot, 2014).

the aggregate specification above. This design has the structure of a shift-share research design (Bartik, 1991; Goldsmith-Pinkham, Sorkin, and Swift, 2020): the commodity supply shock is the common *shift*, while a country's net-trade position in that commodity is the *share* that scales its exposure to the shift. The commodity-by-month fixed effects absorb the shift entirely, while the country-by-commodity fixed effects absorb the *average* level of each country's exposure to each commodity. Because the trade-balance share itself varies over time, however, these fixed effects remove only its time-invariant level and not its month-to-month realization, so identification exploits the interaction of the shift with a country's exposure—including its contemporaneous, time-varying component—rather than fixed cross-sectional shares alone. A causal interpretation of $\gamma(k)$ therefore rests on a share-exogeneity assumption in the spirit of Goldsmith-Pinkham, Sorkin, and Swift (2020): that a country's exposure to a commodity is uncorrelated with other forces moving its inflation or output within the same commodity-month. This assumption is most plausible insofar as trade positions are slow-moving relative to the high-frequency supply innovations, but it is not innocuous, since the contemporaneous share may itself respond to the conditions that drive the outcome.

Figure 6 presents the estimated coefficients $\gamma(k)$ for inflation and industrial production, pooling across all commodities. At the one-year horizon, a country with a one-standard deviation higher exposure to a given commodity ($TB_{j,i,t} = -1$) sees inflation increase by 0.02 percentage points and industrial production fall by approximately 0.05 percentage points more than a country with no net exposure to that commodity ($TB_{j,i,t} = 0$) following a one-standard-deviation negative supply shock in an individual commodity.These effects are largely persistent as the estimated coefficients are largely comparable across horizons.

At this point, it is worth commenting on the magnitudes of the $\gamma(k)$ estimates compared to $\beta(k)$: the coefficients plotted in Figure 6 are a very different object than the coefficients in the country-level regressions in Figure 5 for two reasons: first, the coefficients $\gamma(k)$ correspond to a unit-standard deviation shock to a single commodity whereas the coefficients $\beta(k)$ correspond to a unit standard deviation supply shock to the entire basket of commodities; second, the coefficients $\gamma(k)$ capture relative responses, and therefore their magnitude is influenced by the degree of cross-sectional dispersion in country–commodity exposures, while $\beta(k)$ capture absolute responses.

As before, Figure 7 examines separately the impact of oil supply shocks from shocks to other commodities. Examining the point estimates, we note that the estimated coefficients are somewhat smaller for oil relative to the other commodities, though the difference is not statistically significant. At the one-year horizon, a country with a one-standard deviation higher exposure to a given non-oil commodity ($TB_{j,i,t} = -1$) sees industrial production fall by roughly 0.5 percentage points more than a country with no net exposure to that commodity ($TB_{j,i,t} = 0$) following a one-standard-deviation

negative supply shock; the corresponding difference for oil is about half as large, at roughly 0.25 percentage points. The inflation response mirrors this pattern: maximum exposure to a non-oil commodity raises inflation by about 0.2 percentage points more than no exposure at the two-year horizon, versus roughly 0.06 percentage points for oil. The non-oil differential builds steadily with the horizon and is statistically significant from about the one-year mark onward, so the inflationary transmission of non-oil supply shocks is economically sizable and—at the horizons relevant for the business cycle—clearly different from zero. Notably, this importer premium is not apparent in the coarse country-level split, where non-oil importers and exporters inflate comparably (Appendix Figure A.19); isolating it requires differencing out the common commodity-month shock, precisely what this design does. As in the aggregate specification, then, non-oil commodities generate a differential transmission channel of comparable or greater economic importance than oil, now identified from within-commodity variation in exposure alone.

In sum, commodity supply shocks are global in reach but heterogeneous in impact. Net commodity importers face systematically larger and more persistent output losses *and* stronger inflationary responses than net exporters, who are partially insulated on both margins. The larger inflation response of importers is consistent with the commodity-currency channel dominating the demand channel: the depreciation of importers' currencies passes through to the domestic price of imported goods. The interaction specification, with country-by-commodity and commodity-by-month fixed effects, strips out confounding global forces and—under the share-exogeneity assumption discussed above—attributes this heterogeneity to countries' differential commodity exposure.

3.4 Robustness and Alternative Specifications

As discussed above, GSCI production weights are a practical market-based proxy for the theoretically correct Domar weights (Hulten, 1978). To assess the robustness of our results to the choice of aggregation weights, we construct an alternative set of weights derived from the Bureau of Economic Analysis (BEA) Benchmark Input-Output Accounts. For each of the four benchmark years for which detailed input-output data are available—2002, 2007, 2012, and 2017—we compute each commodity's *Domar weight* as its total intermediate use (summed across all purchasing industries) divided by nominal GDP, using the BEA "Use of Commodities by Industries" table before redefinitions at the detailed level ($\sim$400 commodity rows). Where multiple GSCI commodities share a single BEA commodity code, we allocate the BEA Domar weight proportionally to the GSCI reference production dollar weights within that group. For non-benchmark years we linearly interpolate between the two adjacent benchmark years; years after 2017 hold the 2017 weights fixed. The full commodity-to-BEA-code mapping, the Domar weight calculations, and a comparison against the GSCI weights are reported in Appendix 6.

4 Concluding Remarks

We develop a textual framework to measure supply and demand across the full commodity spectrum—energy, metals, agriculture, and livestock—addressing a challenge that, since Kydland and Prescott (1982), empirical work has been able to meet only for oil. The framework produces daily news-based proxies for supply and demand pressure across 20 commodities, 19 of them previously undocumented, which we further decompose into four underlying drivers—business-cycle developments, geopolitical risk, natural disasters, and climate change—finding that natural disasters and geopolitical risk dominate supply disruptions while the business cycle shapes demand. A battery of validation exercises—an independent human audit, impulse responses that reproduce established oil supply-shock measures, and price reactions concentrated on the day the news breaks—supports interpreting the series as unanticipated shocks. Our central finding is that this breadth matters for the macroeconomy: supply disruptions in non-oil commodities move inflation and industrial production at least as strongly as oil, so the literature's focus on petroleum has understated the importance of commodity supply shocks as a class.

Establishing that this result is not driven by oil alone requires separating causal transmission from the average, time-series responses. Exploiting the full granularity of the data at the country-commodity-month level—absorbing country-by-commodity and commodity-by-month fixed effects—we recover the differential response to a given commodity shock from within-commodity variation in exposure alone, ruling out confounding from countries' broader macroeconomic characteristics. Transmission is sharply heterogeneous: net importers suffer larger and more persistent output contractions and stronger inflation than net exporters, who are partially insulated on both margins, with the inflation gap consistent with the commodity-currency channel. Identified this way, non-oil supply shocks transmit at least as strongly as oil—so the headline result holds precisely where the aggregate evidence is most exposed to confounding.

Our work opens up several avenues for future work. First, the daily indices are a real-time input for monitoring and forecasting the commodity-driven component of inflation and economic activity, complementing existing nowcasting tools. Second, because the costs of commodity supply disruptions fall disproportionately on net importers, our findings raise policy questions we leave open: whether the strategic reserves and supply-security arrangements that importing economies maintain almost exclusively for petroleum should extend to critical non-oil commodities such as industrial metals and key agricultural inputs, and whether central banks in commodity-importing economies should weight non-oil commodity developments, not oil alone, in their reaction functions. Last, the growing role of climate change as a driver—increasingly concentrated in energy and in the industrial metals central

to the energy transition—invites a tighter link between our narrative measures and the literature on transition risk and critical minerals.

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, Enghin. 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.

Baqaee, 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.

Bartik, Timothy J. 1991. *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

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.

Campa, José Manuel and Linda S. Goldberg. 2005. "Exchange Rate Pass-Through into Import Prices." *The Review of Economics and Statistics* 87 (4):679–690.

Cashin, Paul, Luis Cespedes, and Ratna Sahay. 2004. "Commodity currencies and the real exchange rate." *Journal of Development Economics* 75 (1):239–268.

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.

Chen, Yu-chin, Stephen J. Turnovsky, and Eric Zivot. 2014. "Forecasting inflation using commodity price aggregates." *Journal of Econometrics* 183 (1):117–134.

Chiaie, Simona Delle, Laurent Ferrara, and Domenico Giannone. 2022. "Common factors of commodity prices." *Journal of Applied Econometrics* 37 (3):461–476.

Datta, D. Deepa and A. Daniel Dias. 2019. "Oil Shocks: A Textual Analysis Approach." Tech. rep.

Deaton, Angus and Ron Miller. 1996. "International Commodity Prices, Macroeconomic Performance and Politics in Sub-Saharan Africa." *Journal of African Economies* 5 (3):99–191.

Drechsel, Thomas, Michael McLeay, and Silvana Tenreyro. 2019. "Monetary Policy for Commodity Booms and Busts." Tech. rep., CEPR Discussion Paper No. DP14030.

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.

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78 (1):35–71.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8):2586–2624.

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.

Hulten, Charles R. 1978. "Growth Accounting with Intermediate Inputs." *The Review of Economic Studies* 45 (3):511–518.

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.

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 Pragidis. 2019. "Assimilation of oil news into prices." *International Review of Financial Analysis* 63 (C):105–118.

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.

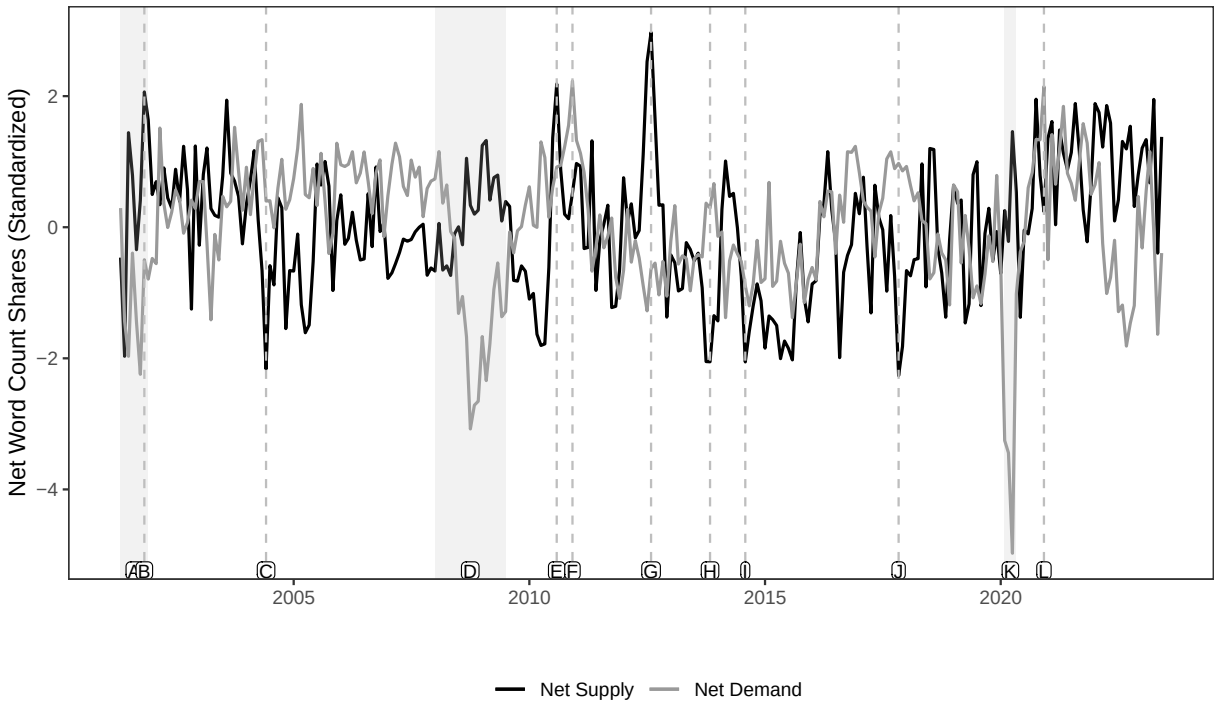
Tetlock, Paul C. 2007. "Giving Content to Investor Sentiment: The Role of Media in the Stock Market." *Journal of Finance* 62 (3):1139–1168.

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.

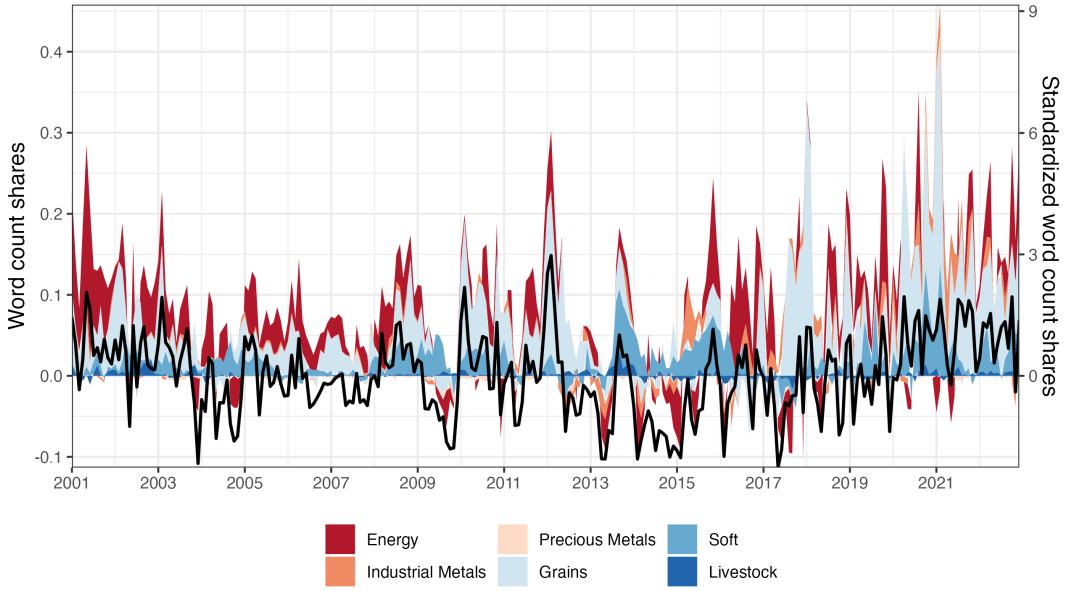
Figure 1: Net supply and net demand indices across all commodity-related news articles



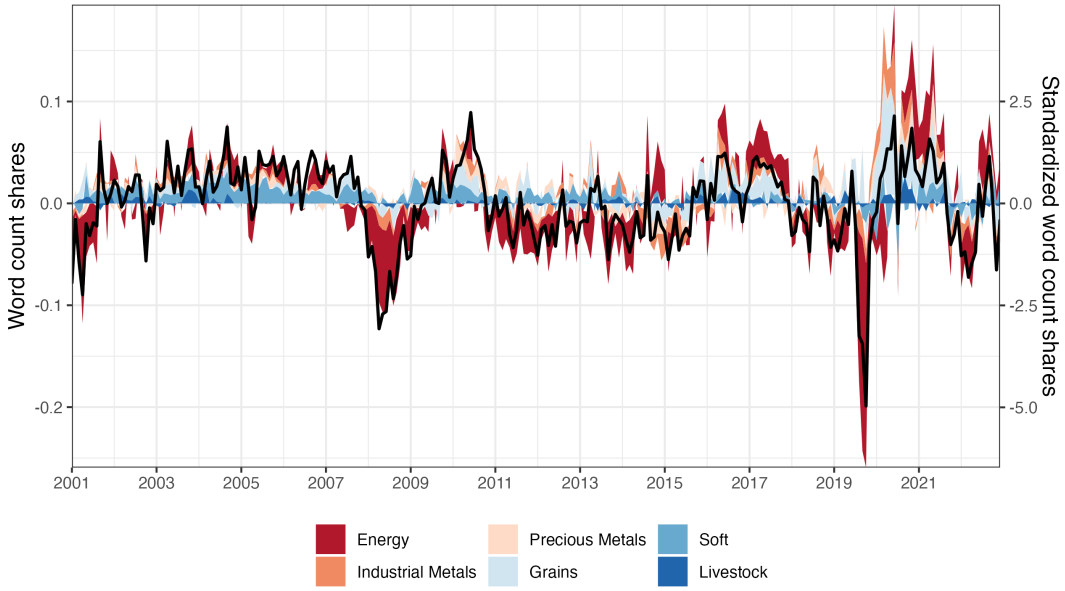
Note: This figure plots the standardized net supply and net demand indicators from 2001 to 2023, constructed using all commodity-related news articles in the Reuters corpus. The bars and lines map a number of well-known commodity-wide developments that exceeded ± 2 S.D. These events are: [A] US Recession 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
is this using GSCI weights? If yes, should say, if no it should in order to be consistent with the rest. Split into net supply (top) net demand (bottom)

Figure 2: Decomposition of composite indices across categories

Panel A: Net supply



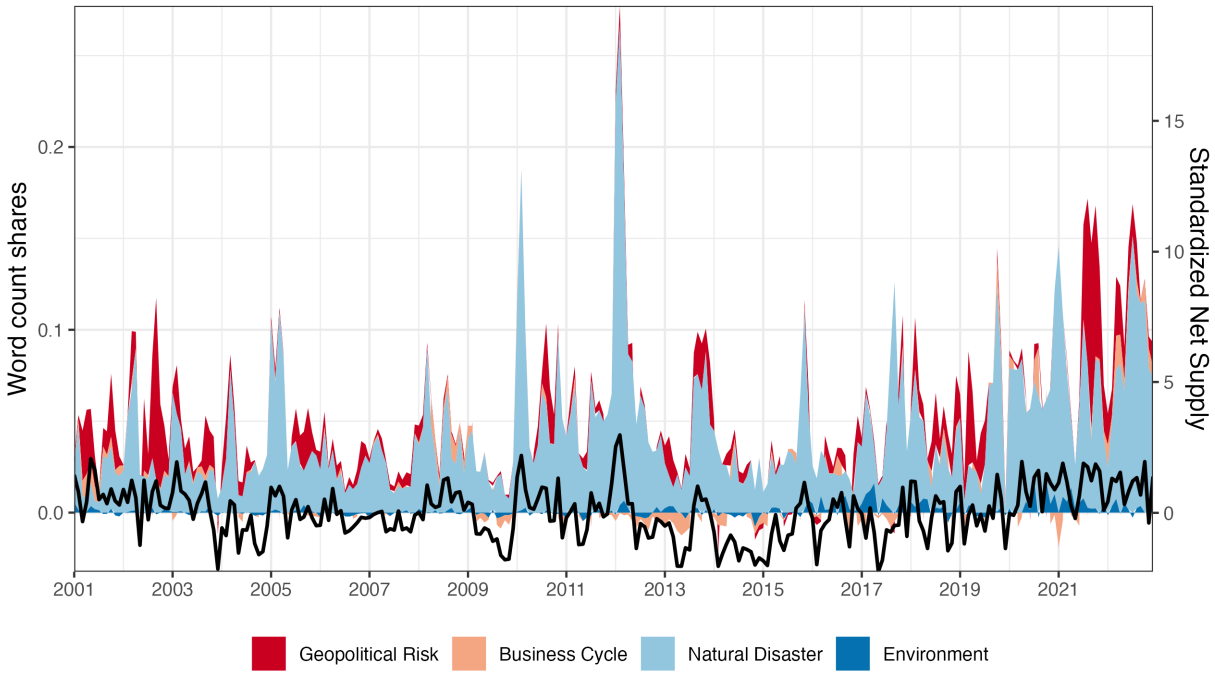
Panel B: Net demand



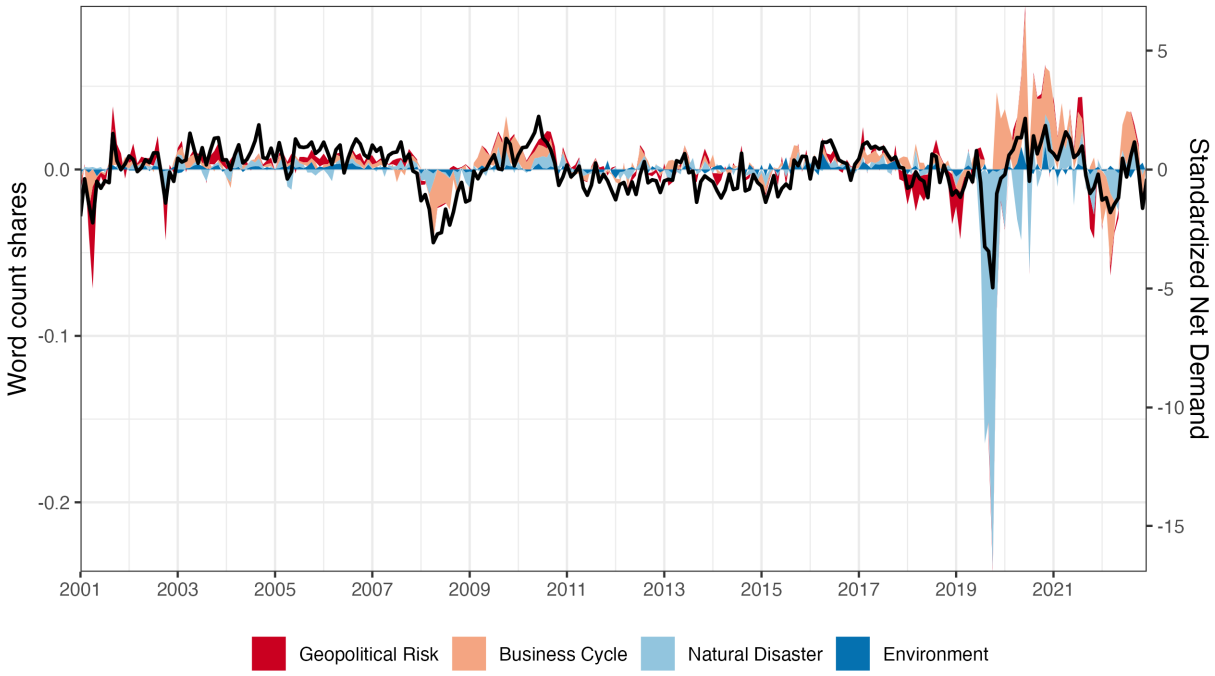
missing caption. not understanding exactly how the decomp works...

Figure 3: Factor decomposition: Net supply and net demand

A. Net supply

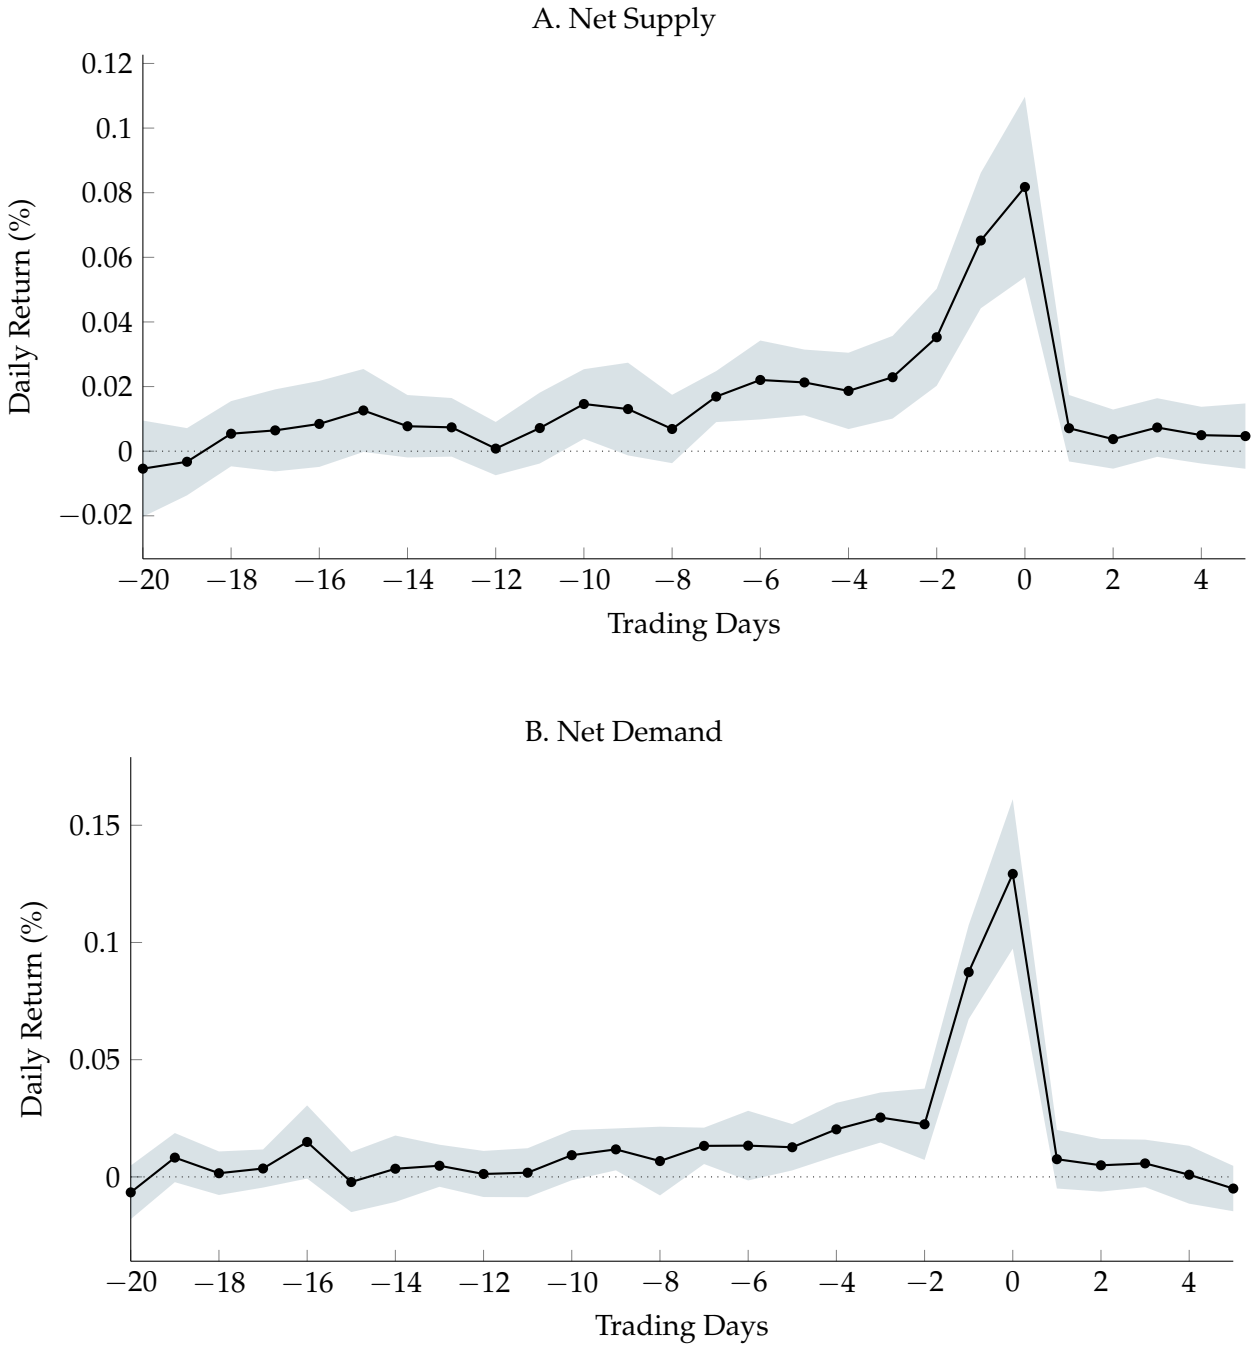


B. Net demand



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.

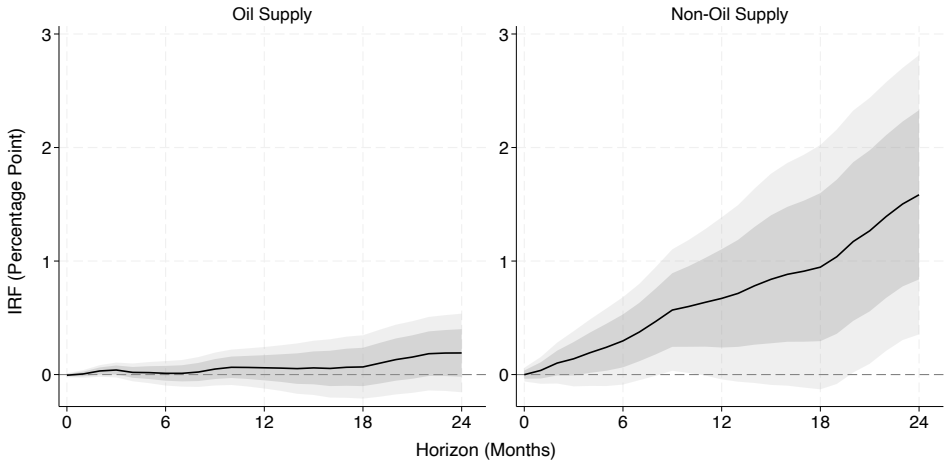
Figure 4: Commodity supply and demand shocks and high-frequency commodity returns



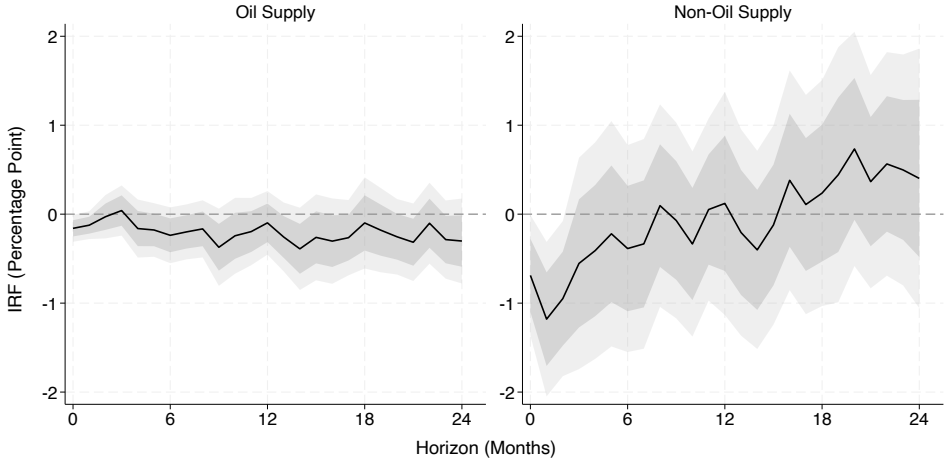
Note: This figure plots the estimated coefficients $\beta(k)$ and $\gamma(k)$ from Equation 6 for event-time horizons $k = -20, \ldots, 5$. Negative values of k correspond to days preceding the news event, whereas positive values correspond to days following the event. The shaded bands denote 95% confidence intervals. The dependent variable is the standardized AR(20) residuals of our supply and demand disturbances constructed following equation (4).

Figure 5: Impact of commodity shocks on inflation and industrial production

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

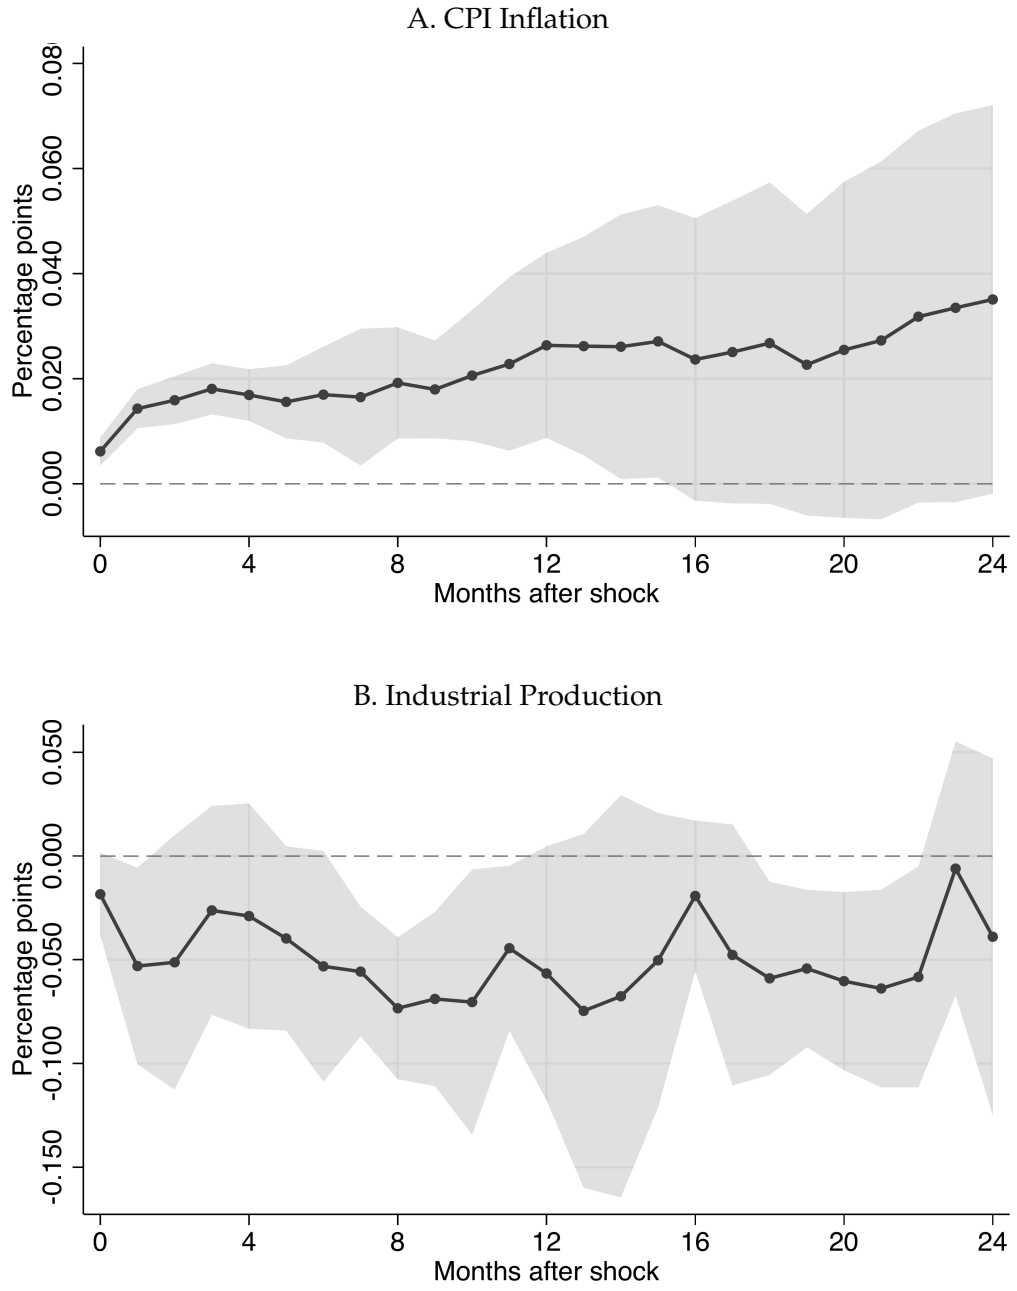


B. Industrial Production (left: Brent; right: weighted non-oil commodity basket)



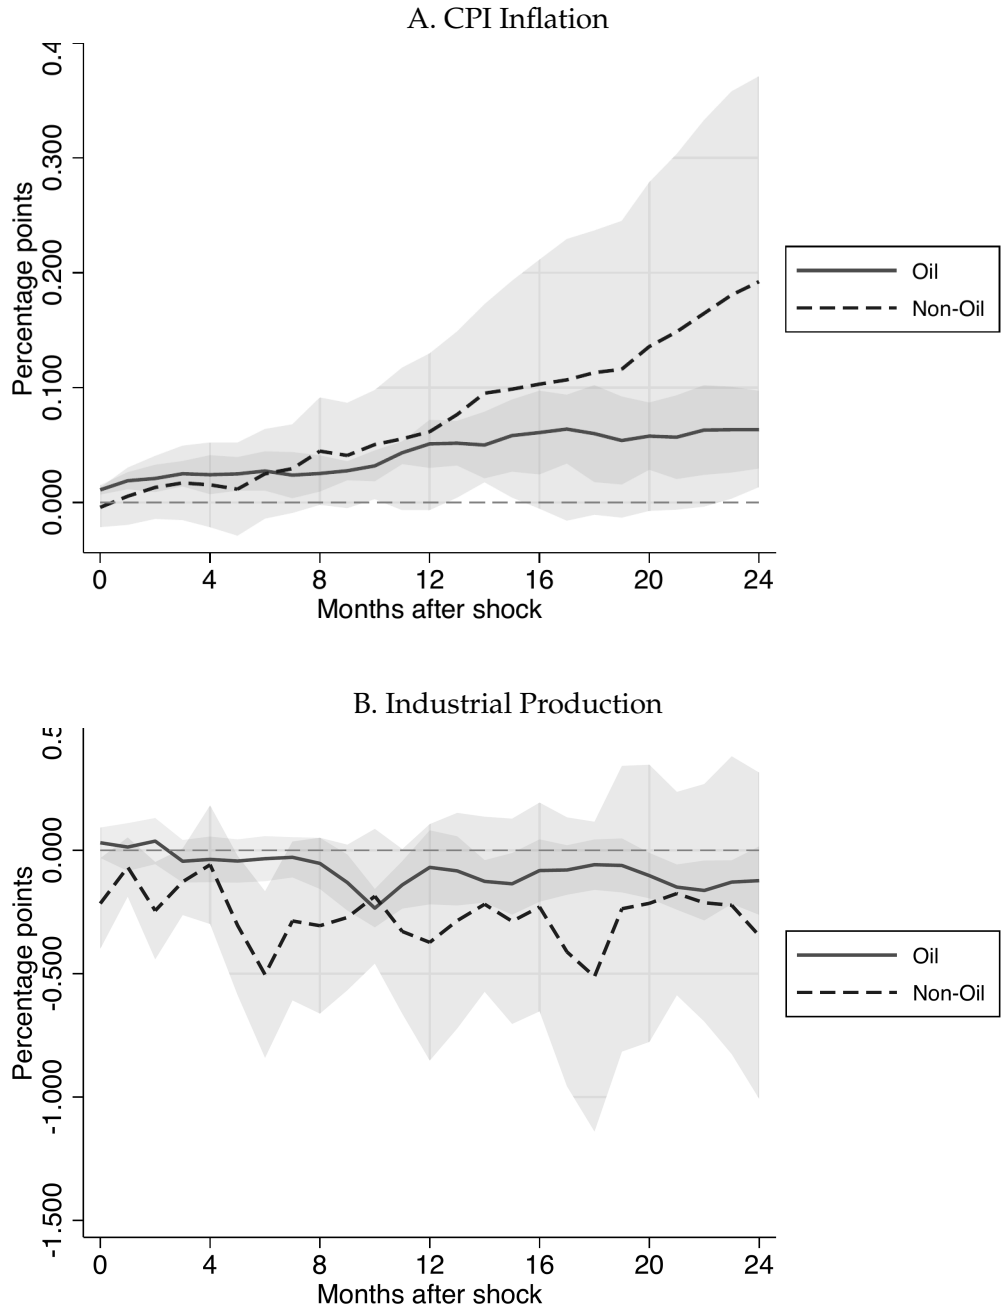
Note: This figure shows the impulse response of inflation (Panel A) and industrial production (Panel B) to a one-standard-deviation commodity supply contraction. In particular,. the figure plots the estimated coefficients $\beta_i(k)$ from equation (8) across all 60 countries. The left column uses the oil supply index; the right column uses the non-oil production-weighted supply index. Add a third panel, ordered first, that combines the oil and non oil into a single supply shock. explain what weights are. Standard errors are adjusted according to the Driscoll-Kraay procedure.

Figure 6: Macroeconomic impact of commodity shocks, accounting for commodity-specific time trends



Note: Figure plots the estimated coefficients $\gamma(k)$ from equation (10) in the main text. These coefficients correspond to the differential response to a unit standard deviation supply shock for a specific commodity for a country whose normalized commodity-specific trade balance (**??**) is one standard deviation higher than the average trade balance. clarify if the standard deviation is computed unconditionally or by commodity. All specifications include country $\times$ month and commodity $\times$ month fixed effects. Consequently, the coefficients $\gamma(k)$ are identified from within commodity-month variation in exposure only. needs better title

Figure 7: Macroeconomic impact of commodity shocks, comparing oil vs other supply shocks



Note: Figure plots the estimated coefficients $\gamma(k)$ from equation (10) in the main text, estimated separately for oil and non-il commodities. These coefficients correspond to the differential response to a unit standard deviation supply shock for a specific commodity for a country whose normalized commodity-specific trade balance (**??**) is one standard deviation higher than the average trade balance. clarify if the standard deviation is computed unconditionally or by commodity. All specifications include country $\times$ month and commodity $\times$ month fixed effects. Consequently, the coefficients $\gamma(k)$ are identified from within commodity-month variation in exposure only. Also put the legend elsewhereas the figure looks a bit ugly and non-centered

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 net supply decrease & net demand increase indicators

Commodity Returns (%)	Narrative Indices (Raw Series)			Narrative Indices (AR20 residuals)		
	(1)	(2)	(3)	(4)	(5)	(6)
Net Supply	0.047***	0.042***	0.043***	0.049***	0.045***	0.045***
	(6.29)	(6.77)	(6.87)	(5.94)	(6.57)	(6.67)
Net Demand	0.085***	0.073***	0.073***	0.087***	0.074***	0.074***
	(7.26)	(6.76)	(7.06)	(7.16)	(6.73)	(6.98)
Observations	106,880	111,940	111,920	106,480	111,540	111,540
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

Commodity Returns (%)	Narrative Indices (Raw Series)			Narrative Indices (AR20 residuals)		
	(1)	(2)	(3)	(4)	(5)	(6)
Net Supply	0.003	-0.004	-0.004	0.003	-0.004	-0.004
	(0.94)	(-1.52)	(-1.50)	(0.85)	(-1.59)	(-1.57)
Net Demand	0.009***	-0.006*	-0.006*	0.009***	-0.007**	-0.007**
	(3.91)	(-1.96)	(-1.97)	(3.40)	(-2.28)	(-2.28)
Observations	96,192	100,746	100,728	95,832	100,386	100,386
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

Note: This table reproduces columns (4) to (6) of Table 2 but this time we estimate Equation (5) for the daily GSCI component returns of commodity i, but replace the commodity-specific supply and demand indicators with those constructed for commodity $j \neq i$.

5 Data Construction, Dictionaries, and Audit Process

5.1 Data Sources

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, mcoff, 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.

5.2 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.

5.3 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.

5.4 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 A.4 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).

5.5 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 1. The algorithm retains sentences that explicitly reference the focal commodity and excludes content centered on competing commodities.

5.6 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.

5.7 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 A.10 to Figure A.13 plot the quarterly comparison for each of the supply and demand components along with the net index performance of the two methods.

6 BEA Input-Output Weights

Data Sources

We use four editions of the BEA Benchmark Input-Output (I-O) Accounts: 2002, 2007, 2012, and 2017. These are the most detailed publicly available I-O tables for the United States, each covering approximately 400 commodity rows and 400 industry columns. In each benchmark the relevant table is

the *Use of Commodities by Industries, Before Redefinitions* (producers' prices), which records the dollar value of each commodity purchased by each industry as an intermediate input, as well as the purchases by final-demand categories (personal consumption, investment, government, exports). We download the following files:

- *2002*: Fixed-width text file `REV_NAICSUseDetail 4-24-08.txt` from the BEA Historical Benchmark I-O archive (https://apps.bea.gov/industry/zip/2002detail.zip).
- *2007, 2012, 2017*: Sheets from the Excel workbook `IOUse_Before_Redefinitions_PRO_2017_Detail.xlsx` (https://apps.bea.gov/industry/xls/IOUse_Before_Redefinitions_PRO_2017_Detail.xlsx), which contains a separate worksheet for each of these three benchmark years.

All values are in millions of current-year US dollars.

Domar Weight Construction

For each commodity row i in the Use table, we sum the producers'-price transactions across all industry-column purchasers, excluding final-demand categories (personal consumption, private investment, government, exports, and changes in inventories, identified by column codes beginning with `F`). This yields U_i, the total intermediate use of commodity i in the economy. The *Domar weight* is then

$$\lambda_i = \frac{U_i}{Y},$$

where Y is nominal GDP in the benchmark year (BEA NIPA Table 1.1.5): \$10,977,514 million (2002), \$14,477,635 million (2007), \$16,197,007 million (2012), and \$19,519,354 million (2017). In the 2007, 2012, and 2017 benchmark Excel files the column `T001` ("Total Intermediate") directly contains U_i for each commodity row, so no summation is required. In the 2002 text file, intermediate use is recovered by summing the producers'-value field across all rows associated with commodity i whose industry code does not begin with `F` or `V` (value added).

We normalize the resulting Domar weights across the 20 GSCI commodity markets to obtain weights that sum to one within our commodity universe:

$$\hat{w}_i^{\mathrm{IO}} = \frac{\lambda_i}{\sum_{j=1}^{20} \lambda_j}.$$

These normalized weights are directly comparable to the GSCI reference production dollar weights (RPDWs), which are also normalized to sum to unity across the same commodity universe.

Commodity-to-BEA Code Mapping

Table A.5 reports the mapping from each of the 20 GSCI commodity markets to the corresponding BEA Input-Output commodity code. Because the BEA classification system (based on NAICS) changed between the 2002 and 2007 benchmark revisions, we report separate codes for the two eras. Several features of the mapping warrant note.

Shared codes. Several GSCI commodities map to the same BEA commodity row, because the BEA aggregates commodity production at a level coarser than the GSCI. In these cases we allocate the BEA row's Domar weight proportionally to each constituent GSCI commodity's share of the group's total GSCI reference production dollar weight (RPDW). Affected pairs are: crude oil and natural gas (both in `211000`: Oil and gas extraction); gasoil, gasoline, and heating oil (all in `324110`: Petroleum refineries); zinc and nickel in 2002 (both in `331419`: Primary smelting of nonferrous metals exc. Cu/Al); and zinc and nickel in 2007–17 (both in `331410`: Nonferrous metal except Al smelting and refining).

Code changes across benchmarks. Several commodity codes differ between the 2002 and 2007–17 benchmarks, reflecting methodological changes in the 2013 BEA comprehensive revision. The most consequential are: (i) aluminum moves from `33131A` (2002) to `331313` (2007–17); (ii) copper's primary smelting splits (`331411` in 2002) are merged into the broader nonferrous smelting category `331410` in 2007–17; and (iii) the separate codes for sugarcane farming (`1119A0`), cotton farming (`111920`), and all other crop farming (`1119B0`) in 2002 are consolidated into a single code `111900` (Other crop farming) in 2007–17. For precious metals, the 2002 code `2122A0` covers gold, silver, and other metal ore mining; in 2007–17, the same code expands to include iron ore mining, which inflates the measured intermediate use of this row and slightly overstates the Domar weights for gold and silver in the later benchmarks.

Commodities without US primary production. Cocoa and coffee are not produced domestically in the United States and therefore have negligible or zero output in the BEA domestic production rows. For the 2002 benchmark we assign both to `1119B0` (All other crop farming), which captures related agricultural imports and domestic processing activity. For 2007–17 both map to `111900` (Other crop farming). Because this category is a broad aggregate capturing many unrelated crops, the resulting Domar weights for cocoa and coffee substantially overstate their importance as US intermediate inputs. Users relying on these estimates for cocoa and coffee should treat the weights as upper bounds.

Temporal Extension

The four benchmark years (2002, 2007, 2012, 2017) span the full sample of the paper (2001–2023). For years not coinciding with a benchmark, we construct annual BEA-derived weights by linear interpolation between the two adjacent benchmark Domar weight vectors:

$$\hat{w}_{i,t}^{\mathrm{IO}} = \left(1 - \frac{t - t_0}{t_1 - t_0}\right) \hat{w}_{i,t_0}^{\mathrm{IO}} + \frac{t - t_0}{t_1 - t_0} \hat{w}_{i,t_1}^{\mathrm{IO}}, \quad t_0 \le t \le t_1,$$

where t_0 and t_1 are the surrounding benchmark years. For years after 2017 (the most recent available benchmark at the detail level), we hold the 2017 Domar weights fixed. The interpolation is re-normalized after each step so that weights sum to one. A separate CSV file for each sample year (2002–2022) as well as a combined long-format file are available in the replication archive.

Comparison of GSCI and BEA Weights

Table A.6 reports the GSCI normalized production weights and the BEA Domar weights (normalized, 2002 benchmark) for all 20 commodities, together with their ratio. The comparison confirms the theoretical predictions in Section 3: commodities with large final-demand fractions (gold, silver, and to a lesser extent the grain crops) are overweighted by the GSCI, while commodities that are disproportionately important as domestic intermediate inputs (live cattle, refined petroleum products) are underweighted. For crude oil and natural gas—the dominant GSCI constituents—the two weighting schemes agree closely. We interpret the broad similarity across the weighting schemes as validation that the GSCI provides a reasonable first-order approximation to the theoretically motivated Domar weights for the purpose of aggregating supply disturbances.

Table A.1: Keywords for the global commodity supply and demand index

Energy	**Industrial metals**	**Precious metals**	**Agricultural**		**Livestock**
			Grains	**Softs**	
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 A.2: Regression of S&P GSCI commodity-specific daily returns on net supply and demand indices

	A. Energy					B. Grains			C. Precious metals	
	Gas	Gasoline	Oil	Gasoil	Heating Oil	Corn	Soybean	Wheat	Gold	Silver
Net Supply (Standardized)	0.0129	0.0281*	0.1295***	0.0382**	-0.0018	0.0684***	0.0602***	0.1326***	0.0029	-0.0007
	(0.0134)	(0.0133)	(0.0133)	(0.0134)	(0.0133)	(0.0133)	(0.0133)	(0.0132)	(0.0134)	(0.0134)
Net Demand (Standardized)	0.0712***	0.0533***	0.0885***	-0.0059	0.0556***	0.0984***	0.1283***	0.1016***	0.0687***	0.0583***
	(0.0134)	(0.0134)	(0.0136)	(0.0134)	(0.0133)	(0.0133)	(0.0133)	(0.0132)	(0.0133)	(0.0134)
Within R^2	0.006	0.010	0.029	0.007	0.008	0.014	0.021	0.029	0.005	0.004
N	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597

	D. Industrial metals				E. Softs				F. Livestock	
	Aluminium	Copper	Nickel	Zinc	Cocoa	Coffee	Cotton	Sugar	Cattle	Hog
Net Supply (Standardized)	0.0511***	0.0442***	0.0250	0.0377**	0.0534***	0.0513***	0.0726***	0.0781***	0.0469***	0.0504***
	(0.0134)	(0.0132)	(0.0134)	(0.0134)	(0.0133)	(0.0134)	(0.0132)	(0.0133)	(0.0130)	(0.0132)
Net Demand (Standardized)	0.0269*	0.1509***	0.0575***	0.0321*	0.0938***	0.0923***	0.1302***	0.0922***	0.2151***	0.1622***
	(0.0134)	(0.0132)	(0.0134)	(0.0134)	(0.0134)	(0.0134)	(0.0132)	(0.0133)	(0.0131)	(0.0132)
Within R^2	0.009	0.030	0.005	0.007	0.012	0.011	0.024	0.014	0.051	0.030
N	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597	5,597

missing caption

Table A.3: Data Sources and Variable Definitions

Variable	Description	Source
Energy		
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
Metals		
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
Agriculture		
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
Livestock		
Cattle	S&P GSCI Live Cattle Spot (GSLCSPT)	Datastream
Lean Hogs	S&P GSCI Lean Hogs Spot (GSLHSPT)	Datastream
Control variables		
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

Variable	Description	Source
VIX Index	CBOE Volatility Index: VIX (VIXCLS)	FRED
Macro variables		
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 (CPI-AUCSL)	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 A.4: Dictionary words: Supply and demand by commodity, and increase/decrease word stems

Panel A: Global Commodity Index

Supply suppl*, produc*, output

Demand demand*, consum*, buy*, purchas*

Panel B: Crude Oil

Supply suppl*, produc*, output, discovery, glut*, reserv*, surplus*, rig*

Demand demand*, consum*, buy*, util*, drain*, deplet*, refin*, purchas*

Panel C: Natural Gas

Supply suppl*, produc*, output, rig*, drill*

Demand demand*, consum*, buy*, purchas*

Panel D: Wheat

Supply suppl*, produc*, output, crop*, planting*, farm*, harvest*

Demand 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*, heigh*, 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*, stumbl*, 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*, dwindl*, plummet*, slip*, decreas*, ebb*, plung*, slow*

Table A.5: Mapping of 20 GSCI Commodity Markets to BEA Input-Output Codes

Commodity	Group	BEA Code (2002)	BEA Code (2007–17)	BEA Description (2002)	BEA Description (2007–17)	Notes
Crude Oil	Energy	211000	211000	Oil & gas extraction	Oil & gas extraction	† Shares code with natural gas; split by GSCI RPDW.
Natural Gas	Energy	211000	211000	Oil & gas extraction	Oil & gas extraction	† Shares code with crude oil; split by GSCI RPDW.
Gasoil	Energy	324110	324110	Petroleum refineries	Petroleum refineries	† Shares code with gasoline and heating oil; split by GSCI RPDW.
Gasoline	Energy	324110	324110	Petroleum refineries	Petroleum refineries	† Shares code with gasoil and heating oil; split by GSCI RPDW.
Heating Oil	Energy	324110	324110	Petroleum refineries	Petroleum refineries	† Shares code with gasoil and gasoline; split by GSCI RPDW.
Copper	Ind. Metals	331411	331410	Primary smelting & refining of copper	Nonferrous metal (exc. Al) smelting & refining	Code change: 2002 copper-specific code merged into broader nonferrous smelting in 2007+.
Aluminum	Ind. Metals	33131A	331313	Alumina refining & primary aluminum production	Alumina refining & primary aluminum production	Code renamed in 2007 revision.
Zinc	Ind. Metals	331419	331410	Primary smelting of nonferrous metals (exc. Cu/Al)	Nonferrous metal (exc. Al) smelting & refining	† Shares code with nickel; split by GSCI RPDW.
Nickel	Ind. Metals	331419	331410	Primary smelting of nonferrous metals (exc. Cu/Al)	Nonferrous metal (exc. Al) smelting & refining	† Shares code with zinc; split by GSCI RPDW.
Gold	Prec. Metals	2122A0	2122A0	Gold, silver, & other metal ore mining	Iron, gold, silver, & other metal ore mining	Code expanded in 2007+ to include iron ore, inflating intermediate use. † Shares code with silver; split by GSCI RPDW.

Table A.5 – *continued*

Commodity	Group	BEA Code	BEA Code	BEA Description	BEA Description	Notes
Silver	Prec. Metals	2122A0	2122A0	Gold, silver, & other metal ore mining	Iron, gold, silver, & other metal ore mining	† Shares code with gold; split by GSCI RPDW.
Corn	Grains	1111B0	1111B0	Grain farming	Grain farming	† Shares code with wheat; split by GSCI RPDW.
Wheat	Grains	1111B0	1111B0	Grain farming	Grain farming	† Shares code with corn; split by GSCI RPDW.
Soybean	Grains	1111A0	1111A0	Oilseed farming	Oilseed farming	Dedicated code in all benchmark years.
Sugar	Softs	1119A0	111900	Sugarcane & sugar beet farming	Other crop farming	† Shares 2007+ code with cocoa, coffee, and cotton.
Cocoa	Softs	1119B0	111900	All other crop farming	Other crop farming	‡ No US production; mapped to nearest aggregate. Weight is an upper bound.
Cotton	Softs	111920	111900	Cotton farming	Other crop farming	Dedicated code in 2002; merged into 111900 in 2007+.
Coffee	Softs	1119B0	111900	All other crop farming	Other crop farming	‡ No US production; mapped to nearest aggregate. Weight is an upper bound.
Hogs	Livestock	112A00	112A00	Animal production, exc. cattle and poultry	Animal production, exc. cattle and poultry	Dedicated code; includes hog & pig farming.
Live Cattle	Livestock	1121A0	1121A0	Cattle ranching & farming	Beef cattle ranching & farming, incl. feedlots	Dedicated code; description clarified in 2007 revision.

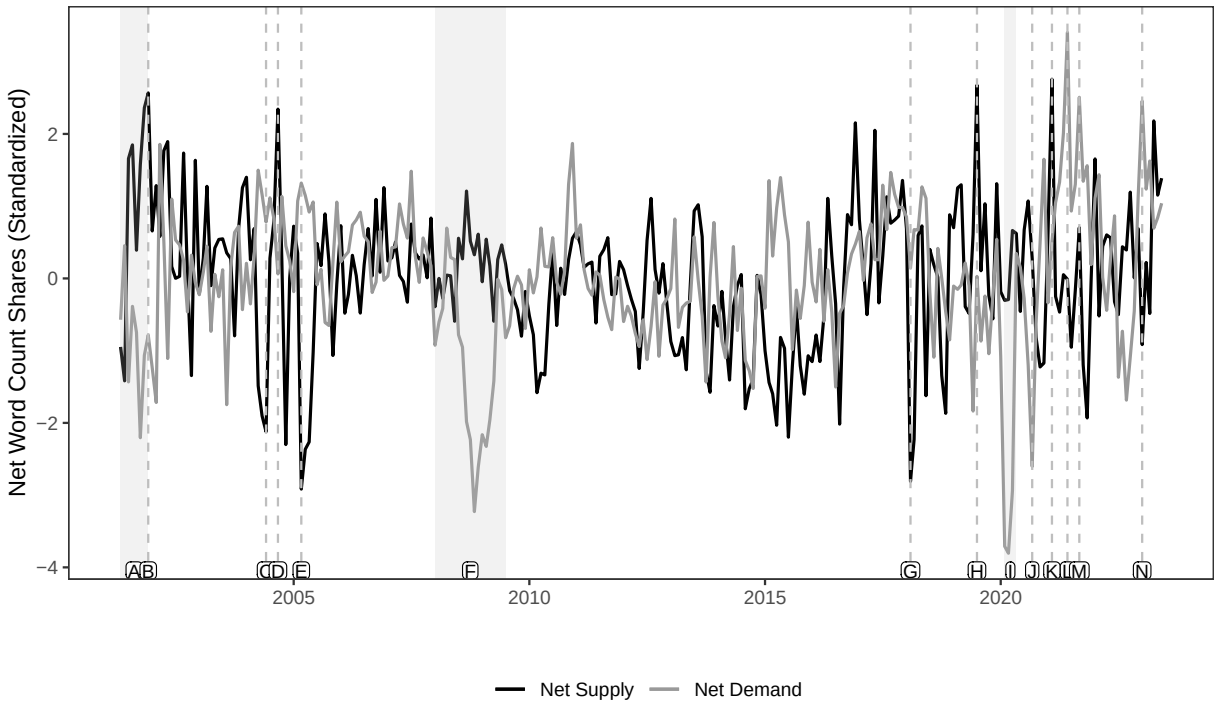
Notes. BEA commodity codes follow the NAICS-based classification of the Benchmark Input-Output Accounts. Where multiple GSCI commodities share a BEA code (indicated by the † symbol), the BEA row's Domar weight is allocated to constituent commodities proportionally to their GSCI reference production dollar weights within the group. For cocoa and coffee (‡), no dedicated US-production BEA row exists; both are mapped to the nearest available "other crop farming" aggregate. Code changes across benchmarks reflect the 2013 BEA comprehensive revision of the Industry Economic Accounts.

Table A.6: GSCI Production Weights vs. BEA Domar Weights: 2002 Benchmark

Commodity	**Group**	**GSCI Weight** (normalized, %)	**BEA Domar Weight** (normalized, %)	**Ratio (GSCI/BEA)**
Energy				
Crude Oil		33.45	30.91	1.08
Natural Gas		8.66	8.00	1.08
Gasoil		8.78	10.97	0.80
Gasoline		8.19	10.23	0.80
Heating Oil		7.24	9.05	0.80
Industrial Metals				
Copper		2.73	1.02	2.67
Aluminum		3.44	2.81	1.22
Zinc		0.71	0.74	0.96
Nickel		1.07	1.11	0.96
Precious Metals				
Gold		2.37	0.41	5.78
Silver		0.36	0.06	5.78
Grains				
Corn		5.10	2.58	1.98
Wheat		5.10	2.58	1.98
Soybean		2.25	2.12	1.06
Softs				
Sugar		1.42	0.44	3.24
Cocoa		0.59	2.33†	0.25
Cotton		1.54	0.51	3.02
Coffee		0.47	1.86†	0.25
Livestock				
Hogs		2.85	3.03	0.94
Live Cattle		3.68	9.23	0.40
Total		100.0	100.0	—

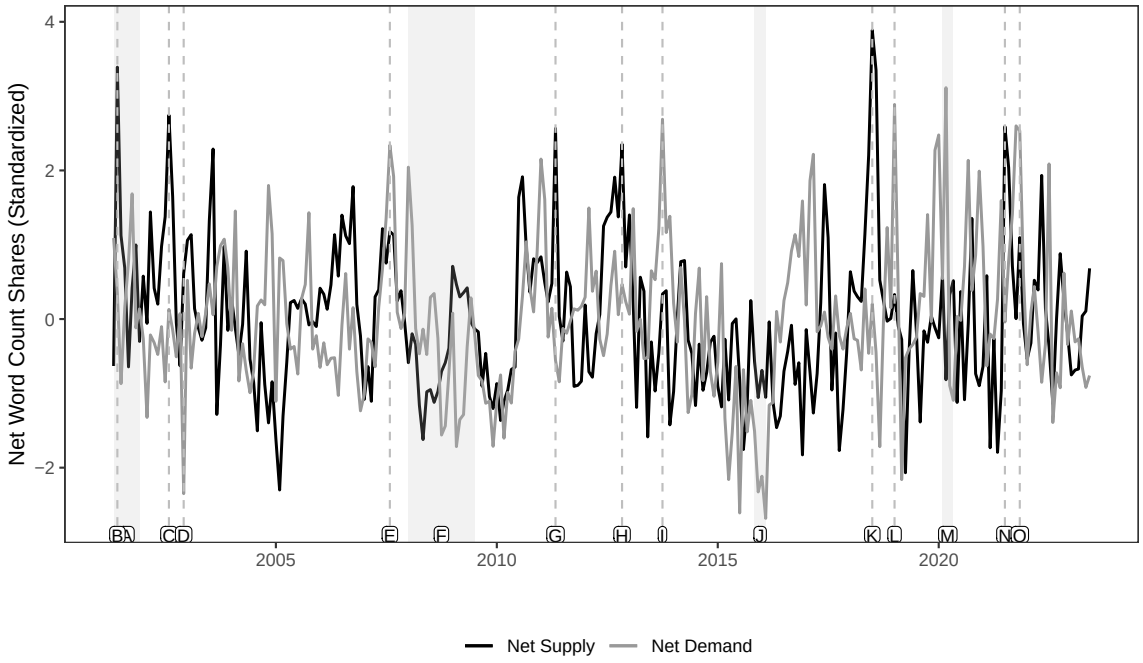
Notes. Weights are normalized to sum to 100 within the 20-commodity universe. GSCI weights are S&P GSCI Reference Production Dollar Weights (RPDWs) for 2002. BEA Domar weights equal each commodity's total intermediate use (summed across all purchasing industries in the BEA 2002 Benchmark Use Table before redefinitions, producers' prices) divided by US nominal GDP ($10,977,514 million), then normalized. Where multiple GSCI commodities share a BEA code, the code's Domar weight is distributed proportionally to the constituent GSCI RPDWs. † Cocoa and coffee are mapped to the broad "All other crop farming" aggregate (`1119B0`), which has no dedicated US production row; these weights overstate the true intermediate-use share and should be treated as upper bounds.

Figure A.1: Crude oil index: Net supply and net demand



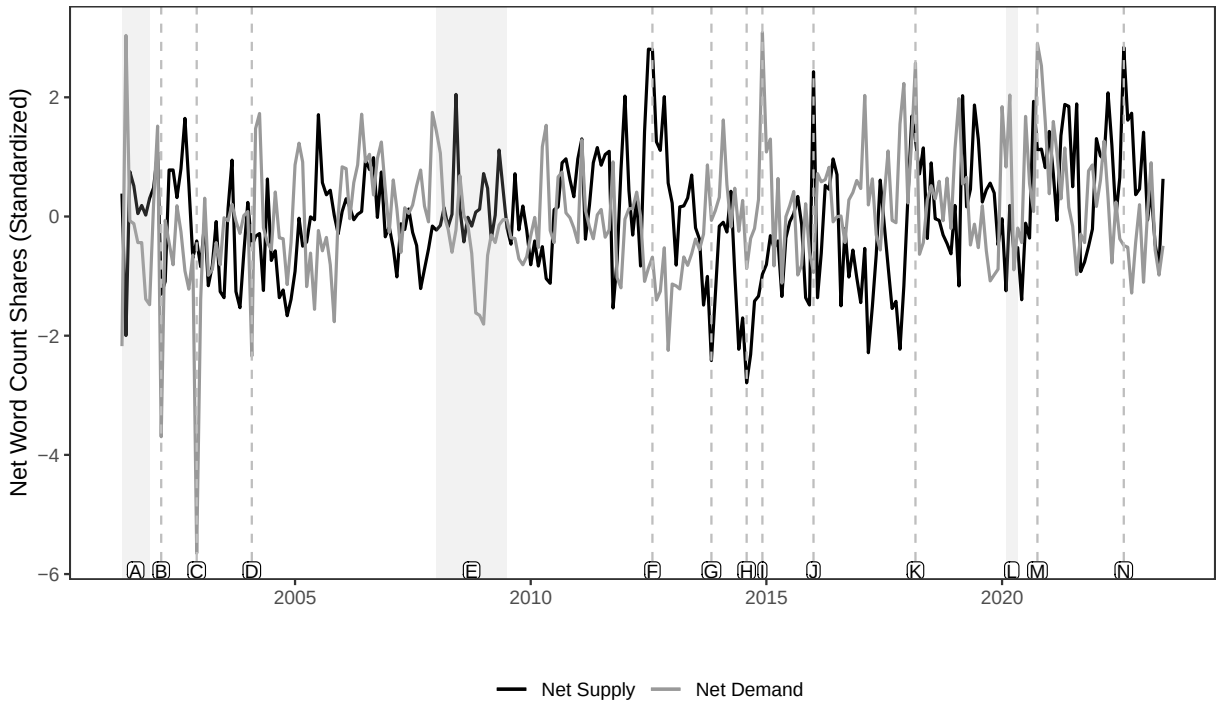
Note: This figure plots the standardized net supply and demand indicators for the period between 2021 and 2023. The bars map a number of well-known oil sector developments that are in the top 1% (two-tailed) of events. These events are: [A] US Recession, [B] OPEC 118th Extraordinary Meeting (production 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 A.2: Wheat index: Net supply and net demand



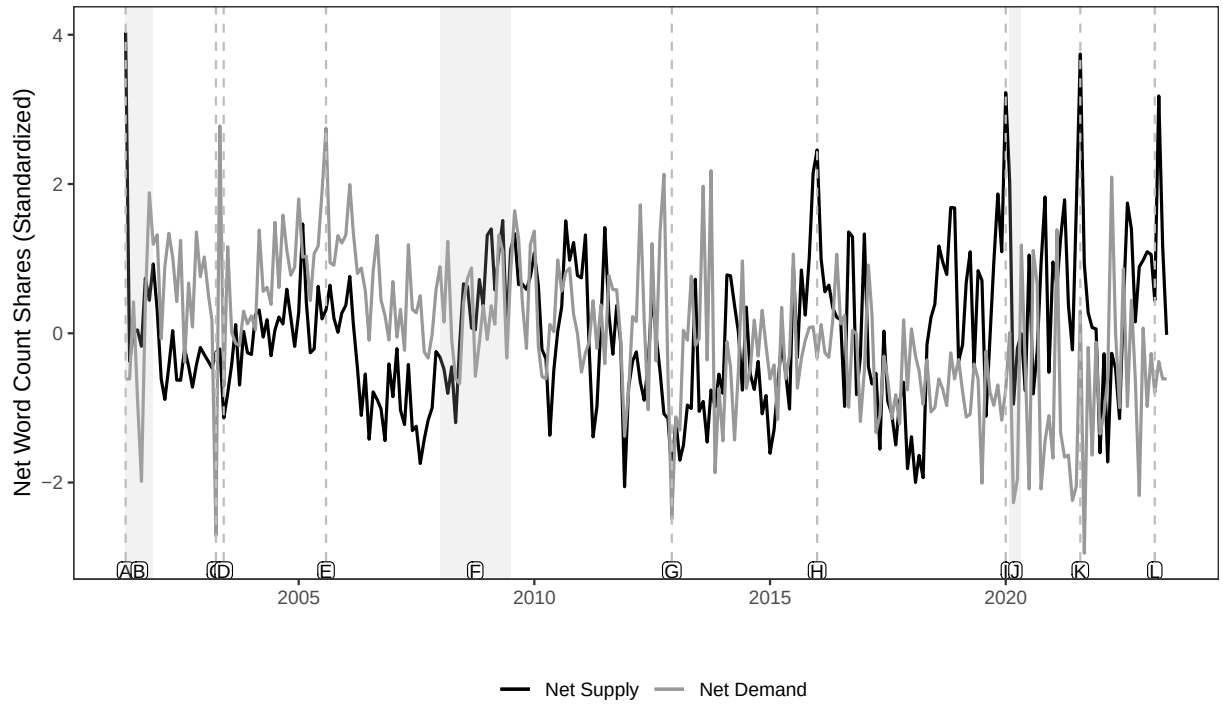
Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the wheat sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] US wheat production 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 A.3: Corn index: Net supply and net demand



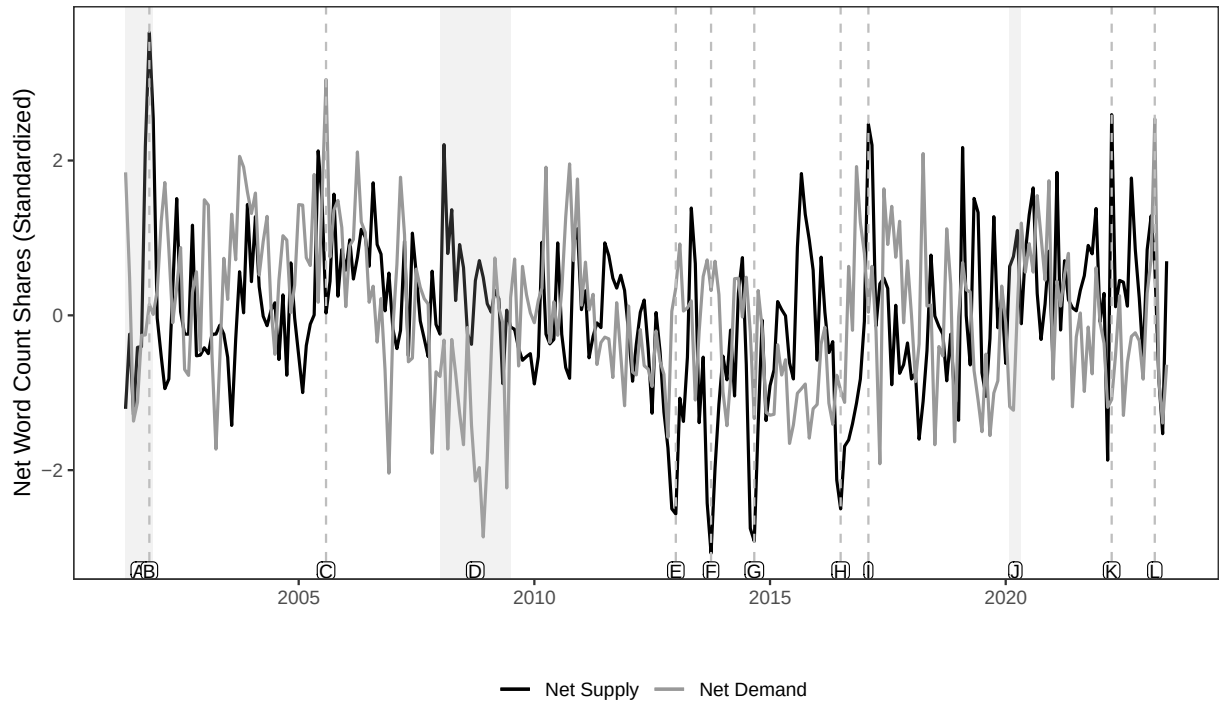
Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the corn sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] Weak US export 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 A.4: 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)

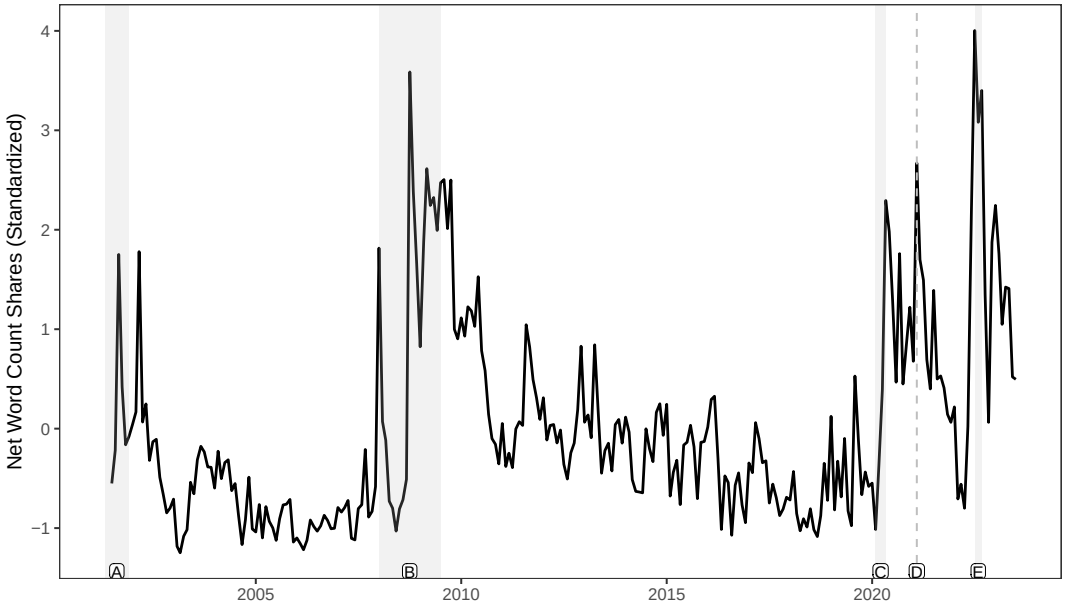
Figure A.5: Copper index: Net supply and net demand



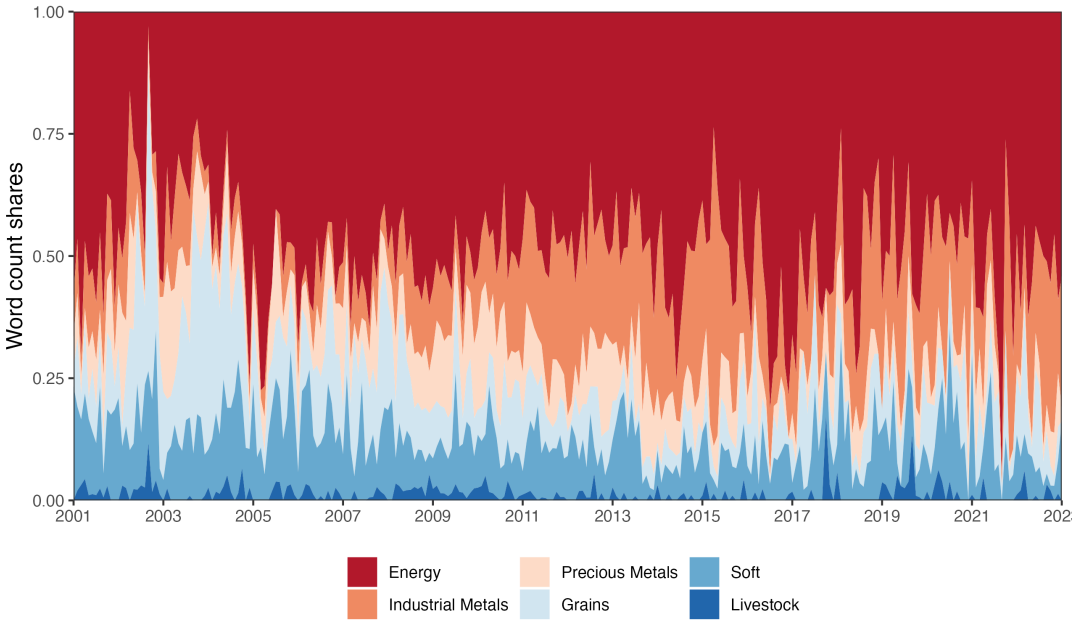
Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2023. The bars map a number of important developments in the copper sector that are in the top 1% (two-tailed) of events. These events respectively are: [A] US Recession and 9/11, [B] roducers in 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

Figure A.6: Business cycle

A. Occurrence in news articles



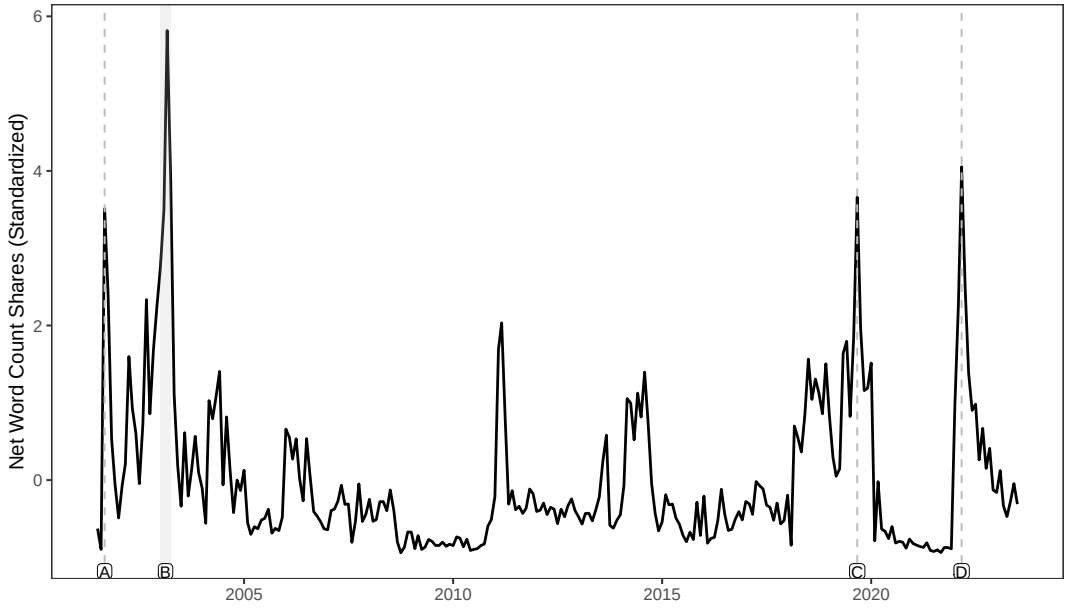
B. Decomposition across commodity categories



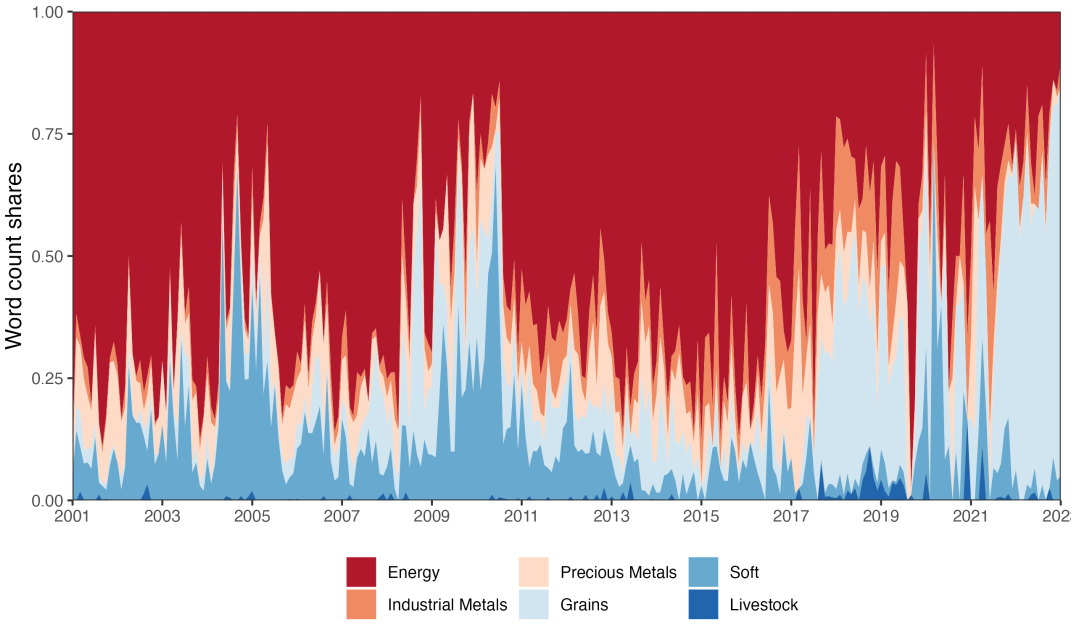
Note: This figure plots the Business Cycle Theme for the period between 2001 and 2023. The bars map a number of important economic developments (i.e. in the top 1% (two-tailed) of events) reflecting the state of the business cycle. These events respectively are: [A] US Recession and 9/11, [B] Global 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 A.7: Geopolitical risk

A. Occurrence in news articles



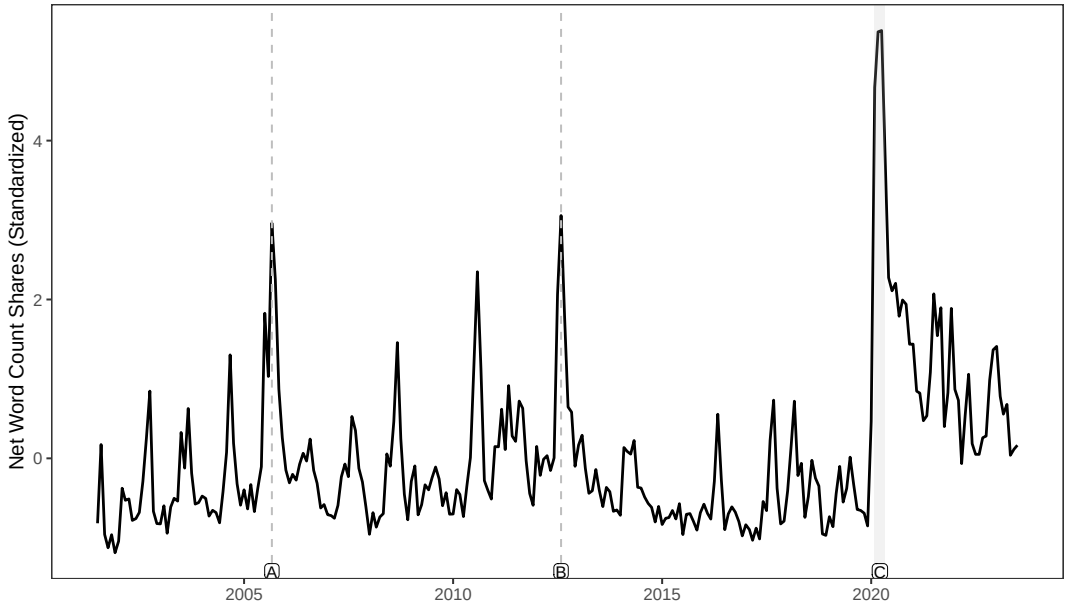
B. Decomposition across commodity categories



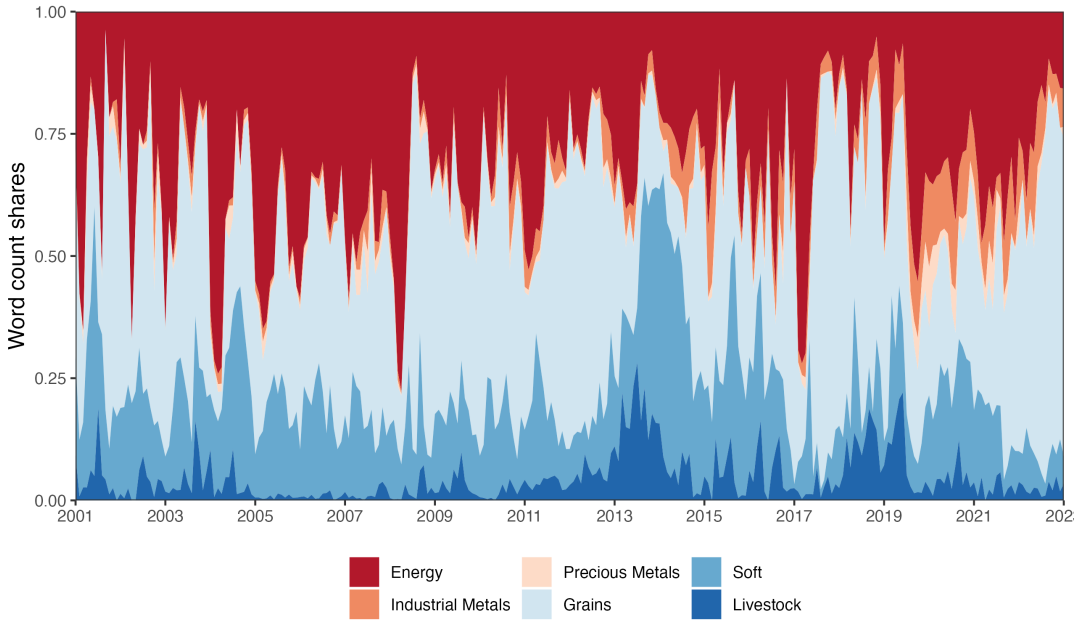
Note: This figure plots the Geopolitical Risk Theme for the period between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) geopolitical events. These events respectively are: [A] 9/11 attacks, [B] 2003 Iraq invasion, [C] China-United States trade war, [D]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

Figure A.8: Natural disaster

A. Occurrence in news articles



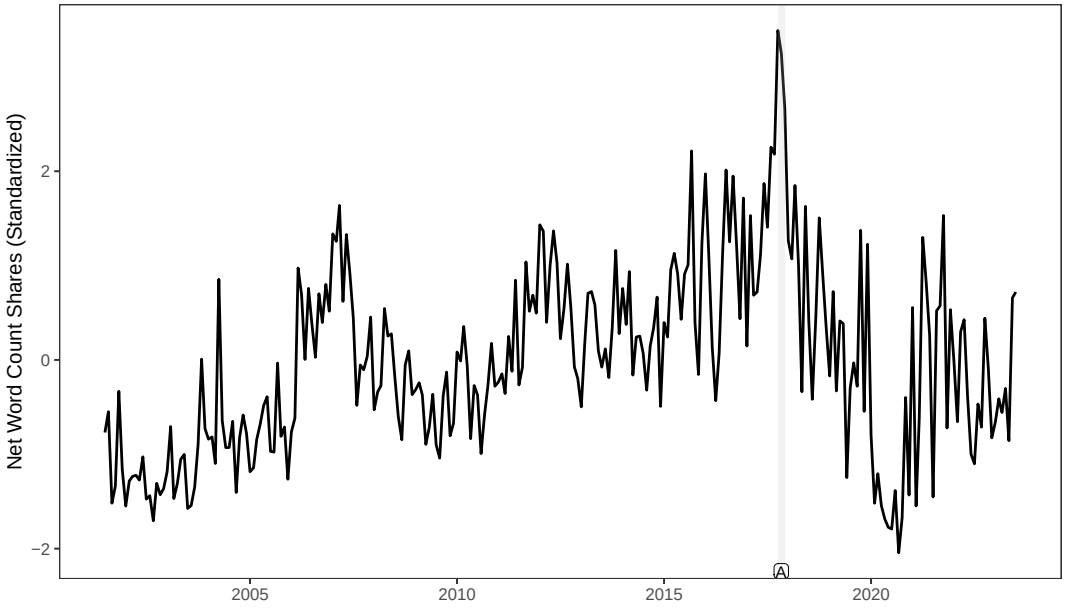
B. Decomposition across commodity categories



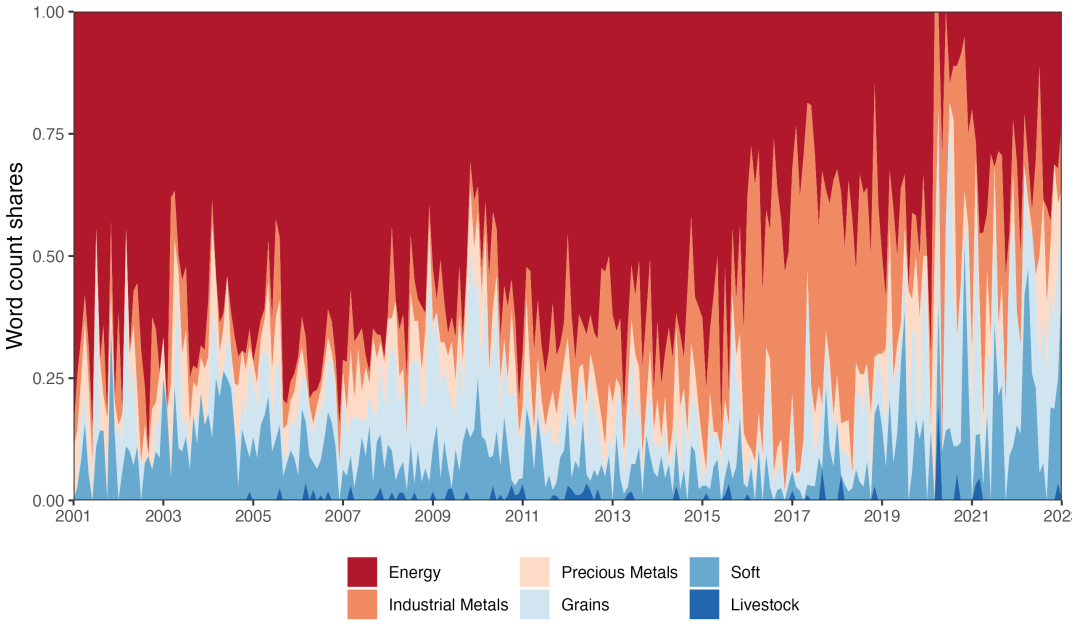
Note: This figure plots the Natural Disasters Theme for the period between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) natural disasters. These events respectively are: [A] 2005 Atlantic hurricane season (Katrina, Emily, Rita, Wilma), [B] US Summer 2012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

Figure A.9: Climate change

A. Occurrence in news articles

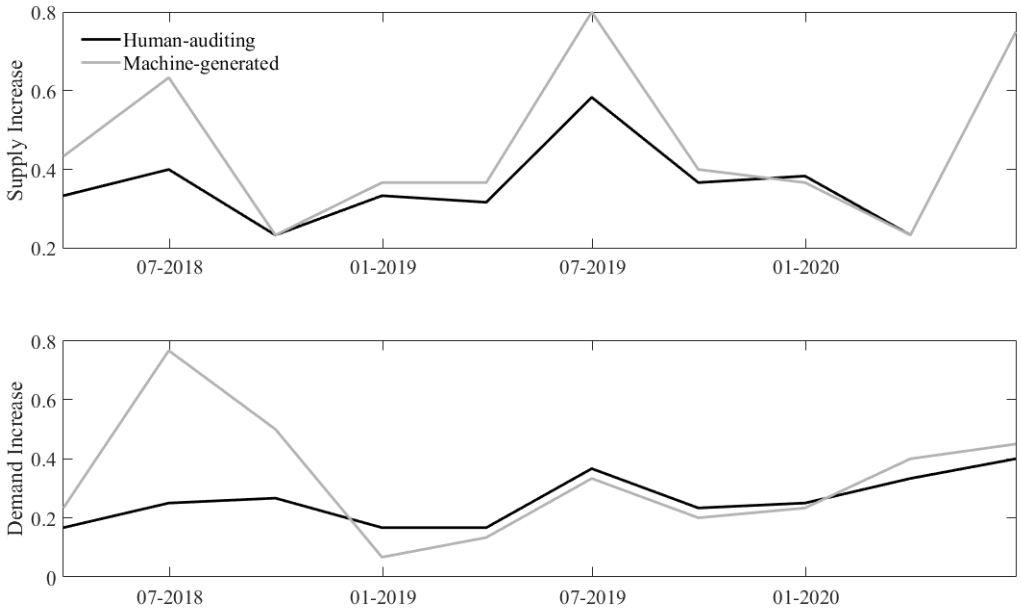


B. Decomposition across commodity categories



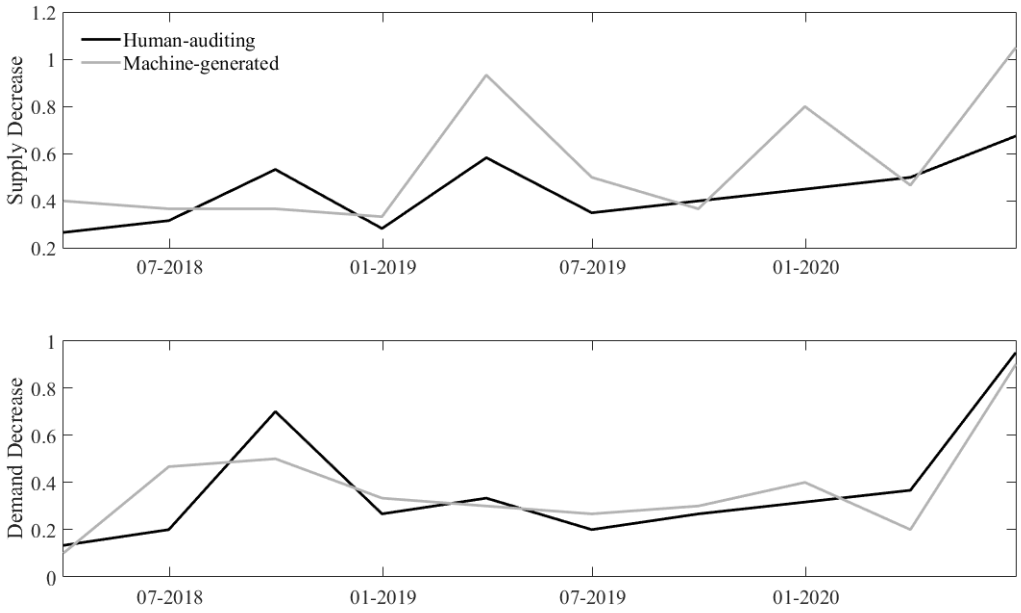
Note: This figure plots the Climate Change Theme between 2001 and 2023. The bars map a number of important (i.e. in the top 1% (two-tailed) of events) climate change events. [A] The U.S. Environmental Protection Agency requires fuel companies to blend slightly more renewable fuels into the nation's fuel 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 A.10: Human- and computer-based increase indices



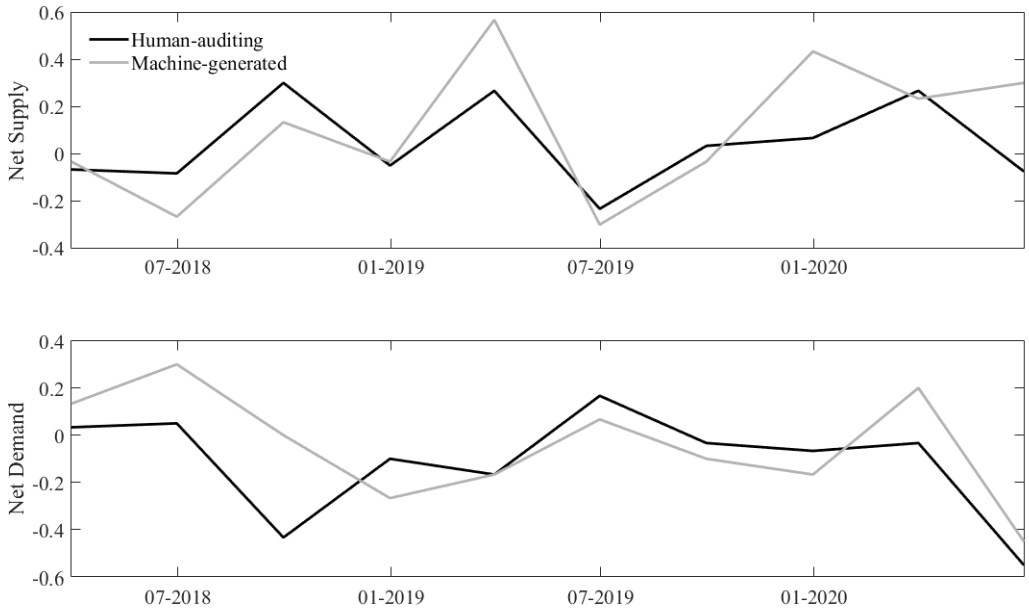
Note: The figure compares our algorithm-generated indices with a human-coded benchmark at the quarterly frequency. The benchmark is built from a fixed number of articles drawn at random each month and independently classified by at least two auditors, following the procedure described in Section 5.7. The plotted series are the supply-increase and demand-increase indices constructed by the two methods. Across the four supply and demand components, the quarterly correlation between the human- and algorithm-based indices ranges from 0.64 (demand decrease) to 0.85 (demand increase). Please verify against the actual plot: this note assumes the figure shows the supply-increase and demand-increase indices for both the human and algorithmic methods.

Figure A.11: Human- and computer-based decrease indices



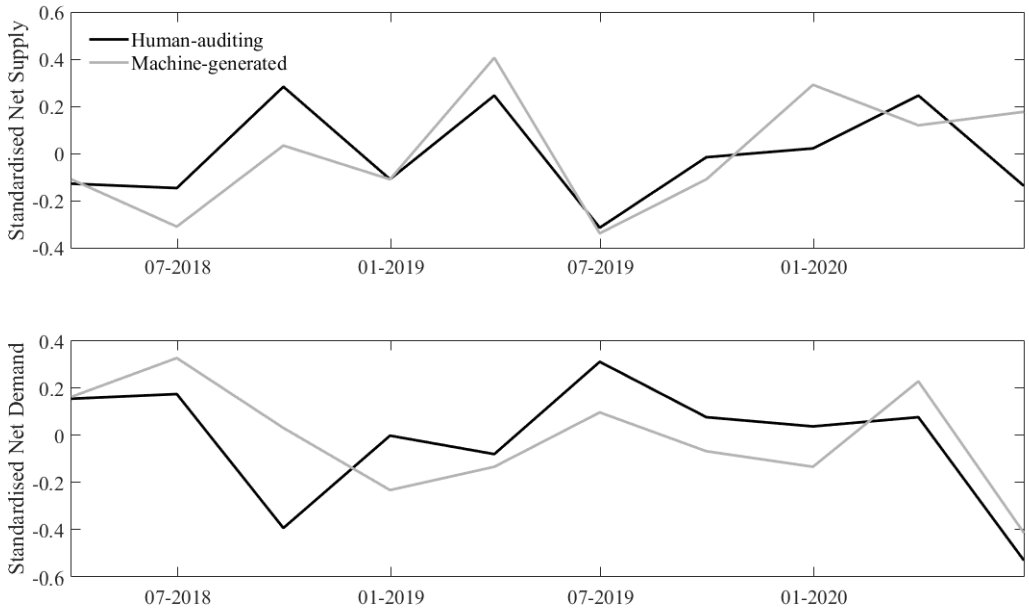
Note: The figure plots the human- and algorithm-based supply-decrease and demand-decrease indices at the quarterly frequency. See the note to Figure A.10 and Section 5.7 for the construction of the human-coded benchmark. Please verify against the actual plot: this note assumes the figure shows the supply-decrease and demand-decrease indices.

Figure A.12: Human- and computer-based net indices



Note: The figure plots the human- and algorithm-based net supply and net demand indices at the quarterly frequency, where each net index combines the corresponding increase and decrease components. See the note to Figure A.10 and Section 5.7 for the construction of the human-coded benchmark. Please verify against the actual plot: this note assumes the figure shows net supply and net demand indices, and confirm the sign convention for the net index (increase minus decrease, versus the contractions-net-of-expansions convention used in the macro section).

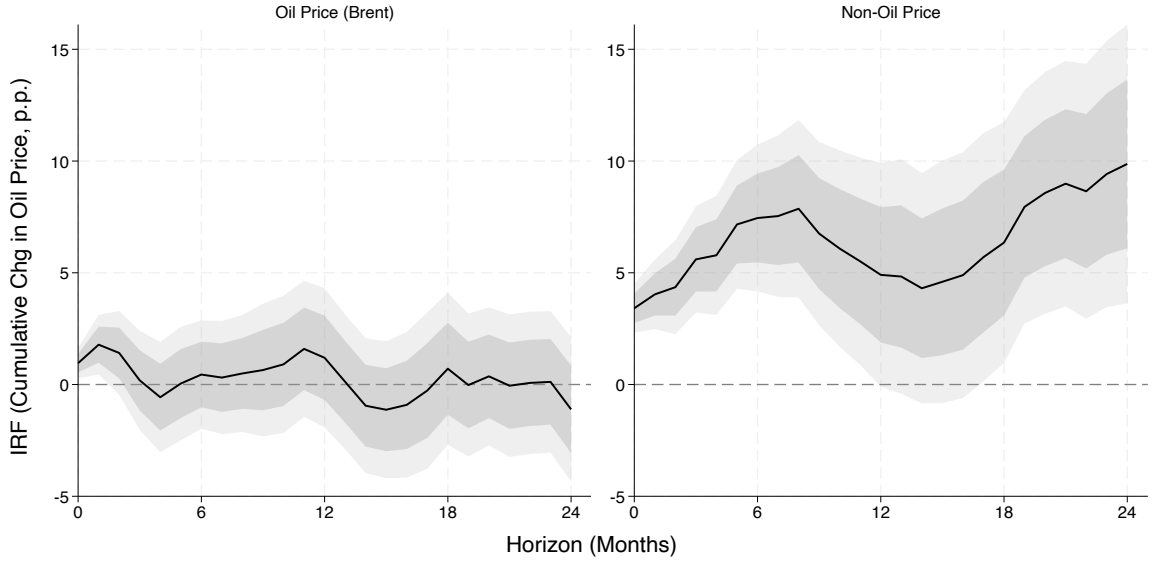
Figure A.13: Human- and computer-based standardized net indices



Note: The figure repeats Figure A.12 for the net indices standardized to unit standard deviation. See the note to Figure A.10 and Section 5.7 for the construction of the human-coded benchmark. Please verify against the actual plot: this note assumes the figure shows the standardized net supply and net demand indices.

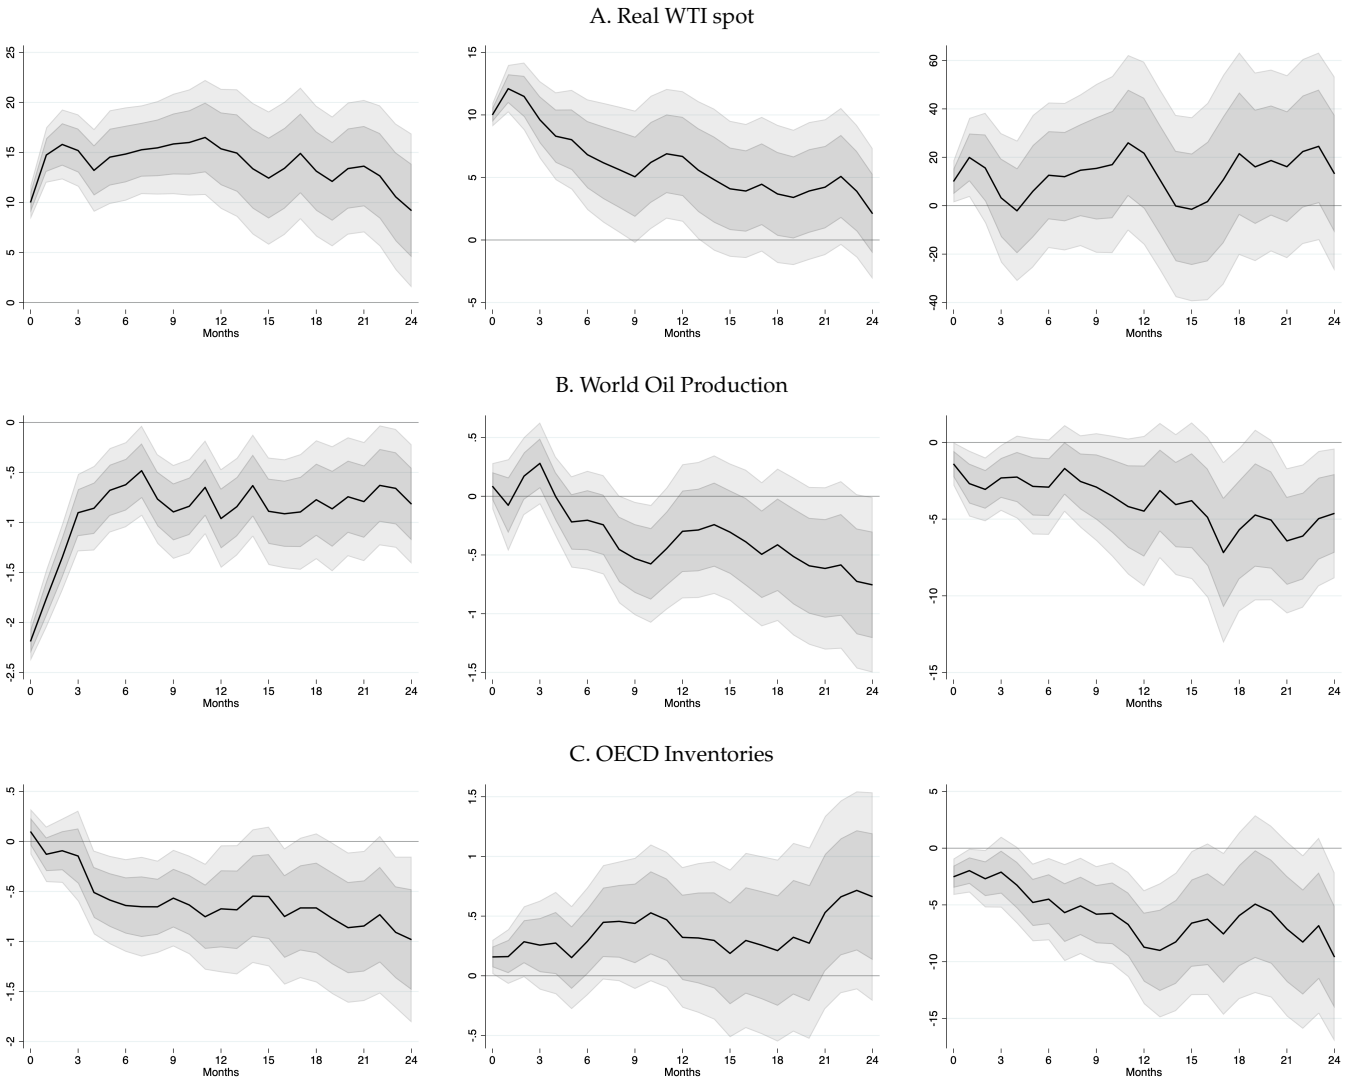
Figure A.14: Long-run response of commodity prices to supply distrubrances

(left: Brent; right: weighted non-oil commodity basket)



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.

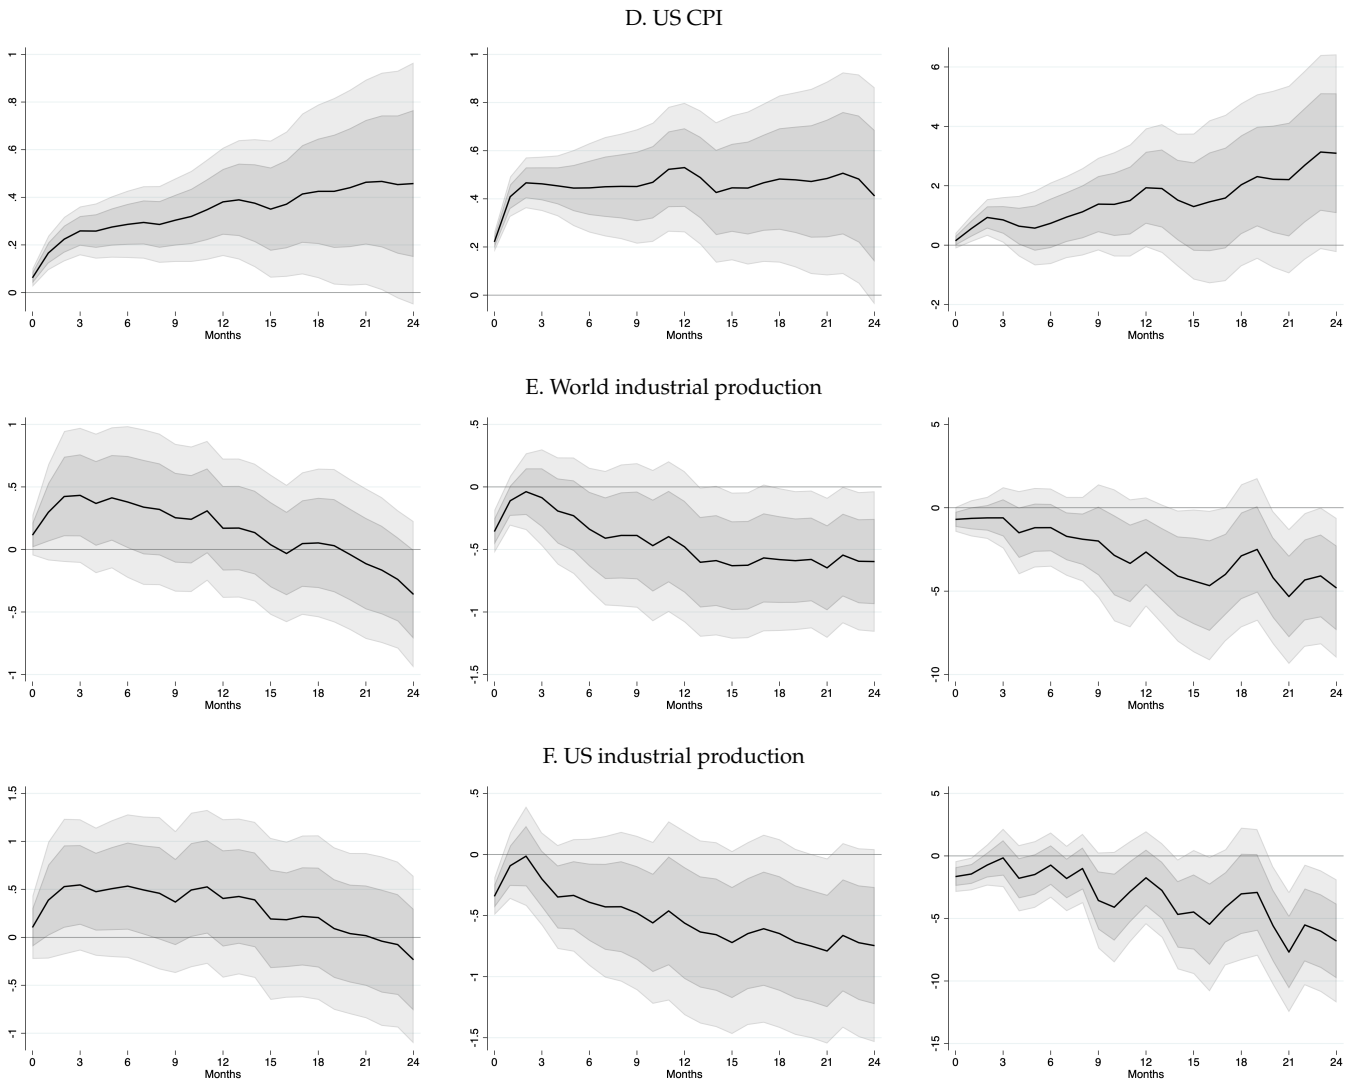
Figure A.15: Comparison to Existing Indicators of Oil Supply Shocks (part 1)



(I) Baumeister & Hamilton, 2019 Kanzig, 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 Oil Production, and OECD Inventories, 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 A.16: Comparison to Existing Indicators of Oil Supply Shocks (part 2)

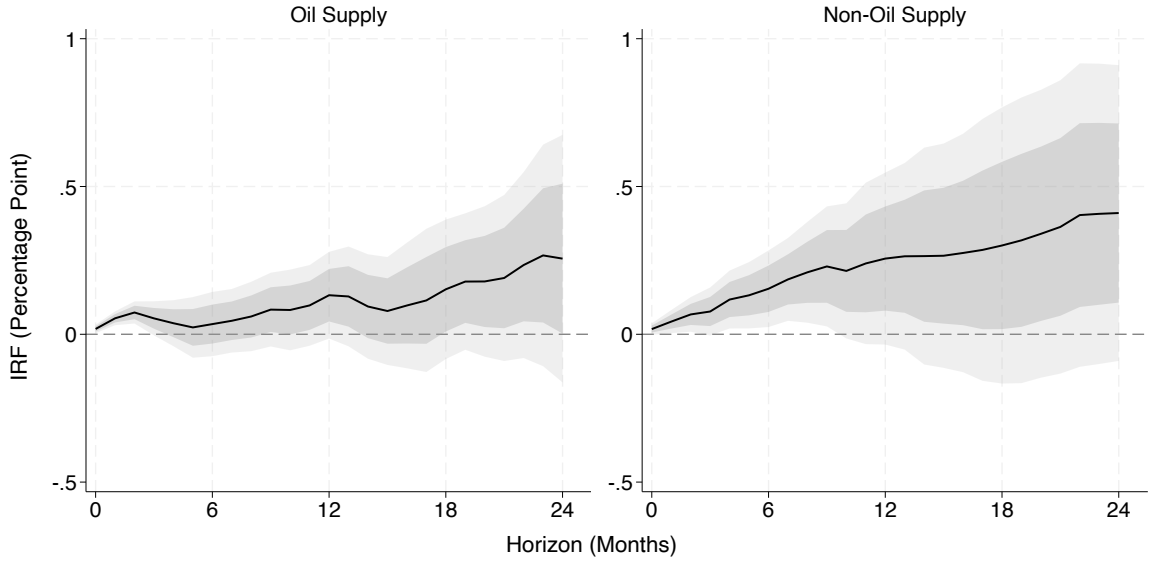


(I) Baumeister & Hamilton, 2019 Kanzig, 2021 (III) LMPRP, 2025.

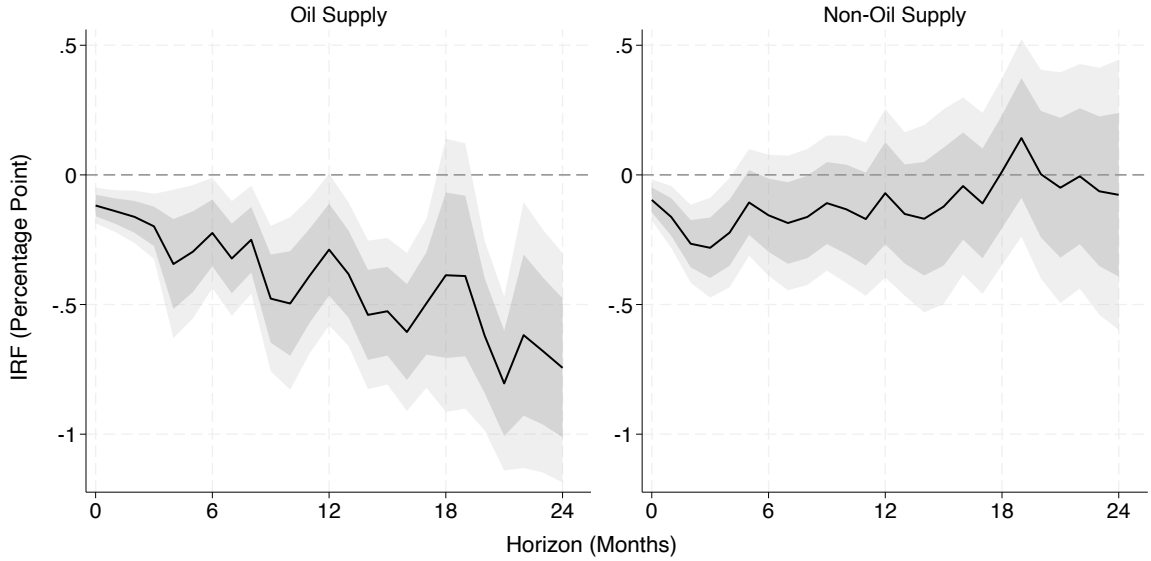
Note: This figure shows the impulse responses to a one standard deviation oil supply shock for the US CPI and World and US Industrial Production, employing the oil supply shock series of Baumeister and Hamilton (2019), Känzig (2021) and a "surprise" series, calculated from our Net Supply Crude Oil Index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

Figure A.17: Local projections: US impulse responses to oil and non-oil supply shocks

A. US CPI (left: oil; right: non-oil)



B. US industrial production (left: oil; right: non-oil)

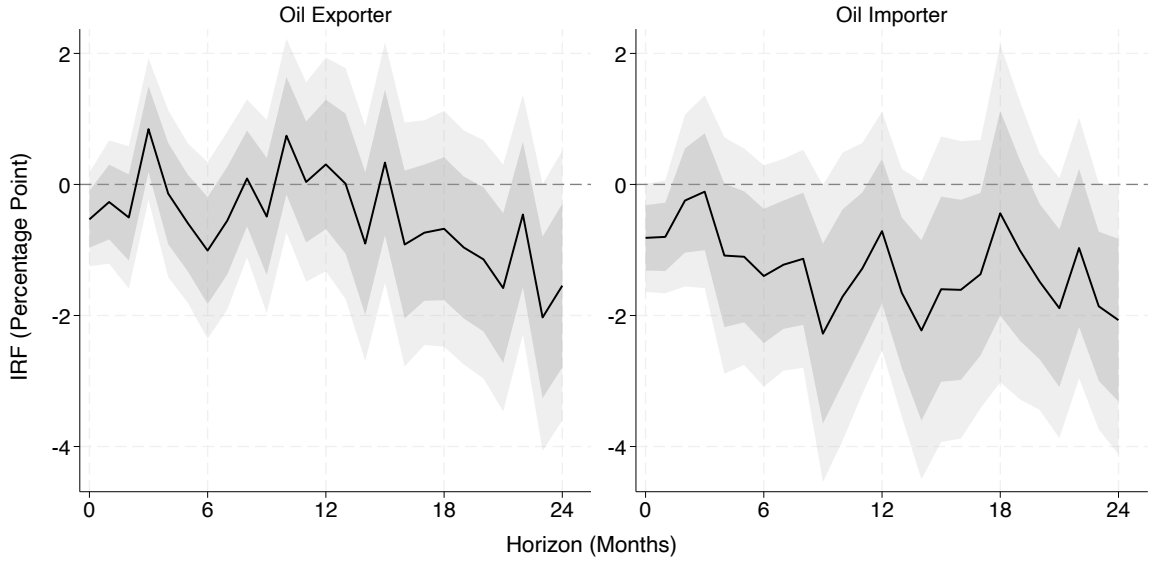


Note: This figure shows the impulse responses of US CPI and US industrial production to a one-standard-deviation supply contraction, employing the oil supply index (left) and the non-oil production-weighted supply index (right), constructed as AR(20) residuals aggregated to monthly frequency.

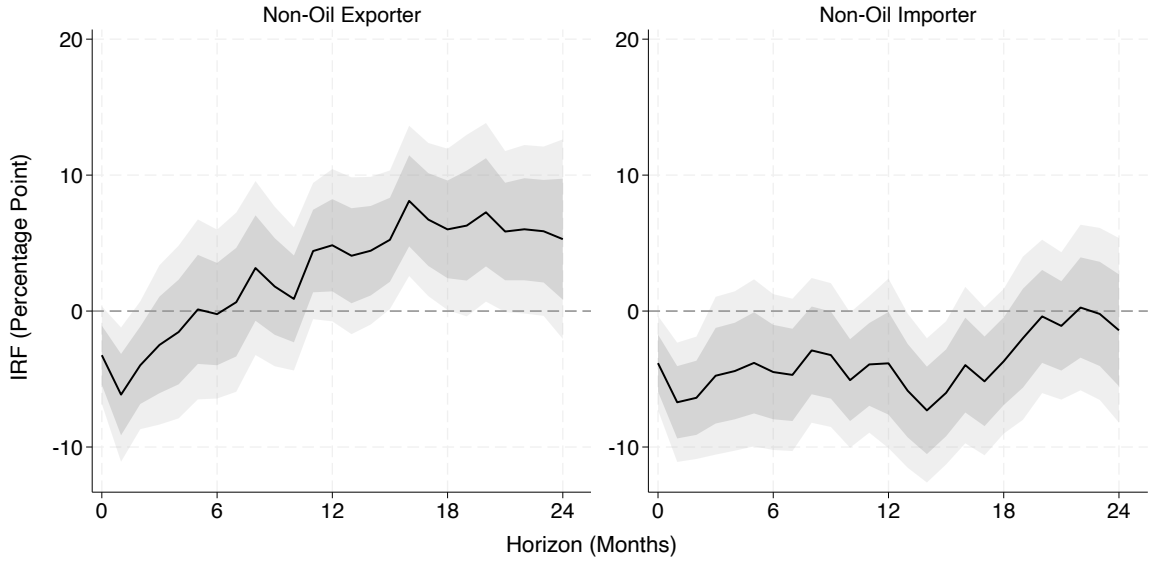
is this helping you to show? Can you also show the composite? If the composite looks good, maybe consider showing it next to Fig 5?

Figure A.18: Impact of commodity shocks on production: commodity net importers vs net exporters

A. Net Oil Importers and Exporters, Industrial Production



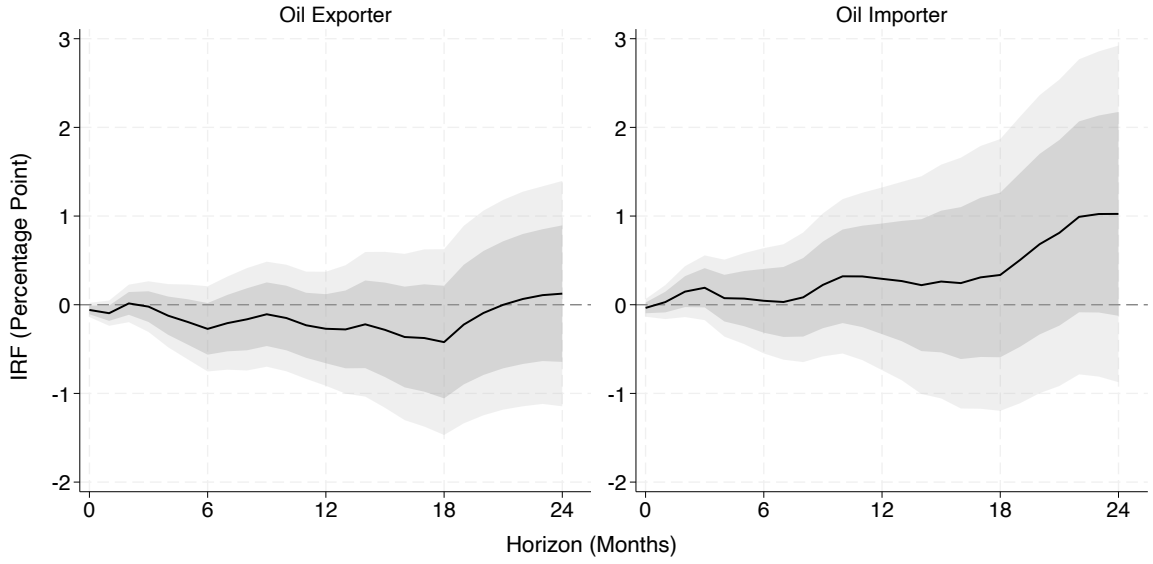
B. Net Non-oil Importers and Exporters, Industrial Production



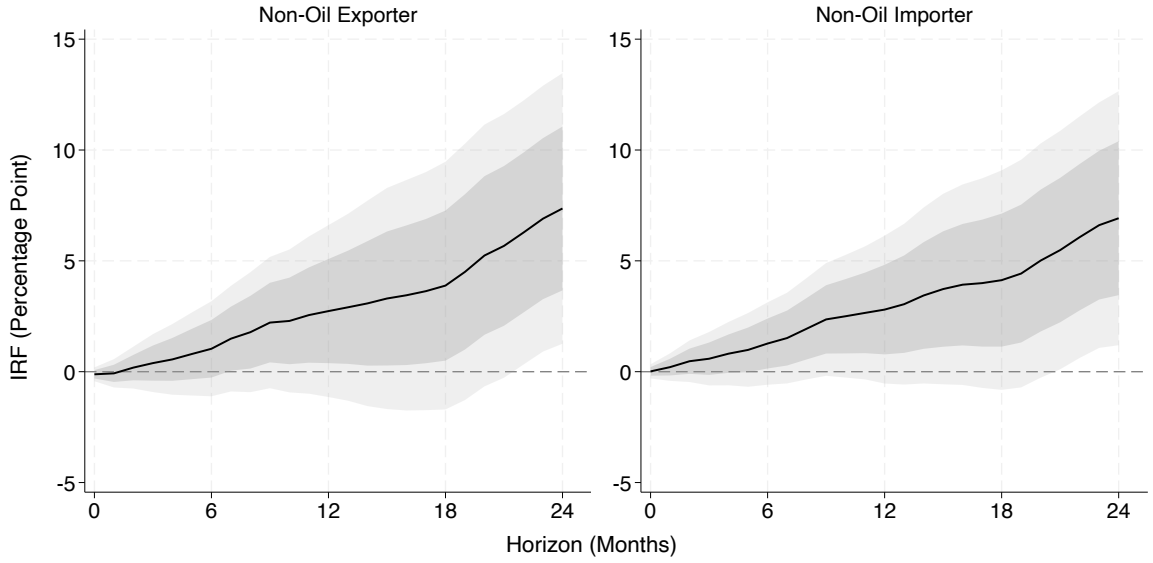
can we see composite results too?

Figure A.19: Impact of commodity shocks on inflation: commodity net importers vs net exporters

A. Net Oil Importers and Exporters, Inflation



B. Net Non-oil Importers and Exporters, Inflation



can we see composite results too?