

Not All Oil Narrative Shocks Are Alike: The Macroeconomic Effects of News-Driven Sentiment and Uncertainty Shocks

Gabriel A. Abeyie

Cameron School of Business
University of North Carolina Wilmington

April 21, 2026

Abstract

Oil-market narratives affect the macroeconomy not only through the direction of the news they convey, but also through the uncertainty they create. I construct monthly indices of oil narrative sentiment and uncertainty using coverage from *The Wall Street Journal* and *The New York Times* from 1953 to 2022, and use type-specific narrative measures as external instruments in a Proxy-SVAR. A positive sentiment shock raises real oil prices persistently and increases world industrial production and world oil production, consistent with a demand-side information channel. By contrast, an uncertainty shock reduces world oil production and world industrial production while generating only a muted real-oil-price response, consistent with a supply-side real-options channel. Together, the two shocks imply a stagflationary configuration in which oil prices and inflation rise while real activity responds weakly. The results show that separating directional tone from uncertainty is essential for understanding how oil news transmits to the macroeconomy.

Keywords: oil-market narratives; sentiment; uncertainty; text-as-data; external instruments; oil-macroeconomy; Proxy-SVAR; stagflation

JEL Classification: C32, E32, Q41, Q43, G14

1 Introduction

Recent turbulence in the oil market has renewed interest in the longstanding question of how oil prices affect the macroeconomy. But oil price movements reflect more than contemporaneous physical disruptions or realized shifts in global demand. They also reflect changing beliefs about future market conditions. News about wars in oil-producing regions, OPEC decisions, speculative activity, refinery outages, and energy-transition policy can alter how market participants interpret future oil scarcity, demand, and investment conditions before any physical barrel changes hands (Kilian and Murphy, 2014; Hamilton, 2009). Yet the macroeconomic consequences of such oil-market news need not be uniform. Some narratives provide directional information that shifts beliefs about future market conditions. Others primarily increase ambiguity about the future path of the oil market. This paper argues that this distinction matters for macroeconomic transmission.

The structural oil-market literature has made substantial progress in distinguishing between supply shocks, global demand shocks, and precautionary or inventory-demand shocks (Kilian, 2009; Kilian and Murphy, 2014; Baumeister and Hamilton, 2019). But these frameworks do not directly separate oil-market news that moves expectations in a directional way from oil-market news that primarily raises ambiguity. As a result, episodes of intense oil-market coverage that differ in economic content may be absorbed into broad residual demand-side or uncertainty-related disturbances. This omission is important because directional narrative shifts and ambiguity shocks imply different adjustment mechanisms, different responses of production and activity, and potentially different interpretations of oil-price movements.

This paper studies whether oil-market narratives have macroeconomic effects, and whether those effects differ depending on whether narratives become more *directional* or more *uncertain*. I distinguish between two dimensions of oil-market news. The first is *narrative sentiment*, which captures whether oil-market coverage is directionally bullish or bearish. The second is *narrative uncertainty*, which captures the degree of ambiguity, unpredictability, or lack of clarity embedded in that coverage. The key hypothesis is that these two dimensions are economically distinct and propagate through different channels. A rise in bullish narrative sentiment may shift expectations about future scarcity or future demand, increasing the willingness of market participants to hold inventories and generating a demand-side information effect (Beaudry and Portier, 2006; Barsky and Kilian, 2002). By contrast, a rise in narrative uncertainty may raise the option value of waiting and delay partially irreversible production and investment decisions (Bernanke, 1983; Dixit and Pindyck, 1994; Bloom, 2009), generating a supply-side delay effect.

I measure these two dimensions using new monthly indices of oil-market narrative

sentiment and uncertainty constructed from *Wall Street Journal* and *New York Times* articles from 1953 to 2022. A domain-specific oil dictionary identifies more than 141,000 oil-relevant articles from a raw corpus of nearly two million business and market news stories. Each article is scored separately for directional tone and uncertainty content and then aggregated to the monthly frequency. I also assign oil-related coverage to five narrative types: geopolitical conflict, supply disruption, OPEC coordination, speculation/financialization, and energy transition. This produces both aggregate narrative indices and type-specific series over seven decades of oil-market history.

The data show that narrative sentiment and narrative uncertainty are not simply two versions of the same oil-news factor. Although the series are correlated, they behave differently across major historical episodes and are associated with different macroeconomic response patterns. During the 1986 OPEC collapse, for example, sentiment is strongly negative while uncertainty remains low, consistent with a market that largely understood the directional implications of Saudi Arabia's shift away from swing-producer behavior.¹ During the Arab Spring, by contrast, bullish sentiment coincides with sharply elevated uncertainty, reflecting upward pressure from perceived supply risk alongside substantial ambiguity about the future course of geopolitical events.² These differences indicate that oil-market narratives contain separable information about directional beliefs and ambiguity rather than loading on a single latent news factor.

I then estimate the macroeconomic effects of these two narrative dimensions in a Proxy-SVAR (Mertens and Ravn, 2013; Stock and Watson, 2018). The objective is to identify shocks to aggregate oil-market narratives rather than physical oil shocks directly. Four type-specific indices serve as external instruments for the aggregate narrative series: geopolitical events, speculative positioning, supply disruptions, and energy transition policy. I exclude OPEC-related narratives from the baseline instrument set because OPEC's institutional role is to respond to oil-market conditions in real time, creating a contemporaneous feedback relationship with oil prices that has no close analogue in the other four categories.³ The remaining instruments capture variation from economically distinct sources of oil-market news, each operating through a different channel. In this respect, the approach differs from high-frequency event-window designs such as Känzig (2021) and from event-based media-shock instruments such as Gavriilidis et al. (2026).⁴ Instead, it exploits heterogeneity within the

¹In late 1985, Saudi Arabia abandoned its role as swing producer and began competing for market share, precipitating a price collapse from roughly \$27 to below \$10 per barrel in 1986. See Gately (1986) and Hamilton (2013).

²The Arab Spring (2010–2012) disrupted production in Libya and raised concerns about contagion to major Gulf producers. See Hamilton (2011) for a discussion of the oil-market implications.

³Appendix C.3 provides formal diagnostic evidence for this exclusion and reports the OPEC-inclusive specification as a robustness check.

⁴See also Caldara and Iacoviello (2022) for the Geopolitical Risk index and Baker et al. (2016) for the

narrative environment itself. Because identification rests on the plausibility of the underlying components, I report instrument-specific first-stage diagnostics following [Montiel Olea et al. \(2021\)](#), forecastability and serial-correlation tests, over-identification tests ([Sargan, 1958](#); [Hansen, 1982](#)), and exclusion exercises. Confidence bands are constructed using the moving block bootstrap following [Jentsch and Lunsford \(2022\)](#).

The empirical results reveal two distinct macroeconomic signatures. A bullish narrative sentiment shock raises the real oil price persistently, increases world industrial production, raises world oil production, and gradually passes through to consumer prices. The positive co-movement of oil prices and activity is difficult to reconcile with an adverse supply interpretation ([Kilian, 2009](#)) and is instead consistent with a demand-side information channel in which bullish oil-market narratives shift beliefs about future market tightness or future demand. By contrast, a narrative uncertainty shock reduces world oil production and world industrial production, while producing only a muted and statistically weak response of the real oil price and consumer prices. This pattern is consistent with a supply-side delay channel: heightened ambiguity discourages adjustment in drilling and production, suppressing real activity while leaving the net price effect weak because lower demand and tighter supply work in opposite directions ([Bloom, 2009](#); [Elder and Serletis, 2010](#)). Additional evidence from drilling activity supports this interpretation, as narrative uncertainty lowers the Baker Hughes rig count whereas sentiment raises it.

A further implication of the two-channel framework is that simultaneous sentiment and uncertainty shocks can generate a macroeconomic pattern that neither shock produces alone. When bullish oil-market sentiment and elevated uncertainty occur together, oil prices rise because stronger demand expectations and supply concerns move in the same direction, while oil production remains weak because the incentive to expand output is offset by the incentive to delay investment. The result is a stagflationary combination of higher oil prices, subdued production, weaker real activity, and higher inflation. This perspective helps interpret episodes such as the Gulf War (1990–91), the Arab Spring (2011), and the Russia-Ukraine conflict (2022), in which oil-market narratives likely became both more bullish and more uncertain. The standard three-shock framework of [Kilian \(2009\)](#) cannot separate these dimensions because it does not distinguish between the directional and ambiguity components of oil-market news.

These findings imply that oil-market narratives contain macroeconomically relevant variation beyond standard physical oil shocks and that separating directional tone from ambiguity is important for understanding how oil news transmits to the broader economy. More broadly, the results suggest that standard oil-shock frame-

Economic Policy Uncertainty index, both of which use newspaper text to construct aggregate indices but do not decompose them into type-specific instruments for structural identification.

works (Kilian, 2009; Baumeister and Hamilton, 2019) may conflate two distinct forms of oil-market news: one that behaves like a demand-side expectations shock and another that behaves more like a supply-side uncertainty shock. Distinguishing between them helps explain why some episodes of intense oil-market coverage generate price-output co-movement, while others primarily suppress production and activity without producing a comparably strong price response.

The paper contributes to several literatures. First, it contributes to the structural oil-market literature following Kilian (2009) by showing that oil-market narratives contain macroeconomically relevant variation not well summarized by physical supply, global demand, or inventory-demand shocks alone. Second, it contributes to narrative economics (Shiller, 2017) by showing that a single domain-specific news corpus can recover distinct dimensions of narrative content with different macroeconomic consequences. Third, it contributes to the growing literature using text-based measures to identify economically meaningful shocks (Baker et al., 2016; Caldara and Iacoviello, 2022; Shapiro et al., 2022; Larsen and Thorsrud, 2019) by showing how decomposition within a corpus can generate multiple proxies for a common aggregate narrative disturbance. In this sense, the paper is not only about measuring oil-market narratives, but also about identifying how different forms of narrative variation map into different macroeconomic transmission channels.

The paper also contributes methodologically. Many text-based measures are treated as reduced-form indicators whose economic content is difficult to interpret. The approach here shows how decomposition within a corpus can be used to generate multiple proxies for a common aggregate narrative shock, conceptually related to multi-instrument designs in the IV literature (Angrist and Krueger, 1991; Stock and Yogo, 2005). In this respect, the strategy is conceptually related to multi-component instrument designs: identification hinges not on a single black-box index, but on the plausibility of the underlying components and on whether the results are robust to excluding problematic ones. In this application, the diagnostics indicate strong first-stage relevance in both systems and broadly support the validity of the type-specific instruments.

The remainder of the paper proceeds as follows. Section 2 develops the conceptual framework. Section 3 describes the newspaper corpus, article scoring, and narrative classification. Section 4 presents the Proxy-SVAR and the instrument construction. Sections 5.1 and 5.2 report the main impulse responses, Section 5.3 examines the combined stagflationary pattern, Section 5.4 presents the rig-count evidence for the real options channel, and Sections 5.5 and 5.6 report the historical and variance decompositions. Section 6 presents robustness checks. Section 7 concludes.

2 Conceptual Framework

The paper's central claim is that oil-market narratives affect the macroeconomy through two distinct dimensions of news content. One dimension is *narrative sentiment*, which captures whether coverage points in a bullish or bearish direction for oil-market conditions. The other is *narrative uncertainty*, which captures how much that coverage leaves the future path of the oil market unclear. The key idea is that these two dimensions need not operate through the same economic channel.

The structural oil-market literature, beginning with [Kilian \(2009\)](#), focuses on shocks that originate in the *physical* oil market: realized supply disruptions, shifts in global demand, and changes in precautionary demand driven by concerns about future scarcity (see also [Kilian and Murphy, 2014](#); [Baumeister and Hamilton, 2019](#)). These shocks are traced through observable outcomes such as production, consumption, inventories, and prices. But the same events that move the physical oil market also generate narratives that can shape beliefs before any physical quantity adjusts. A conflict in the Middle East, for example, can matter not only because it threatens current supply, but also because it changes how market participants think about future scarcity, volatility, and investment conditions ([Hamilton, 2003](#)).

What existing frameworks do not separate cleanly is oil-market news that moves expectations in a direction from oil-market news that mainly makes the future harder to assess.⁵ That distinction matters because the two forms of narrative variation imply different adjustment mechanisms, especially for production and investment. This paper does not identify physical oil shocks directly. Instead, it identifies shocks to aggregate oil-market narratives and interprets their macroeconomic effects through the joint responses of oil prices, oil production, real activity, and consumer prices. The purpose of this conceptual framework is therefore to explain why sentiment and uncertainty should be treated separately and to derive the qualitative predictions that guide the empirical analysis.

2.1 Narrative Sentiment and the Demand-Side Information Channel

Narrative sentiment captures the directional framing of oil-market conditions. When coverage becomes more bullish, emphasizing stronger demand, tighter supply, recovering activity, or persistent price pressure, market participants may revise upward their expectations of future oil scarcity or future oil demand. This mechanism is closely related to the oil-specific demand channel in [Kilian \(2009\)](#): forward-looking agents become more willing to hold inventories and bid up current prices in anticipation of

⁵The closest precedents are [Jo \(2014\)](#), who distinguish oil-price uncertainty from oil supply shocks, and [Baumeister et al. \(2018\)](#), who separate expected and unexpected components of energy-price changes. Neither, however, extracts distinct narrative dimensions from text data.

tighter conditions ahead.

This channel does not require that media narratives reveal fundamentally new information in the strict rational-expectations sense. Narratives can matter because they affect salience, framing, and repetition. That view is consistent with narrative economics (Shiller, 2017), with models of limited attention (Sims, 2003), and with information rigidities (Mankiw and Reis, 2002). In these environments, agents update beliefs gradually, and the way information is presented affects how quickly and in what direction expectations adjust. A newspaper article about record speculative long positions in crude oil, for example, may not reveal facts that are unavailable elsewhere, but it may still change how those facts are interpreted and weighted by market participants (Tetlock, 2007). At the aggregate level, repeated shifts in the directional tone of oil-market coverage can therefore move beliefs about future market conditions and generate macroeconomically relevant information shocks.

This mechanism implies a clear empirical prediction. If a bullish narrative sentiment shock works mainly through a demand-side information channel, it should raise the real price of oil and, over the relevant horizon, generate positive co-movement between oil prices and real activity. Oil production may also rise with a delay as producers respond to stronger price incentives, and consumer prices may increase gradually through oil-price pass-through (Hamilton, 2013; Blanchard and Galí, 2007).

2.2 Narrative Uncertainty and the Supply-Side Delay Channel

Narrative uncertainty captures the degree of ambiguity surrounding future oil-market conditions. When coverage emphasizes conflicting signals, unclear outcomes, geopolitical instability, or unpredictable policy and market developments, the narrative environment may widen the perceived range of future oil-market outcomes. Such a rise in uncertainty can matter even when the directional implications of the news remain unclear.

A natural mechanism linking this uncertainty to macroeconomic outcomes is the wait-and-see logic developed by Bernanke (1983) and Bloom (2009). When decisions are costly to reverse, greater uncertainty raises the option value of waiting. In the oil sector, this logic is especially relevant because drilling, exploration, and infrastructure investment are capital-intensive and partly irreversible (Dixit and Pindyck, 1994). Kellogg (2014) shows that uncertainty about future oil prices can reduce drilling activity, consistent with this mechanism.⁶ This mechanism implies that narrative uncertainty operates as a supply-side disturbance, not because it directly reduces productive ca-

⁶Agerton et al. (2017) extend this evidence by documenting that oil-price uncertainty affects employment decisions in U.S. oil-producing counties. Elder and Serletis (2010) find that oil-price uncertainty depresses aggregate investment and output.

capacity, but because it delays the adjustment of partially irreversible production decisions.

In the present context, the idea is not that narrative uncertainty is identical to firm-level investment uncertainty, but that heightened ambiguity in oil-market narratives may generate a macroeconomic environment in which producers delay action. If so, a narrative uncertainty shock should depress oil production persistently, weaken real activity, and exert only an ambiguous effect on the real oil price. The reason is that two forces operate in opposite directions: weaker activity reduces oil demand and pushes prices down, while delayed production and investment constrain supply and push prices up. Consumer prices may therefore respond only modestly, if at all.

2.3 Two Dimensions of the Same Narrative Environment

A central feature of the paper is that both dimensions are extracted from the same newspaper corpus. The same article can contain directional content and uncertainty content at the same time. For example, a report on an OPEC production cut may be bullish in tone for oil prices while also emphasizing uncertainty about compliance, duration, or the broader geopolitical response. Conversely, coverage can be sharply directional with little ambiguity, or highly uncertain without a strong directional signal.⁷

This distinction is useful because it implies that sentiment and uncertainty are not simply alternative labels for the same narrative shock. They describe different features of the narrative environment. The empirical analysis therefore treats them as separate aggregate objects and asks whether shocks to each dimension generate different macroeconomic response patterns.

Table 1 illustrates the distinction. The four quadrants combine directional tone (bullish or bearish) with the level of ambiguity (low or high). The examples are stylized, but they capture the basic idea that narrative tone and narrative uncertainty can vary independently.

2.4 Predicted Macroeconomic Signatures

The framework yields a set of qualitative predictions that guide the empirical analysis.

First, a bullish narrative sentiment shock should be associated with a rise in the real price of oil. If the dominant mechanism is demand-side information, real activity should also increase, producing positive co-movement between oil prices and output.

⁷This feature distinguishes the approach from designs that extract a single latent factor from text, such as the news sentiment index of [Shapiro et al. \(2022\)](#) or the topic-based approach of [Larsen and Thorsrud \(2019\)](#). Here, two distinct dimensions are scored from each article.

Table 1: Illustrative Combinations of Sentiment and Uncertainty in Oil-Market Narratives

	Low uncertainty	High uncertainty
Bearish sentiment	“OPEC’s decision to flood the market sent prices into freefall; analysts say the cartel has abandoned supply management.”	“Markets fear severe disruption from the spreading conflict, but the magnitude and duration of any supply loss remain deeply unclear.”
Bullish sentiment	“Output has stabilized and inventories are rebuilding steadily as demand from China continues to exceed forecasts.”	“Producers expect a strong recovery in demand, though the pace and durability of the rebound remain highly uncertain.”

Note: The examples are stylized composites of oil-market reporting and illustrate the distinction between the directional tone of coverage and the ambiguity surrounding that coverage. The scoring procedure described in Section 3.3 extracts these two dimensions separately from each article.

Oil production may rise with a delay as producers respond to stronger price incentives, and consumer prices may increase gradually through pass-through.

Second, a high narrative uncertainty shock should be associated with a persistent decline in oil production if heightened ambiguity induces delay in irreversible decisions. Real activity should weaken, while the response of the real oil price should be muted or ambiguous because the downward demand effect and the upward supply effect work in opposite directions. The response of consumer prices should therefore be smaller and less precisely signed than under a sentiment shock.

Third, the framework also generates a prediction for episodes in which sentiment and uncertainty rise together. When bullish narrative sentiment coincides with elevated narrative uncertainty, the demand-side information channel pushes oil prices upward while the supply-side delay channel restrains production. Because the two forces operate on different margins, the combined effect is a stagflationary pattern in which oil prices and inflation rise, while production and real activity respond weakly relative to a pure sentiment shock. This joint response is not implied by either shock in isolation. Section 5.3 examines this prediction directly.

These predictions provide a disciplined benchmark for interpreting the empirical results. If sentiment shocks operate through a demand-side information channel, they should generate positive co-movement between oil prices and real activity. If uncertainty shocks operate through a supply-side delay channel, they should suppress production and activity while producing a weaker and less precisely signed price response. The empirical analysis evaluates whether the data exhibit these distinct patterns.

Figure 1 summarizes the two channels. The top path links narrative sentiment to a

demand-side information mechanism. The bottom path links narrative uncertainty to a supply-side delay mechanism consistent with wait-and-see behavior. The empirical analysis asks whether the estimated macroeconomic responses follow these distinct qualitative patterns.

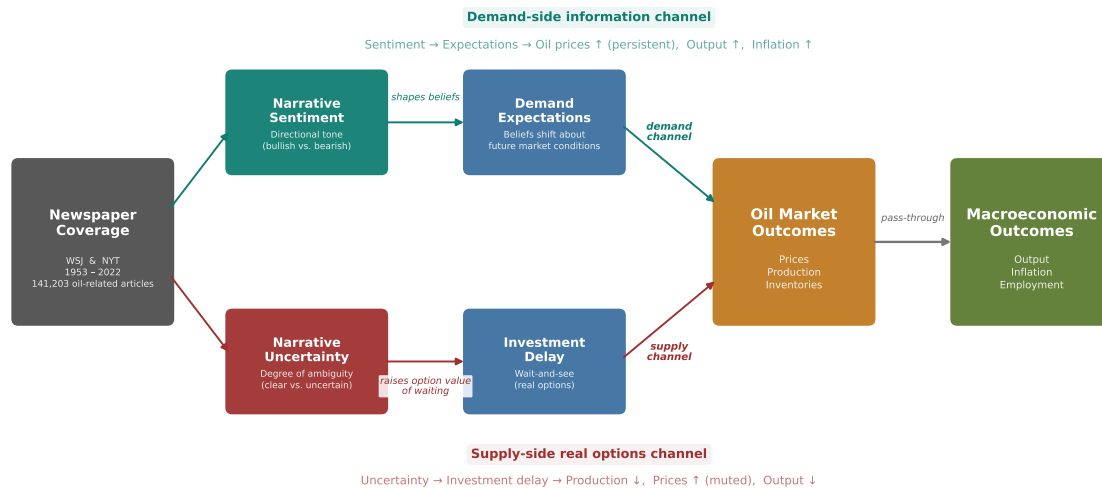


Figure 1: Transmission mechanism: two channels from oil-market narratives to the macroeconomy. The top path traces the demand-side information channel: narrative sentiment shapes expectations about future market conditions, raising oil prices, output, and inflation. The bottom path traces the supply-side uncertainty channel: narrative uncertainty raises ambiguity, delays irreversible decisions, and suppresses oil production with a weaker and less directional price response.

3 Data

This section describes the construction of the narrative indices. The process involves three steps: (i) assembling the newspaper corpus, (ii) scoring articles for sentiment and uncertainty content, and (iii) classifying articles by narrative types.

3.1 Text Corpus

The text data are drawn from ProQuest TDM Studio. I collect all articles from the Business, Finance, Economy, and Markets sections of the *Wall Street Journal* and *New York Times* between January 1953 and December 2022, using section-level metadata rather than keyword searches to avoid selection bias (Gentzkow et al., 2019). The raw corpus comprises 1,973,464 articles; after removing 1,146 with missing text, the final corpus contains 1,972,318 articles spanning 840 months.

3.2 Oil/Energy Article Identification

I identify oil-relevant articles using a domain-specific dictionary organized into five categories: direct oil references (“crude oil,” “petroleum,” “WTI”), OPEC/geopolitics (“OPEC,” “oil embargo”), upstream (“oil drilling,” “shale oil,” “fracking”), downstream (“refinery,” “gasoline price”), and broader energy (“energy crisis,” “energy price”). An article is classified as oil-relevant if it contains at least one core oil term or at least two terms from any category combination, yielding 141,203 articles (7.2%). Figure 2 provides a sample oil-related news article in the corpus.



Figure 2: Sample news article from the *Wall Street Journal*. This illustrates the type of oil narrative content in the corpus

3.3 Text Processing and Scoring

Each oil-related article i is scored along two dimensions, *sentiment* and *uncertainty*, using the formulae:

$$s_i = \frac{N_i^+ - N_i^-}{N_i^{\text{total}}}, \quad u_i = \frac{N_i^{\text{unc}}}{N_i^{\text{total}}}. \quad (1)$$

Here, N_i^+ and N_i^- denote the counts of positive and negative lexicon words in article i , N_i^{unc} denotes the count of uncertainty-related words, and N_i^{total} is the total word count.

The sentiment dictionary is the modifier-augmented oil-specific dictionary of Loughran et al. (2019), who extend the Loughran and McDonald (2011) financial sentiment lexicon to address a construct validity problem specific to oil-market text analysis. Generic financial sentiment dictionaries treat crisis vocabulary (e.g., “embargo,” “conflict,” “shortage,” “disruption”) as unambiguously negative. In oil-market coverage, however, such vocabulary typically signals supply-tightening conditions that are bullish for prices. Applying a generic financial dictionary to oil news therefore systemati-

cally misclassifies geopolitical supply-disruption episodes as bearish sentiment shocks even when the underlying price movement is upward. The Loughran et al. (2019) dictionary contains 59 unambiguously bullish oil keywords, 19 unambiguously bearish oil keywords, and 52 modifier-dependent keywords (e.g., “demand,” “supply,” “output”) whose sign is determined by a ± 4 word window of 291 positive and 536 negative modifiers. This procedure ensures that the directional content of oil-market news is recovered from the perspective of oil prices rather than from the perspective of generic financial sentiment.

The uncertainty dictionary draws on Baker et al. (2016) and Loughran and McDonald (2011) and includes terms such as “uncertain,” “risk,” “unpredictable,” “unclear,” and related variants.⁸

The sentiment index measures the directional tone of oil-market coverage from the perspective of *market conditions*, whether the narrative environment is bullish (prices rising, demand strong, supply tightening) or bearish (prices falling, demand weak, oversupply). A bullish sentiment score does not mean “good news for consumers” or “good news for producers”; it means that the narrative framing of oil-market conditions is directionally positive. This is the relevant dimension for identification because it is the bullish or bearish framing of media coverage, and not the welfare implications for any particular agent, that shapes the demand expectations through which the information channel operates. The same convention is used by Shapiro et al. (2022) for their news sentiment index and by Bybee et al. (2024) for business news topics.

I aggregate article-level scores to monthly frequency by averaging across all oil-relevant articles published in each calendar month, weighting each article equally: $Sent_t = M_t^{-1} \sum_{i=1}^{M_t} s_{i,t}$ and $Unc_t = M_t^{-1} \sum_{i=1}^{M_t} u_{i,t}$, where M_t is the number of oil-relevant articles in month t . The resulting monthly series are standardized to zero mean and unit variance over the structural estimation sample 1953:01–2022:12. Figure 3 displays the two aggregate indices over the full sample.

⁸The distinction between ambiguity (Knightian uncertainty) and risk is formally developed in Ellsberg (1961) and Gilboa and Schmeidler (1989). While the narrative uncertainty index constructed here does not map directly onto the formal ambiguity-aversion framework, the underlying logic is related: elevated narrative uncertainty reflects situations in which agents face a wider and harder-to-characterize range of future oil-market states.




Figure 3: Oil narrative sentiment and uncertainty indices, 1953:01–2022:12. Panel (a): oil narrative sentiment index. Panel (b): oil narrative uncertainty index. Both indices are standardized to zero mean and unit variance over the full sample, with a 0.05 correlation. Shaded vertical bars denote NBER recession dates. Dotted vertical lines mark major oil market events.

Several features of Figure 3 merit discussion. First, the sentiment index (panel a) correctly assigns positive values to supply-tightening crisis episodes and negative values to demand-collapse and oversupply episodes. The Arab Oil Embargo of October 1973 registers as $+0.52\sigma$, the Iranian Revolution of January 1979 as $+0.48\sigma$, the U.S.-led invasion of Iraq in March 2003 as $+0.59\sigma$, the Libyan civil war of March 2011 as $+2.40\sigma$, and the Russia-Ukraine war of February 2022 as $+0.98\sigma$. In each case, news coverage emphasized supply tightness, rising prices, and upward directional pressure on the oil market, and the index registers this directional content as bullish. By contrast, episodes characterized by demand collapse or oversupply register as negative: the OPEC price collapse of January 1986 at -1.27σ , the Asian financial crisis at -1.61σ (August 1998, when the demand impact on oil prices was most severe), the Lehman failure of September 2008 at -1.02σ , the OPEC price war of November 2014 at -0.47σ , and the COVID-19 demand collapse of April 2020 at -3.57σ , the largest negative reading in the sample. The sample-wide pattern of sign assignments validates the directional interpretation of the index.

Second, the uncertainty index (panel b) registers strongly during both supply-tightening crises and demand collapses, consistent with its interpretation as a second-moment object. The largest spike in the sample is the Libyan civil war (March 2011) at $+4.20\sigma$, followed by the Iraqi invasion of Kuwait (August 1990) at $+3.26\sigma$, the U.S.-led invasion of Iraq (March 2003) at $+2.70\sigma$, the Russia-Ukraine war (February 2022)

at $+2.63\sigma$, and the Lehman failure (September 2008) at $+2.46\sigma$. Notably, the uncertainty index also registers the ambiguous-direction events of August 1990 (Iraqi invasion of Kuwait) and September 2001 (9/11) as elevated even when the corresponding sentiment readings are near zero. These were episodes in which oil-market participants faced substantial unresolved questions about the duration, magnitude, and geographic spread of supply disruptions.

Table 2 reports the standardized sentiment and uncertainty indices at twelve major oil-market events spanning the sample, organized by the underlying economic mechanism. The table groups events into three categories: supply-tightening episodes (Panel A), demand-crash and oversupply episodes (Panel B), and ambiguous-direction geopolitical events (Panel C). All five supply-tightening episodes register as positive sentiment shocks. All five demand-crash episodes register as negative sentiment shocks. The two ambiguous-direction events register as near-zero sentiment shocks paired with elevated uncertainty. This pattern provides direct event-level support for the construct validity of the two indices.

Table 2: Narrative Shocks at Major Oil-Market Events

Event	Date	Sentiment (σ)	Uncertainty (σ)
<i>Panel A: Supply-tightening episodes</i>			
Arab Oil Embargo	1973:10	+0.52	-1.46
Iranian Revolution	1979:01	+0.48	-1.00
Iraq Invasion (U.S.-led)	2003:03	+0.59	+2.70
Libyan Civil War	2011:03	+2.40	+4.20
Russia-Ukraine War	2022:02	+0.98	+2.63
<i>Panel B: Demand-crash / oversupply episodes</i>			
OPEC Price Collapse	1986:01	-1.27	-0.82
Asian Financial Crisis	1998:08	-1.61	+1.33
Global Financial Crisis (Lehman)	2008:09	-1.02	+2.46
OPEC Price War	2014:11	-0.47	-0.97
COVID-19 Demand Collapse	2020:04	-3.57	+0.40
<i>Panel C: Ambiguous-direction geopolitical events</i>			
Gulf War (Iraqi invasion of Kuwait)	1990:08	-0.25	+3.26
September 11 attacks	2001:09	+0.16	+1.70

Notes: This table reports standardized monthly oil narrative sentiment and oil narrative uncertainty index values for selected major oil-market events. Values are expressed in standard deviations relative to the full-sample distributions. Higher sentiment denotes more supply-tightening or oil-price-bullish narrative tone, whereas higher uncertainty denotes greater ambiguity in oil-market narratives. Panel classifications reflect the predominant interpretation of each episode and are intended as descriptive rather than definitive.

Third, while the two indices are constructed from the same underlying corpus, they capture distinct dimensions of narrative content. Their Pearson correlation is

+0.053 over the full 1953:01–2022:12 corpus. The divergence between the two indices is most clearly seen in specific contrasts. The OPEC price collapse of January 1986 is a demand-and-oversupply event with substantial directional content but modest residual ambiguity: sentiment registers at -1.27σ while uncertainty registers at -0.82σ . The market understood that Saudi Arabia had abandoned swing-producer status, that supply would expand sharply, and that prices would fall; the directional content of coverage was clear and the residual ambiguity was modest. By contrast, the Libyan civil war of March 2011 is the canonical two-channel event in the sample: sentiment registers at $+2.40\sigma$ (bullish supply-tightening) while uncertainty registers at $+4.20\sigma$ (extreme ambiguity about duration, contagion to other producers, and the broader geopolitical trajectory of the Arab Spring). Both channels are activated simultaneously and at high intensity, reflecting an event that combined sharp directional information with substantial ambiguity. These contrasts, and the broader patterns documented in Table 2, motivate the central identification strategy: sentiment and uncertainty capture distinct dimensions of narrative content that transmit to the macroeconomy through different channels, as evaluated in Section 4.

3.4 Narrative Classification

Each oil-relevant article is assigned to a dominant narrative type based on keyword matching against five domain-specific sub-dictionaries: (i) *geopolitical conflict* (wars, sanctions, military intervention), (ii) *supply disruption* (outages, hurricanes, pipeline failures), (iii) *OPEC coordination* (quotas, production cuts, cartel dynamics), (iv) *speculation/financialization* (futures trading, hedge funds, index investment), and (v) *energy transition* (climate policy, renewables, electric vehicles). Articles not matching any sub-type at a threshold of two keyword occurrences are classified as *general oil*. Title matches are weighted three times body-text matches to reflect the salience of headline framing.⁹

Table 3 reports the classification distribution. The largest classified type is speculation/financial (4.7%), followed by geopolitical (3.9%), OPEC (2.4%), energy transition (1.5%), and supply disruption (1.4%). The remaining 86.1% are general oil articles, as expected, since the sub-dictionaries prioritize precision over recall: the threshold requirement that an article contain at least two type-specific keywords (with title matches weighted three times body-text matches) minimizes false-positive classifications at the cost of a larger residual general category. Monthly type-specific sentiment and uncertainty indices are computed as in Section 3.3 and standardized over the structural estimation sample.¹⁰

⁹The complete keyword lists for each of the five sub-dictionaries are reported in Appendix A.

¹⁰Appendix B validates both the oil-relevance filter and the narrative classification using a large lan-

Table 3: Narrative Type Distribution of Oil-Relevant Articles

Narrative Type	Articles	Share (%)
Speculation / Financial	6,580	4.7
Geopolitical Conflict	5,445	3.9
OPEC Coordination	3,413	2.4
Energy Transition	2,129	1.5
Supply Disruption	2,038	1.4
General Oil	121,598	86.1
Total	141,203	100.0

Notes: Classification based on keyword matching against five sub-dictionaries. Threshold: ≥ 2 matches. Titles weighted $3\times$. An article is assigned to a single narrative type based on its dominant keyword match; articles that do not pass the threshold for any sub-dictionary are classified as *general oil*.

Figure 4 displays the evolution of type-specific narrative sentiment and uncertainty over the full 1953–2022 sample. The patterns track oil-market history closely: speculation narratives surge after 2004 with the financialization of commodity markets, OPEC narratives spike during every major coordination episode from the 1973 embargo through the 2014–2016 price war, geopolitical narratives peak during the Gulf War, the Iraq invasion, and the Russia-Ukraine war, and energy transition narratives emerge gradually after the 2015 Paris Agreement and accelerate sharply after 2020.

Narrative plausibility. As documented in the discussion of Figures 3 and 4, both the aggregate and type-specific indices align closely with qualitative accounts of major oil-market episodes. At the aggregate level (Figure 3), the sentiment index correctly separates supply-tightening episodes from demand-crash episodes. Supply-tightening crises register as positive sentiment shocks: the Arab Oil Embargo ($+0.52\sigma$), the Iranian Revolution ($+0.48\sigma$), the U.S.-led invasion of Iraq ($+0.59\sigma$), the Libyan civil war ($+2.40\sigma$), and the Russia-Ukraine war ($+0.98\sigma$). Demand-crash episodes register as negative sentiment shocks: the OPEC price collapse of 1986 (-1.27σ), the Asian financial crisis (-1.61σ), the Lehman failure (-1.02σ), the OPEC price war of 2014 (-0.47σ), and the COVID-19 demand collapse (-3.57σ , the sample minimum). The uncertainty index spikes during periods of genuine ambiguity about oil-market outcomes, with its largest readings during the Libyan civil war ($+4.20\sigma$), the Iraqi invasion of Kuwait ($+3.26\sigma$), the U.S.-led invasion of Iraq ($+2.70\sigma$), the Russia-Ukraine war ($+2.63\sigma$), and the Lehman failure ($+2.46\sigma$). This sample-wide pattern of sign assignments validates the directional interpretation of the sentiment index and the second-moment interpretation of the uncertainty index.

guage model.

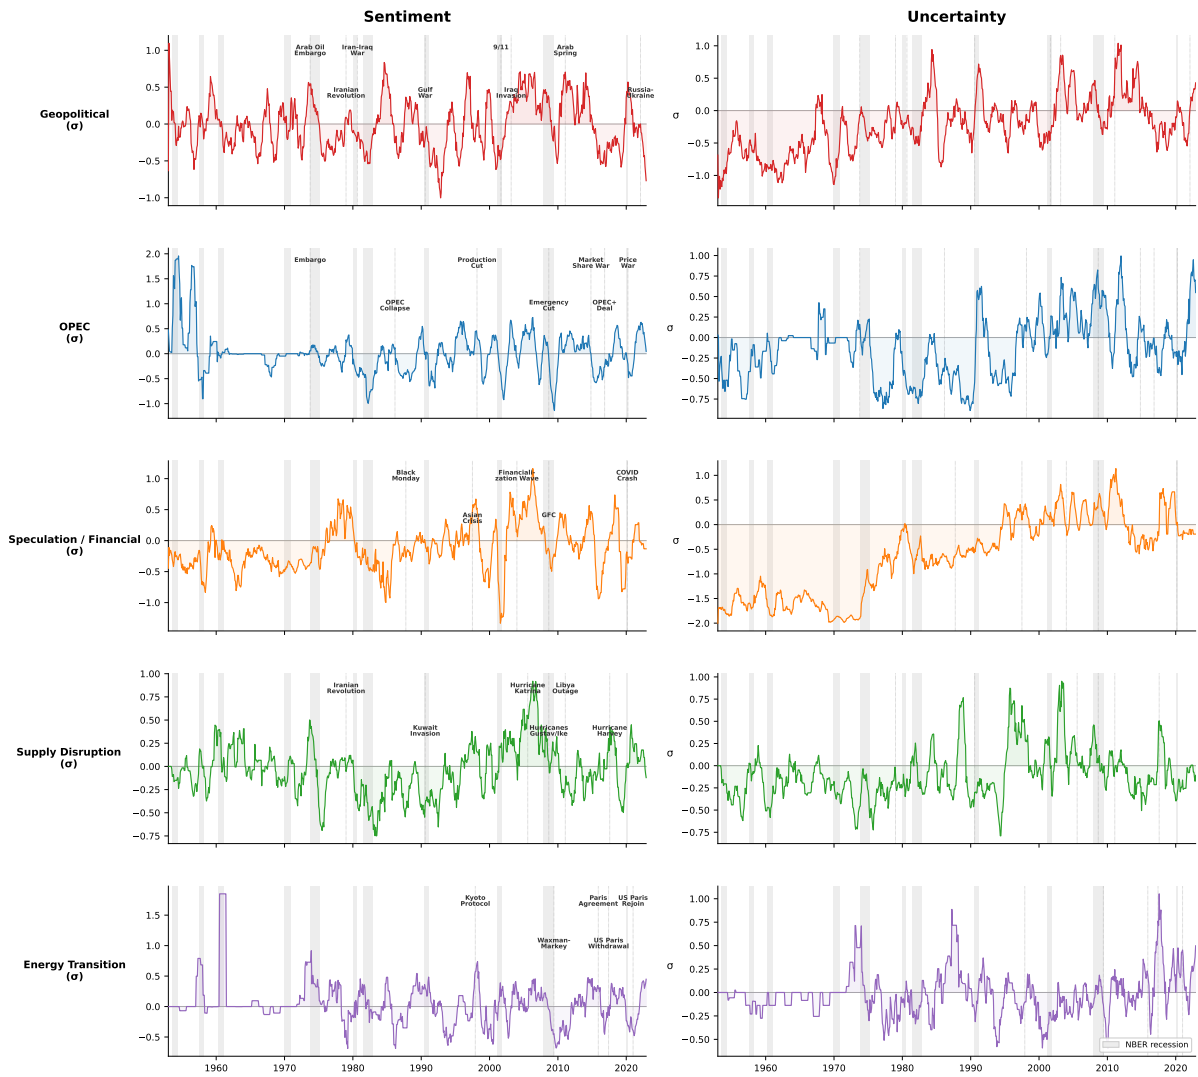


Figure 4: Type-specific narrative indices, 1953–2022. Each row corresponds to one narrative type: geopolitical conflict, OPEC coordination, speculation/financial, supply disruption, and energy transition. Left column: sentiment (directional tone). Right column: uncertainty (degree of ambiguity). All series are 12-month moving averages of the standardized type-specific indices. Dotted vertical lines mark type-relevant events. Shaded vertical bars denote NBER recession dates.

At the type-specific level (Figure 4), each index responds to type-relevant events: geopolitical sentiment moves sharply during armed conflicts, OPEC sentiment swings around coordination episodes, speculation sentiment tracks the financialization wave of the mid-2000s, supply disruption uncertainty rises during hurricanes, and energy transition coverage emerges after the Paris Agreement and accelerates after 2020. Crucially, the type-specific indices respond to *different* events at *different* times: during the Gulf War (1990–91), geopolitical uncertainty spikes sharply while speculation uncertainty remains near zero; during the financialization wave (2004–08), speculation uncertainty rises while supply uncertainty remains subdued. This cross-type heterogeneity is precisely the variation that the identification strategy exploits.

Distinctness from existing measures. Table 4 reports pairwise correlations between the oil narrative indices and six widely used comparison measures: the University of Michigan Consumer Sentiment Index (UMCSent), the [Caldara and Iacoviello \(2022\)](#) Geopolitical Risk index (GPR), the [Abiad and Qureshi \(2023\)](#) Oil Price Uncertainty index (OPU), the [Baker et al. \(2016\)](#) News-Based Economic Policy Uncertainty index (EPU), the CBOE Volatility Index (VIX), and the CBOE Oil Volatility Index (OVX). Correlations are reported over the structural estimation sample 1987:02–2022:12.¹¹

Panel A examines the aggregate indices. Oil narrative uncertainty exhibits its strongest correlation with the oil price uncertainty (OPU) index of [Abiad and Qureshi \(2023\)](#) at $\rho = +0.47$, followed by VIX ($\rho = +0.40$), GPR ($\rho = +0.23$), and EPU ($\rho = +0.27$). These moderate positive correlations indicate that narrative uncertainty shares common variation with established measures of macro-financial and oil-specific uncertainty without being redundant with any single one of them. The strongest link is with OPU, which is the closest comparable measure in scope, yet the $\rho = +0.47$ correlation implies that fewer than 25% of the variation in the two series overlaps. Oil narrative sentiment correlates negatively with oil-market volatility (OVX, $\rho = -0.55$) and with general financial volatility (VIX, $\rho = -0.39$), reflecting the well-established inverse relationship between bullish directional tone and realized volatility, and positively with consumer sentiment ($\rho = +0.15$). The sentiment index is essentially uncorrelated with the [Caldara and Iacoviello \(2022\)](#) Geopolitical Risk index ($\rho = +0.02$), which is economically meaningful: it confirms that the sentiment index does not simply track the frequency of geopolitical crises. Supply-tightening geopolitical events register as bullish sentiment shocks while non-oil geopolitical events do not move the index, so aggregate correlation with GPR averages to near zero. The cross-correlation between the aggregate sentiment and uncertainty indices is $\rho = +0.14$, compared to typical pairs of news-based macroeconomic indices that correlate at 0.4 to 0.6 (for example, EPU with VIX at approximately 0.4, or the [Jurado et al. \(2015\)](#) macro and financial uncertainty indices at approximately 0.6). The near-orthogonality of sentiment and uncertainty is therefore a substantive feature of the decomposition, not an artifact of construction.

Panels B and C report correlations for the type-specific indices that serve as candidate instruments in the Proxy-SVAR. The type-specific correlations are uniformly weaker than the aggregate correlations, as expected when narrower sub-components are compared against broad market-wide measures. The strongest type-specific sentiment correlation is between OPEC sentiment and OVX ($\rho = -0.32$), reflecting that OPEC-related narrative tone moves with oil-market price cycles that OVX reflects.

¹¹The VIX and OVX series begin in 1990 and 2007, respectively. OVX correlations are therefore computed over a 187-month window weighted toward high-volatility episodes (2014 oil price war, 2020 COVID demand collapse, 2022 Russia-Ukraine war).

The strongest type-specific uncertainty correlations are between OPEC uncertainty and OPU ($\rho = +0.26$) and between transition uncertainty and EPU ($\rho = +0.23$), both consistent with the interpretation that OPEC-related narratives overlap partially with oil-policy uncertainty and that energy-transition narratives overlap partially with broader policy uncertainty. Geopolitical uncertainty correlates moderately with GPR ($\rho = +0.21$), which confirms domain overlap without redundancy: the two measures both capture geopolitical-crisis intensity but the oil-specific framing differs. Supply-disruption uncertainty correlates just $+0.03$ with GPR and $+0.11$ with OPU, reinforcing the interpretation that physical supply events are economically distinct from geopolitical or policy-driven shocks. All other type-specific correlations fall below $|0.20|$, confirming that each type-specific index captures variation that no existing market-wide measure isolates at this level of granularity.

Table 4: Correlation of Narrative Indices with Comparison Measures (Standardized Indices)

	UMCSent	GPR	OPU	EPU	VIX	OVX	Sent/Unc
<i>Panel A: Aggregate indices</i>							
Oil narrative sentiment	+0.15	+0.02	-0.04	-0.25	-0.39	-0.55	+0.14
Oil narrative uncertainty	-0.36	+0.23	+0.47	+0.27	+0.40	+0.29	—
<i>N</i>	431	431	431	431	396	187	431
<i>Panel B: Type-specific sentiment indices</i>							
Geopolitical	+0.01	+0.05	+0.05	-0.07	-0.05	-0.13	—
OPEC	+0.05	-0.08	-0.01	-0.11	-0.16	-0.32	—
Speculation	+0.04	+0.04	-0.03	-0.20	-0.16	-0.28	—
Supply	+0.01	+0.02	+0.14	-0.15	-0.07	-0.13	—
Transition	-0.07	+0.10	+0.24	+0.08	-0.00	-0.12	—
<i>Panel C: Type-specific uncertainty indices</i>							
Geopolitical	-0.16	+0.21	+0.18	+0.05	+0.07	+0.04	—
OPEC	-0.15	+0.17	+0.26	-0.03	+0.18	+0.13	—
Speculation	+0.02	+0.06	+0.05	-0.16	+0.11	+0.01	—
Supply	+0.03	+0.03	+0.11	-0.14	+0.03	-0.10	—
Transition	-0.17	+0.08	+0.21	+0.23	+0.11	+0.12	—

Notes: Pairwise Pearson correlations between the oil narrative indices and six widely used comparison measures, computed over the structural estimation sample 1987:02–2022:12. UMCSent: University of Michigan Consumer Sentiment Index. GPR: [Caldara and Iacoviello \(2022\)](#) Geopolitical Risk index. OPU: [Abiad and Qureshi \(2023\)](#) Oil Price Uncertainty index. EPU: [Baker et al. \(2016\)](#) News-Based Economic Policy Uncertainty index. VIX: CBOE Volatility Index (daily averages aggregated to monthly). OVX: CBOE Oil Volatility Index (daily averages aggregated to monthly). Sent/Unc: cross-correlation between aggregate sentiment and uncertainty. Panel A reports correlations for the aggregate (standardized) indices. Panels B and C report correlations for the ten type-specific indices that serve as candidate instruments in the Proxy-SVAR. *N* reports the number of monthly observations; VIX and OVX begin in 1990 and 2007, respectively.

These results establish that the oil narrative indices capture economically mean-

ingful variation that is distinct from existing measures of sentiment, uncertainty, and risk. The aggregate uncertainty index correlates moderately with each standard uncertainty measure, with OPU being the closest but far from redundant. The aggregate sentiment index correlates inversely with financial volatility, consistent with the established relationship between bullish directional tone and realized market calm, and is nearly orthogonal to the Geopolitical Risk index, which establishes that sentiment is not simply a crisis-frequency measure. The type-specific indices are uniformly weakly correlated with all six comparison measures, confirming that the type decomposition isolates source-specific variation (geopolitical events, OPEC decisions, speculative activity, physical supply disruptions, and energy transition policy) that no existing index captures at this level of granularity. These properties support the use of the type-specific indices as external instruments in the Proxy-SVAR estimated in Section 4.

4 Econometric Framework

4.1 Proxy-SVAR Identification

I model the joint dynamics of narrative and macroeconomic variables using a reduced-form VAR of order p :

$$\mathbf{Y}_t = \boldsymbol{\mu} + \sum_{\ell=1}^p \mathbf{B}_\ell \mathbf{Y}_{t-\ell} + \mathbf{u}_t, \quad \mathbf{u}_t \sim (0, \boldsymbol{\Sigma}_u), \quad (2)$$

where \mathbf{Y}_t is an $n \times 1$ vector of endogenous variables with $n = 6$. The structural relationship is $\mathbf{u}_t = \mathbf{A}_0^{-1} \boldsymbol{\varepsilon}_t$, where $\boldsymbol{\varepsilon}_t$ is the vector of mutually uncorrelated structural shocks and $\boldsymbol{\Sigma}_u = \mathbf{A}_0^{-1} (\mathbf{A}_0^{-1})'$.

Identification of a single structural shock follows the Proxy-SVAR framework of [Mertens and Ravn \(2013\)](#) and [Stock and Watson \(2018\)](#). Let $\varepsilon_{1,t}$ denote the structural shock of interest and let \mathbf{Z}_t denote the $k \times 1$ vector of external instruments ($k = 4$). The instruments must satisfy:

$$E[\mathbf{Z}_t \varepsilon_{1,t}] = \boldsymbol{\phi} \neq 0 \quad (\text{relevance}), \quad (3)$$

$$E[\mathbf{Z}_t \varepsilon_{j,t}] = 0 \quad \text{for } j \neq 1 \quad (\text{exogeneity}). \quad (4)$$

The first-stage regression tests relevance by projecting the reduced-form residual of the instrumented variable, $u_{1,t}$, on the instruments:

$$u_{1,t} = \mathbf{Z}_t' \boldsymbol{\gamma} + v_t. \quad (5)$$

The F -statistic from (5) serves as the primary diagnostic for instrument strength. The

fitted values $\hat{u}_{1,t}$ recover the structural impulse vector \mathbf{b}_1 (the column of \mathbf{A}_0^{-1} corresponding to $\varepsilon_{1,t}$) up to scale and sign (Mertens and Ravn, 2013). Because $k = 4 > 1$, the system is over-identified, which strengthens the first stage relative to any single instrument.

I estimate two separate systems. In the **sentiment system**, the shock of interest is the narrative sentiment shock ($\varepsilon_{1,t} = \varepsilon_t^{\text{sent}}$), and the four sentiment-type instruments serve as \mathbf{Z}_t . In the **uncertainty system**, the shock of interest is the narrative uncertainty shock ($\varepsilon_{1,t} = \varepsilon_t^{\text{unc}}$), and the four uncertainty-type instruments serve as \mathbf{Z}_t . Importantly, both aggregate indices (*sentiment* and *uncertainty*) enter both VARs as endogenous variables. This allows me to test whether the sentiment shock moves the uncertainty index and vice versa, providing direct evidence on the orthogonality of the two identified shocks.

Confidence bands are constructed using the moving block bootstrap (MBB) of Stock and Watson (2018) with block length $\lceil 5.03T^{1/4} \rceil$, where T is the effective sample size. I report 90% and 68% confidence intervals based on 10,000 bootstrap replications. Forecast error variance decompositions and historical decompositions are computed following Kilian and Lütkepohl (2017). As a specification check, I also report impulse responses from local projections with external instruments (LP-IV) following Stock and Watson (2018), which do not impose VAR dynamics.

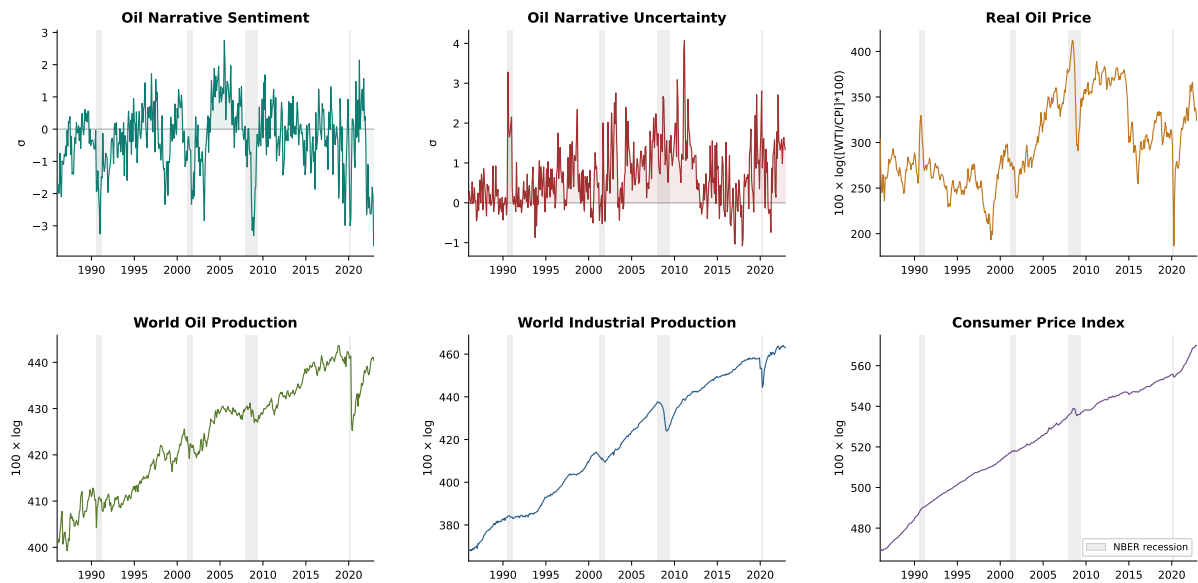


Figure 5: VAR variables, 1986:01–2022:12. Top row: oil narrative sentiment (standardized), oil narrative uncertainty (standardized), real oil price ($100 \times \log([\text{WTI}/\text{CPI}] \times 100)$). Bottom row: world crude oil production ($100 \times \log$), world industrial production ($100 \times \log$), consumer price index ($100 \times \log$). Shaded vertical bars denote NBER recession dates.

4.2 Data and Variables

The VAR includes six variables:

$$Y_t = \left(\text{SENT}_t \quad \text{UNC}_t \quad \text{ROP}_t \quad \text{WOP}_t \quad \text{WIP}_t \quad \text{CPI}_t \right)', \quad (6)$$

where SENT_t and UNC_t are the standardized aggregate narrative sentiment and uncertainty indices (constructed in Section 3.3), $\text{ROP}_t = 100 \times \log\left(\frac{\text{WTI}_t}{\text{CPIAUCSL}_t} \times 100\right)$ is the real oil price, $\text{WOP}_t = 100 \times \log(\text{wOILPRO}_t)$ is world oil production, $\text{WIP}_t = 100 \times \log(\text{wINDPRO}_t)$ is world industrial production, and $\text{CPI}_t = 100 \times \log(\text{CPIAUCSL}_t)$ is the consumer price index. For the uncertainty system, the ordering of the first two variables is reversed.

The real oil price is constructed using the WTI spot price from the U.S. Energy Information Administration (EIA), available from January 1986, deflated by the CPI-AUCSL. World industrial production (wINDPRO) is from [Baumeister and Hamilton \(2019\)](#), and the consumer price index (CPIAUCSL) is from the FRED-MD database ([McCracken and Ng, 2016](#)). World crude oil production (wOILPRO) is from the EIA. All variables enter in $100 \times \log$ levels except the text indices, which are standardized. The lag order is $p = 6$.

The estimation sample runs from 1987:02 to 2022:12 ($T = 425$ monthly observations after purging lags), determined by the intersection of EIA oil price availability (1986:01), text index coverage, and the loss of initial observations from instrument purging. Figure 5 displays the six VAR variables over the estimation sample.

4.3 Instrument Construction

The identification strategy requires external instruments that are correlated with the aggregate narrative indices but capture variation from distinct, identifiable sources. I propose a novel approach: using the *within-corpus narrative type decomposition* as the source of instrumental variation.

The five type-specific sentiment indices (geopolitical, OPEC, speculation, supply disruption, and energy transition) each measure the directional tone of coverage from a different narrative source. By construction, the aggregate sentiment index reflects both the type-specific components and the much larger residual general-oil component. The type-specific indices are therefore mechanically relevant for the aggregate. The key insight is that each type captures variation from a *different driver* of oil-market news: geopolitical events operate through conflict and sanctions, OPEC decisions operate through cartel coordination, speculation operates through financial market positioning, supply disruptions operate through physical infrastructure, and energy transition operates through policy and technology. These are economically distinct sources

of narrative variation, which makes them suitable for joint identification: each captures a different sub-channel of the same underlying narrative shock rather than the same shock measured with different noise.

A useful clarification is that the type-specific instruments are not close replicas of the aggregate narrative indices. The aggregate indices are constructed from the full set of oil-relevant articles, whereas the headline instrument set is drawn only from a small classified subset of source-specific narratives. As shown in Table 3, 86.1% of oil-relevant articles fall into the residual general-oil category, while the four headline instrument categories; geopolitical conflict, speculation/ financialization, supply disruption, and energy transition, together account for only 11.5% of oil-relevant coverage. The instruments therefore do not mechanically reproduce the aggregate index. Rather, they isolate source-specific variation within a much broader oil-news environment in which most coverage is not assigned to any headline type. This distinction is important for identification: the role of the type-specific instruments is not to proxy for the aggregate index with minor measurement differences, but to provide distinct source-level narrative surprises that shift the broader aggregate narrative disturbance.

This “narrative heterogeneity as identification” approach differs from existing strategies in the oil-market literature. [Känzig \(2021\)](#) identifies oil supply news shocks using high-frequency price movements around OPEC announcements (a single-event, single-instrument strategy). [Gavriilidis et al. \(2026\)](#) construct a narrative instrument from the media impact of 146 classified climate policy events. The approach here replaces narrow event windows with the full richness of continuous narrative variation, decomposed by source, yielding multiple instruments per shock that jointly produce strong first stages. The approach is also portable: any text-based index that can be decomposed by source or topic can use the same strategy.

Exclusion of OPEC from the headline instrument set. Of the five type-specific narratives constructed in Section 3.3, four (geopolitical, speculation, supply disruption, and energy transition) are rooted in events that evolve on timescales substantially independent of monthly oil-price movements. The fifth, OPEC, differs in an important respect: OPEC’s institutional purpose is to respond to oil-market conditions in real time. Reporters write more about OPEC when prices move, OPEC itself communicates more when prices move, and market commentary about OPEC responds directly to prices. The result is a tight contemporaneous feedback loop between OPEC narratives and oil prices that has no parallel in the four other categories. The purging step described in Section 4.4 removes predictable feedback at lagged horizons but cannot remove contemporaneous feedback within the same month. Appendix C.3 (Table 6) shows that including OPEC alongside the other four instruments produces a borderline rejection of the over-identification restrictions in the sentiment system,

driven entirely by the real-oil-price equation, while excluding OPEC delivers clean over-identification across both systems. The headline specification therefore uses the four non-OPEC type-specific instruments; Appendix E.5 reports the OPEC-inclusive specification as a robustness check, and the structural conclusions are qualitatively unchanged.

The four purged sentiment instruments form the instrument vector for the sentiment system in the headline specification:

$$\mathbf{Z}_t^{\text{sent}} = \left(\hat{\eta}_t^{\text{geo,sent}} \quad \hat{\eta}_t^{\text{spec,sent}} \quad \hat{\eta}_t^{\text{supply,sent}} \quad \hat{\eta}_t^{\text{trans,sent}} \right)'. \quad (7)$$

Analogously, four purged uncertainty instruments form $\mathbf{Z}_t^{\text{unc}}$.

4.4 Instrument Purging

Following Romer and Romer (2004) and Miranda-Agrippino and Ricco (2021), each type-specific index is purged of predictable variation before use as an instrument. The purging regression for instrument j in dimension $d \in \{\text{sent}, \text{unc}\}$ is:

$$s_t^{j,d} = c + \sum_{\ell=1}^{12} \alpha_{\ell} s_{t-\ell}^{j,d} + \sum_{\ell=1}^{12} \beta'_{\ell} \mathbf{x}_{t-\ell} + \eta_t^{j,d}, \quad (8)$$

where $\mathbf{x}_t = (\Delta \log \text{ROP}_t, \Delta \log \text{WOP}_t, \Delta \log \text{WIP}_t, \Delta \log \text{CPI}_t)'$ is the vector of macroeconomic state variables. The residual $\hat{\eta}_t^{j,d}$ replaces the raw index as the purged instrument.

This step is essential. It ensures that the instruments capture genuinely *new* narrative information (the surprise component of type-specific media coverage) rather than responding to lagged macroeconomic conditions. The purging is conceptually related to the high-frequency identification approach in Känzig (2021): just as Känzig's OPEC announcement returns isolate the surprise component of OPEC decisions by removing variation predictable from prior market conditions, the purging regression removes the predictable component of narrative variation. The two strategies achieve exogeneity at different temporal resolutions: Känzig's design within a narrow announcement window, and the present approach at all lags up to twelve months.

4.5 Instrument Validation

I assess the validity of the purged instruments along two dimensions in this section: first-stage strength and over-identification. Two additional diagnostic exercises: Granger non-causality tests verifying the success of the purging procedure and Ljung-Box tests for serial correlation in the purged instruments, are reported in Appendix C.

First-stage strength. Table 5 reports the first-stage regression results for both systems. For the sentiment shock, the joint F -statistic from the four purged type-specific instruments is 16.7 (HC1-robust: 16.8), with $R^2 = 13.7\%$. Among individual instruments, transition sentiment is the strongest (Std. $F = 19.3$, Eff. $F = 22.5$), followed by geopolitical ($F = 18.1$) and speculation ($F = 15.7$). Supply sentiment is individually weak ($F = 5.7$) but contributes to the joint first stage.

For the uncertainty shock, the joint F -statistic is 27.0 (HC1-robust: 13.9), with $R^2 = 20.5\%$. The composition of instrument strength differs markedly: geopolitical uncertainty dominates (Std. $F = 47.5$), followed by speculation ($F = 41.0$) and supply ($F = 12.4$). Transition uncertainty is individually weak ($F = 3.5$). The gap between the standard and HC1-robust joint F -statistics in the uncertainty system (27.0 vs. 13.9) indicates heteroskedasticity in the first-stage residuals, which motivates the use of the heteroskedasticity-robust effective F -statistic below.

I supplement the conventional diagnostics with the effective F -statistic of [Montiel Olea and Pflueger \(2013\)](#), which accounts for heteroskedasticity and serial correlation in the first-stage regression via a Newey-West HAC covariance estimator. The effective F is 15.8 for the sentiment system and 12.1 for the uncertainty system. The sentiment system comfortably exceeds conventional weak-instrument benchmarks under the [Montiel Olea and Pflueger \(2013\)](#) framework. The uncertainty system is somewhat more borderline: its effective F falls slightly below the 10% worst-case Nagar-bias critical value of 12.56 for the four-instrument case, but remains above the corresponding 20% worst-case bias critical value of 8.59. Additionally, it beats the traditional $F > 10$ criteria. This assessment is also consistent with the HC1-robust joint first-stage statistic of 13.9 and the broader set of instrument-validation diagnostics reported below.

The shift in instrument ranking across systems is economically informative. Sentiment is driven by narratives about *what is happening and what market participants expect*: geopolitical events that signal supply tightening, speculative positioning that reflects directional bets, and energy transition policy that shifts long-run demand expectations. Uncertainty is driven by narratives about *sources of unpredictability*: geopolitical instability whose duration and scope are unknown, speculative froth that widens the range of possible outcomes, and physical supply disruptions whose resolution timeline is uncertain. This asymmetry is consistent with the two-channel framework developed in Section 2: sentiment operates through first-moment belief updates while uncertainty operates through second-moment option-value effects.

Over-identification. Table 6 reports Sargan J -statistics for both systems. Panel A shows that the four sentiment instruments pass the over-identification test at all conventional levels: the system J -statistic is 14.46 ($p = 0.491$, distributed $\chi^2(15)$ under H_0), and no individual equation rejects at the 5% level. The real-oil-price equa-

Table 5: First-Stage Diagnostics with Effective F -Statistic

Panel A: Joint Instruments (4 instruments, 1 endogenous variable)

	Standard F	HC1-Robust F	Effective F	R^2	Adj. R^2	T
Sentiment	16.7	16.8	15.8	13.7%	12.9%	425
Uncertainty	27.0	13.9	12.1	20.5%	19.7%	425

Montiel Olea–Pflueger critical values ($k = 4$)

$\tau = 5\%$ max bias	20.15
$\tau = 10\%$ max bias	12.56
$\tau = 20\%$ max bias	8.59

Panel B: Individual Instruments

Instrument	Sentiment			Uncertainty		
	Std. F	Eff. F	R^2	Std. F	Eff. F	R^2
Geopolitical	18.1	17.2	4.1%	47.5	17.6	10.1%
Speculation	15.7	15.2	3.6%	41.0	17.5	8.8%
Supply	5.7	6.3	1.3%	12.4	10.1	2.8%
Transition	19.3	22.5	4.4%	3.5	3.1	0.8%

Montiel Olea–Pflueger critical values ($k = 1$)

$\tau = 10\%$ max bias	23.11
$\tau = 20\%$ max bias	15.06

Notes: Panel A reports the standard, heteroskedasticity-robust (HC1), and effective F -statistics from the joint first-stage regression of the reduced-form VAR residual on all 4 purged type-specific instruments (geopolitical, speculation, supply, transition; OPEC excluded from the headline specification, see Section 4.3). The effective F -statistic of [Montiel Olea and Pflueger \(2013\)](#) uses a Newey–West HAC covariance estimator (5 lags) to account for both heteroskedasticity and serial correlation. Under the null of weak instruments, the effective F is compared to critical values $c_\alpha(\tau)$ from [Montiel Olea and Pflueger \(2013\)](#); τ denotes the worst-case Nagar bias as a fraction of the OLS bias. Panel B reports individual (bivariate) first-stage results for each instrument separately, with $k = 1$ critical values. Sample: 1987:02–2022:12.

tion, which is the most economically sensitive test of instrument validity because it is the variable most likely to exhibit contemporaneous feedback with narrative content, yields $J = 0.89$ ($p = 0.829$). Panel B shows an equally clean result for the uncertainty system: the system J -statistic is 10.22 ($p = 0.805$), with all equation-level p -values exceeding 0.37.

Panel C reports instrument-exclusion C -tests for the sentiment system, testing the marginal validity of each instrument conditional on the remaining three being valid. All four C -statistics are small and non-significant (maximum $C = 0.85$, $p = 0.356$), confirming that no individual instrument introduces identification heterogeneity. The four instruments are mutually consistent proxies for a single underlying narrative sentiment shock.

These diagnostics support the exclusion of OPEC from the headline instrument set. When OPEC is included as a fifth instrument (Appendix C.3), the system J -statistic for the sentiment shock rises to 30.20 ($p = 0.067$), with the rejection driven entirely by the real-oil-price equation ($J = 14.31$, $p = 0.006$). The C -test isolates OPEC as the source of this heterogeneity ($C = 13.43$, $p < 0.001$), while all other instruments have C -statistics below 3.25 (all $p > 0.07$). This pattern is consistent with the contemporaneous-feedback argument developed in Section 4.3: OPEC news responds to oil prices within the month, generating a projection onto the real-oil-price residual that differs from the other four instruments. Excluding OPEC resolves this tension cleanly while preserving strong first-stage relevance in both systems.

5 Empirical Results

5.1 Narrative Sentiment Shock

Figure 6 displays impulse responses to a one-standard-deviation *bullish* narrative sentiment shock, identified using the four purged type-specific sentiment instruments in the Proxy-SVAR framework. The first-stage statistics indicate meaningful instrument relevance: $F = 16.7$ (robust: 16.8), effective $F = 15.8$, with $R^2 = 13.7\%$ and adjusted $R^2 = 12.9\%$.

Real oil price. The real oil price rises sharply on impact, reaching roughly 3% in the first few months, with a peak of about 3.2% around months 1–3. The response then declines gradually but remains positive throughout the horizon, settling at a modest positive level at longer horizons. This pattern is consistent with the idea that bullish oil-market narratives shift beliefs about future fundamentals and raise current willingness to pay for oil.

Table 6: Over-Identification Tests*Panel A: Narrative Sentiment Shock*

Equation	<i>J</i> -statistic	<i>p</i> -value
Other index	1.011	0.799
Real oil price	0.886	0.829
World oil prod	3.941	0.268
World IP	6.099	0.107
CPI	2.526	0.471
System	14.463	0.491

Panel B: Narrative Uncertainty Shock

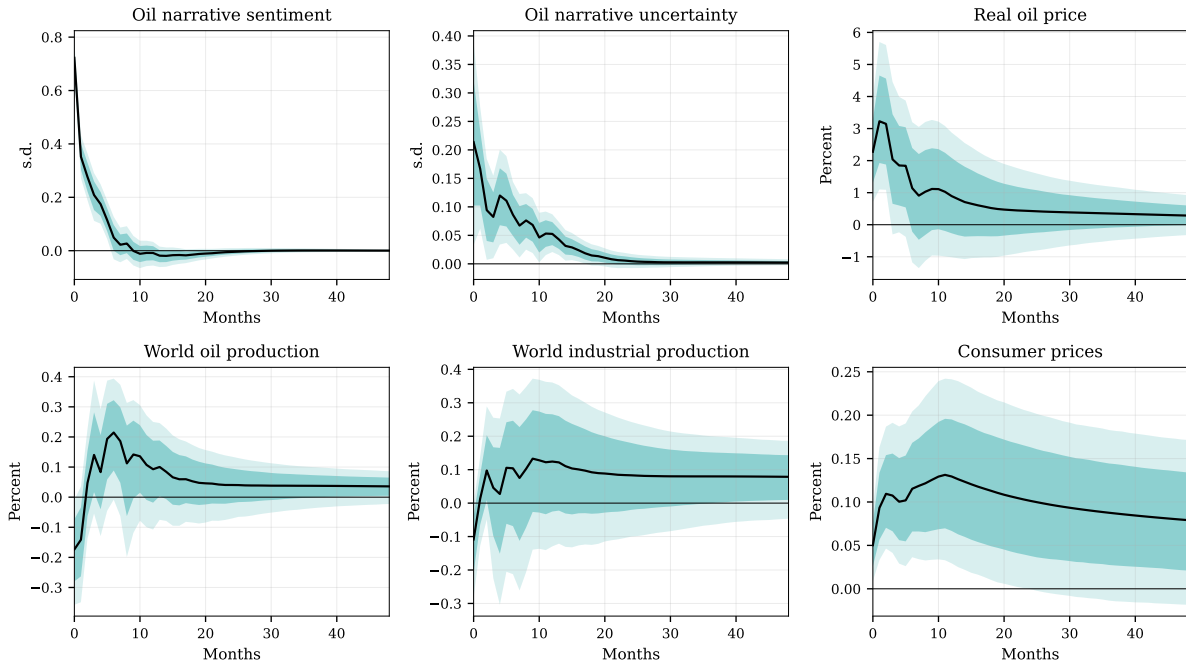
Equation	<i>J</i> -statistic	<i>p</i> -value
Other index	3.113	0.375
Real oil price	2.923	0.404
World oil prod	0.979	0.806
World IP	2.908	0.406
CPI	0.300	0.960
System	10.223	0.805

Panel C: Instrument Exclusion (C-test), Sentiment Shock

Dropped instrument	<i>F</i> (remaining 3)	<i>C</i> -statistic	<i>p</i> -value
Geopolitical	14.4	0.128	0.721
Speculation	16.0	0.056	0.813
Supply	20.7	0.852	0.356
Transition	12.9	0.001	0.980

Notes: Panels A and B report Sargan *J*-statistics from 2SLS regressions of each reduced-form VAR residual $u_{j,t}$ on the 4 purged type-specific instruments (geopolitical, speculation, supply, transition; OPEC excluded from the headline specification, see Section X), using the fitted value $\hat{u}_{1,t}$ from the first-stage regression. Under H_0 (all instruments valid), $J = T \times R^2 \sim \chi^2(3)$, where 3 is the number of over-identifying restrictions ($k - 1 = 4 - 1$). The system statistic sums the equation-level *J*-statistics; under H_0 it is distributed $\chi^2(15)$ (four equations \times 3 restrictions per equation, plus the cross-index equation). Panel C reports the *C*-statistic (difference-in-*J*) for the sentiment system, testing the marginal validity of each instrument conditional on the remaining 3 being valid; $C \sim \chi^2(1)$ under H_0 . * $p < 0.05$; ** $p < 0.01$.

Impulse responses to a one s.d. narrative sentiment shock



First-stage: $F = 16.7$ (robust 16.8), $Eff. F = 15.8$, $R^2 = 13.7\%$, $Adj. R^2 = 12.9\%$

Figure 6: Impulse responses to a one standard deviation narrative sentiment shock. Dark shaded region: 68% MBB confidence bands. Light shaded region: 90% MBB confidence bands. Solid line: point estimate. The shock is identified using four purged type-specific sentiment instruments in the Proxy-SVAR framework. Sample: 1987:02–2022:12, VAR(6), 10,000 bootstrap replications.

World industrial production. World industrial production responds positively, though modestly. After some short-run oscillation, the response rises to roughly 0.12–0.14% around months 8–12 and then gradually settles to a still-positive level of about 0.08%. The key feature is the positive co-movement of oil prices and global output. That sign pattern is difficult to reconcile with an adverse supply disturbance, which would tend to raise prices while reducing activity, and instead points toward a demand-side information channel.

World oil production. World oil production shows a brief negative impact response, but this is quickly reversed. By about month 2 the response turns positive, peaks near 0.2% around months 5–7, and then decays gradually toward a small but positive long-run effect. This hump-shaped response suggests that production adjusts with a delay to narrative-driven price pressure. The delayed expansion in output is more consistent with supply responding endogenously to stronger expected demand conditions than with the sentiment shock reflecting an exogenous physical supply contraction.

Consumer prices. Consumer prices increase on impact and continue to rise gradually for about a year, reaching a peak near 0.13% around months 9–12. The response then declines slowly but remains positive through the end of the horizon, ending around 0.08%. This pattern indicates pass-through from oil-market narratives into consumer prices operating primarily through the oil price channel.

Cross-effect on uncertainty. A bullish sentiment shock also raises oil narrative uncertainty temporarily. Uncertainty increases on impact by roughly 0.25 standard deviations, remains positive through roughly the first year, and then gradually returns toward zero by around 20–24 months. This implies that the sentiment shock is not purely directional in the short run: bullish narrative episodes also tend to coincide with temporarily greater narrative intensity, disagreement, or ambiguity. Still, the uncertainty response is smaller and less persistent than the own response of sentiment, so the shock remains predominantly a directional sentiment disturbance.

The impulse response pattern of a persistent increase in the real oil price, positive co-movement with world industrial production, a delayed increase in world oil production, persistent CPI pass-through, and a temporary positive spillover into uncertainty, is most consistent with the signature of a demand-side information shock. Bullish narrative sentiment shifts expectations about future oil-market conditions, raises current oil prices, and is followed by stronger real activity rather than contraction. These results support the demand-side information channel developed in Section 2 and are consistent with the view that oil-market narratives embed information about future demand conditions.

5.2 Narrative Uncertainty Shock

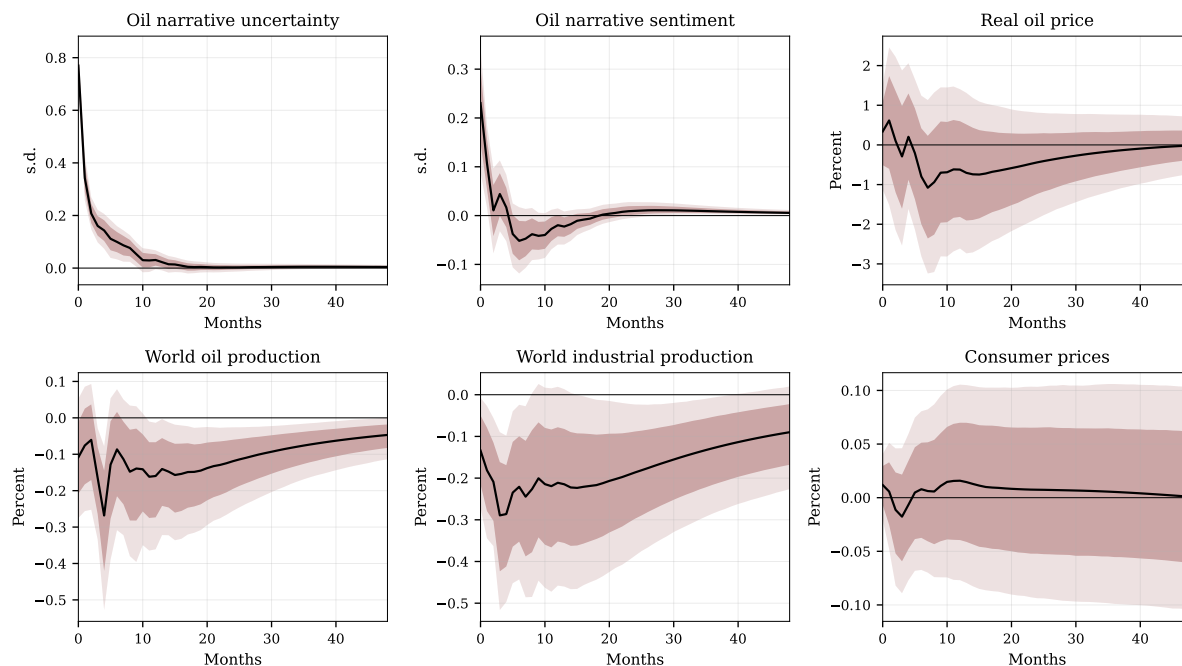
Figure 7 displays impulse responses to a one-standard-deviation narrative uncertainty shock, identified using the four purged type-specific uncertainty instruments in the Proxy-SVAR framework. The first-stage statistics indicate meaningful instrument relevance: $F = 27.0$ (robust: 13.9), effective $F = 12.1$, with $R^2 = 20.5\%$ and adjusted $R^2 = 19.7\%$.

The uncertainty shock generates a response pattern that is qualitatively distinct from the sentiment shock and is consistent with the real-options channel developed in Section 2.

World oil production. World oil production declines on impact and remains below baseline throughout the horizon. The response reaches its trough of roughly -0.25% around months 4–5, remains materially negative through the medium run, and then only gradually recovers, ending still slightly below zero at long horizons. This persis-

tent contraction is consistent with a wait-and-see response: when uncertainty about future oil-market conditions rises, producers delay drilling, investment, and other irreversible decisions. The response is therefore in line with the real-options mechanism emphasized in [Bernanke \(1983\)](#); [Bloom \(2009\)](#); [Kellogg \(2014\)](#).

Impulse responses to a one s.d. narrative uncertainty shock



First-stage: $F = 27.0$ (robust 13.9), Eff. $F = 12.1$, $R^2 = 20.5\%$, Adj. $R^2 = 19.7\%$

Figure 7: Impulse responses to a one standard deviation narrative uncertainty shock. Dark shaded region: 68% MBB confidence bands. Light shaded region: 90% MBB confidence bands. Solid line: point estimate. The shock is identified using four purged type-specific uncertainty instruments in the Proxy-SVAR framework. Sample: 1987:02–2022:12, VAR(6), 10,000 bootstrap replications.

World industrial production. World industrial production falls sharply and persistently following the uncertainty shock. The decline reaches roughly -0.28% to -0.30% in the first year, remains near -0.2% through the medium run, and then recovers only gradually, still remaining below baseline at the end of the horizon. This sustained contraction indicates that oil-market narrative uncertainty spills over to broader global activity, consistent with both reduced energy-sector investment and a wider confidence channel that dampens expenditure decisions.

Real oil price. The response of the real oil price is non-monotonic and imprecisely estimated, but its point estimate becomes negative after the first few months. Following some short-run oscillation around zero, the price declines to roughly -1% around months 6–8 and then gradually reverts toward zero from below. This pattern suggests

that the contractionary demand effects of uncertainty dominate any upward pressure that might arise from lower oil production. Accordingly, the uncertainty shock does not generate the clear positive price-output co-movement that characterizes the sentiment shock.

Consumer prices. The CPI response is small throughout and remains close to zero at all horizons. After a brief negative dip on impact, consumer prices show a slight positive movement around months 8–12, but the effect is economically small and fades back toward zero thereafter. The uncertainty shock therefore has little discernible pass-through to aggregate consumer prices, indicating that its dominant effects operate through real activity rather than through broad inflationary pressure.

Cross-effect on sentiment. The uncertainty shock has only a modest and transitory effect on oil narrative sentiment. Sentiment rises briefly on impact, turns mildly negative over roughly the next year, and then gradually returns toward zero by about months 18–24. This pattern suggests that heightened uncertainty can temporarily erode bullish tone, but the spillover is limited in magnitude and not especially persistent. The sentiment and uncertainty dimensions therefore remain empirically distinct, even though they are not completely orthogonal in the short run.

The impulse response pattern of persistent declines in world oil production and world industrial production, a weak and mostly negative real oil price response after the initial months, a negligible CPI response, and only a modest temporary spillover to sentiment, is consistent with the signature of a contractionary uncertainty shock operating through a real-options channel. Elevated narrative uncertainty causes producers and firms to defer activity, weakening both oil-sector supply and broader global output without generating the clear inflationary effects associated with a directional demand-side narrative shock.

5.3 When Sentiment and Uncertainty Rise Together: An Illustrative Stagflationary Pattern

The two preceding subsections show that narrative sentiment and narrative uncertainty shocks propagate through distinct channels. A bullish narrative sentiment shock raises oil prices, supports production and activity, and passes through gradually to consumer prices. A positive narrative uncertainty shock depresses production and activity, has only a weak and non-monotonic effect on oil prices, and generates little CPI pass-through. Many major oil-market episodes, however, are unlikely to reflect only one of these narrative dimensions in isolation. Instead, episodes such as wars, major geopolitical disruptions, or politically contentious supply events may simultaneously

generate a bullish directional revision in oil-market beliefs and heightened ambiguity about future market conditions.

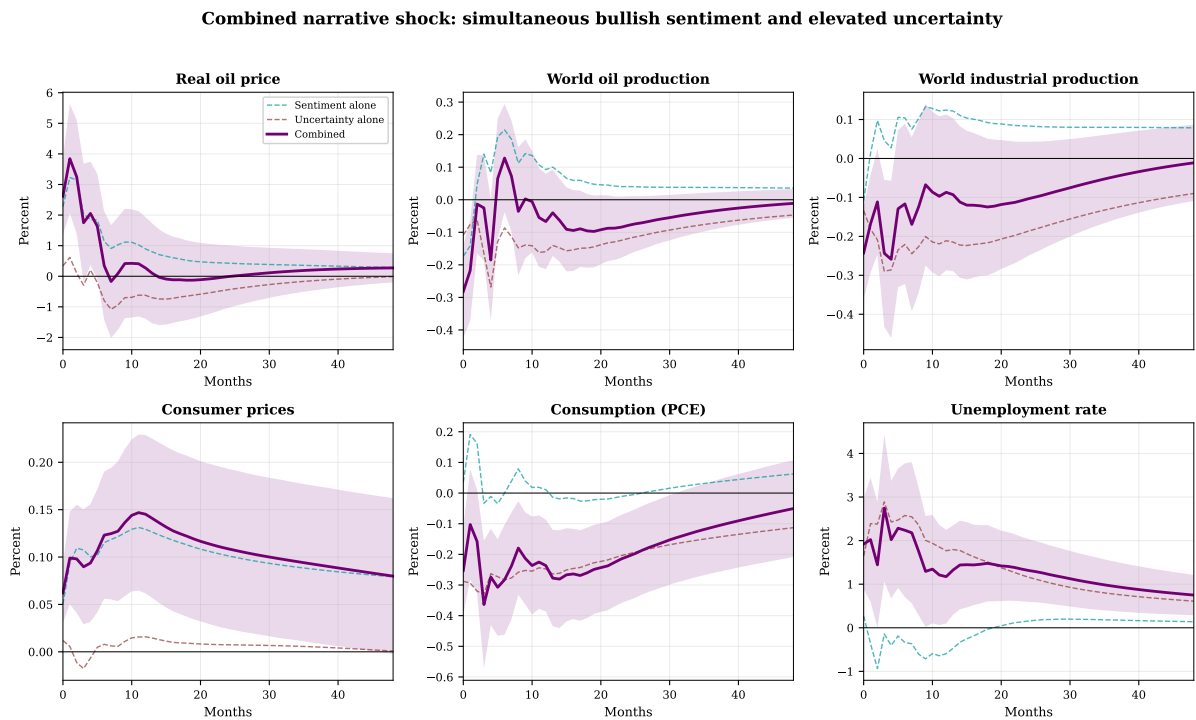


Figure 8: Illustrative linear aggregation of the separately estimated narrative sentiment and narrative uncertainty responses. Solid purple line: pointwise sum of the two individual impulse responses to a simultaneous one s.d. bullish sentiment shock and one s.d. uncertainty shock. Dashed teal: sentiment shock alone. Dashed dark red: uncertainty shock alone. Shaded region: approximate 68% confidence band for the aggregated response. The four baseline macro variables (real oil price, world oil production, world industrial production, consumer prices) are from the 6-variable Proxy-SVAR. Consumption (PCE) and the unemployment rate are from [Gertler and Karadi \(2015\)](#) marginal VARs augmenting the baseline system with one additional variable at a time. The figure is intended as an illustration of joint channel activation, not as a separately identified third structural shock. Sample: 1987:02–2022:12, VAR(6).

To illustrate the joint implications of these two forces, Figure 8 reports a linear aggregation of the separately estimated impulse responses to a one-standard-deviation bullish narrative sentiment shock and a one-standard-deviation narrative uncertainty shock. Specifically, the figure adds the two estimated responses horizon by horizon. This exercise is intended as an economic illustration rather than as a separate structural identification. It is useful for showing how the two estimated channels can combine during episodes in which both narrative sentiment and narrative uncertainty rise at the same time. For the four baseline macro variables, the responses come from the Proxy-SVAR. For personal consumption expenditure and the unemployment rate, the figure reports analogous responses from the [Gertler and Karadi \(2015\)](#) marginal VARs. The resulting pattern is stagflationary: inflation-sensitive variables rise, while real activity, spending, and labor-market conditions weaken.

Real oil price. The real oil price rises sharply in the short run, reaching nearly 4% within the first two months. The response then declines quickly, remains positive through the early months, dips slightly below zero around months 7–9, and gradually returns toward zero thereafter. This front-loaded pattern suggests that the directional component initially dominates price formation, while the contractionary effects associated with heightened uncertainty subsequently offset part of the initial price pressure.

World oil production. World oil production falls on impact, then rebounds and turns briefly positive around months 5–7, peaking at roughly 0.1%. This rebound is temporary. After the first year, production turns mildly negative again and remains below baseline through much of the medium run before slowly converging back toward zero. The production response therefore reflects the tension between the two estimated channels: stronger expected oil-market conditions encourage expansion, while heightened uncertainty raises the option value of waiting and restrains the supply response.

World industrial production. World industrial production declines throughout the horizon. The response reaches roughly -0.25% in the early months, partially recovers, but remains negative even at long horizons. This persistent weakness indicates that the uncertainty channel dominates the real-activity response. Thus, even when oil prices rise in the short run, the broader global economy contracts rather than expands.

Consumer prices. Consumer prices rise throughout the horizon. The response starts near $0.09\text{--}0.10\%$, peaks around $0.14\text{--}0.15\%$ at roughly months 10–12, and then declines only gradually, remaining positive even at the end of the horizon. This sustained CPI response is what gives the aggregated pattern its stagflationary character rather than the purely contractionary profile of the uncertainty shock in isolation.

Consumption and unemployment. The marginal VAR evidence sharpens this interpretation. Personal consumption expenditure falls immediately, reaches a trough of roughly -0.35% in the early months, and remains below baseline throughout the horizon, although the decline narrows gradually over time. The unemployment rate rises sharply, peaks near $2.7\text{--}2.8\%$ in the first several months, and remains elevated even at long horizons. Together, these responses suggest that episodes combining bullish oil-market narratives with heightened uncertainty can weaken household demand and labor-market conditions in a persistent way.

Interpretation. Taken as an illustrative aggregation, the response pattern is stagflationary: oil prices and consumer prices rise, while industrial production and con-

sumption fall and unemployment rises. Unlike a pure sentiment shock, the aggregated response does not generate jointly rising prices and activity. Unlike a pure uncertainty shock, it produces a clearer and more persistent CPI response. The figure therefore suggests that when bullish oil-market sentiment and elevated uncertainty are activated together, the two estimated channels can combine to generate a stagflationary macroeconomic pattern. This interpretation is consistent with the two-channel framework developed above, but it should not be read as a separately identified third structural shock.

5.4 The Real Options Channel: Evidence from Drilling Activity

The conceptual framework in Section 2 suggests that narrative uncertainty may transmit to the macroeconomy through a supply-side delay channel: heightened ambiguity raises the option value of waiting and leads firms to postpone irreversible drilling and investment decisions (Bernanke, 1983; Bloom, 2009; Kellogg, 2014). To assess this mechanism more directly, I augment the baseline VAR with the Baker Hughes U.S. oil-directed rig count following the marginal VAR approach of Gertler and Karadi (2015). Rig activity is a useful margin for this exercise because it lies closer to drilling and investment decisions than aggregate crude oil production, which adjusts more slowly and reflects a broader set of contemporaneous forces.




Figure 9: Rig count response to narrative oil market shocks (marginal VAR). Left panel: response to a one standard deviation narrative uncertainty shock. Right panel: response to a one standard deviation narrative sentiment shock. Dark shaded region: 68% MBB confidence bands. Light shaded region: 90% MBB confidence bands. Sample: 1987:08–2022:12. Robust first-stage F -statistics: 14.7 for the uncertainty-shock specification and 16.5 for the sentiment-shock specification.

Figure 9 reports the rig-count response to each narrative shock. The evidence for the uncertainty channel is strong. Following a one-standard-deviation narrative un-

certainty shock, rig activity declines quickly and persistently. The impact response is near zero, but the decline accelerates after two to three months, reaching roughly -2.5% within the first half-year and a trough near -4% around months 10–12, before recovering only gradually. Even at the end of the horizon, rig counts remain below baseline. This delayed but persistent contraction is consistent with a wait-and-see response: firms do not immediately shut down drilling, but as uncertainty remains elevated they scale back new drilling activity and defer investment.

The response to the sentiment shock is qualitatively different. Rig counts fall slightly on impact, then rise modestly over the next several months, peaking at about 1% around months 5–7. That increase is temporary: the response declines toward zero, turns mildly negative in the medium run, and only gradually returns toward baseline at longer horizons. Thus, bullish oil-market narratives do stimulate drilling activity in the short run, but the effect is considerably smaller and less persistent than the response to uncertainty.

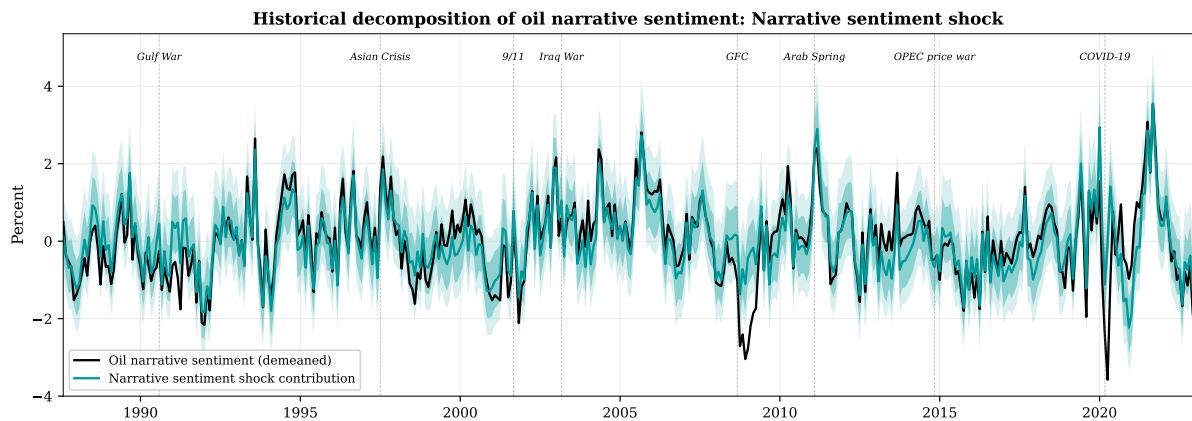
The rig-count responses provide direct support for the real-options interpretation of the uncertainty shock. Narrative uncertainty suppresses drilling activity in a large and persistent way, consistent with firms delaying irreversible investment when the oil-market outlook becomes more ambiguous. By contrast, narrative sentiment generates only a modest, temporary increase in rig deployment. This asymmetry strengthens the broader interpretation of the paper's main results: the contraction in oil production following an uncertainty shock reflects a genuine reduction in drilling activity, while the sentiment channel appears to operate more through expectations and prices than through a large sustained expansion of the drilling margin itself.

5.5 Historical Decomposition

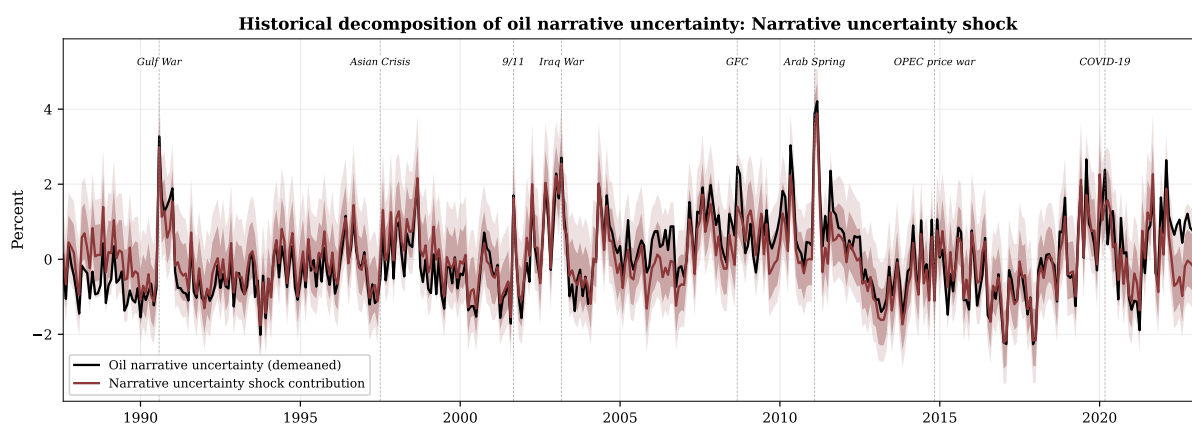
Figure 10 displays the historical contribution of each identified narrative shock to its own instrumented variable: the aggregate oil narrative sentiment index in Panel (a) and the aggregate oil narrative uncertainty index in Panel (b). These decompositions provide an event-validation exercise. If the identified shocks are economically meaningful, their estimated historical contributions should co-move with the realized narrative indices during major oil-market episodes.

Panel (a) shows the contribution of the narrative sentiment shock to the demeaned aggregate oil narrative sentiment index. The shock contribution tracks the realized index closely over much of the sample and captures the major swings in oil-market tone associated with key events, including the Gulf War, the Asian Financial Crisis, 9/11, the Iraq War, the Global Financial Crisis, the Arab Spring, the OPEC price war, and the COVID-19 collapse. The close co-movement between the identified contribution and the realized series suggests that the Proxy-SVAR isolates a quantitatively important

component of variation in aggregate oil-market sentiment rather than merely fitting a small residual fragment of the index.



(a) Narrative sentiment shock



(b) Narrative uncertainty shock

Figure 10: Historical decomposition of oil narrative indices. Solid black line: demeaned aggregate index. Colored line: estimated contribution of the identified structural shock. Shaded regions: 68% and 90% MBB confidence bands. Panel (a): contribution of the narrative sentiment shock to the aggregate oil narrative sentiment index. Panel (b): contribution of the narrative uncertainty shock to the aggregate oil narrative uncertainty index. Vertical dashed lines mark major oil-market events. Sample: 1987:02–2022:12, VAR(6), 10,000 bootstrap replications.

Panel (b) shows the contribution of the narrative uncertainty shock to the demeaned aggregate oil narrative uncertainty index. The decomposition highlights the episodic, spike-driven nature of oil-market uncertainty. Large positive contributions appear around the Gulf War, 9/11, the Iraq War, the Global Financial Crisis, the Arab Spring, and COVID-19, with the Arab Spring standing out as the largest uncertainty episode in the sample. Between major crises, the uncertainty shock contribution is typically smaller and less persistent, which is consistent with the idea that large movements in narrative uncertainty are concentrated around discrete geopolitical and crisis events.

The historical decompositions support the structural interpretation of the identi-

fied shocks. The sentiment shock explains a substantial share of the medium-run fluctuations in aggregate oil-market tone, while the uncertainty shock captures the major crisis-driven spikes in ambiguity and instability. This event-level alignment strengthens the interpretation that the estimated shocks recover economically meaningful narrative disturbances rather than purely statistical innovations.

5.6 Forecast Error Variance Decomposition

Figure 11 reports the forecast error variance shares associated with the narrative sentiment shock and the narrative uncertainty shock for each variable in the baseline system. Because the two shocks are identified in separate Proxy-SVAR specifications, the two lines in each panel should be interpreted as comparative variance shares rather than as components of a single decomposition that sum to one. The results nevertheless provide a useful summary of the relative quantitative importance of the two narrative channels across variables and horizons.

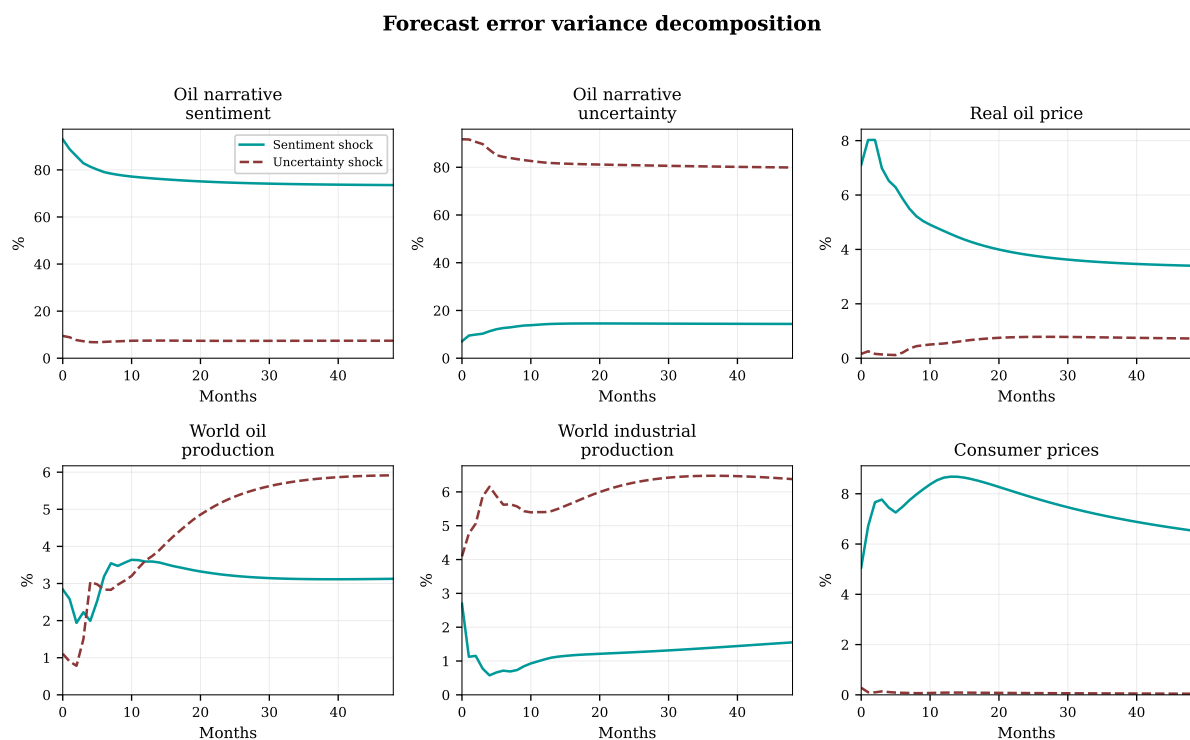


Figure 11: Forecast error variance decomposition. Each panel shows the share of forecast error variance attributable to the narrative sentiment shock (solid teal) and the narrative uncertainty shock (dashed dark red) at horizons 0–48 months. The sentiment and uncertainty shocks are identified from separate Proxy-SVAR systems. Sample: 1987:02–2022:12, VAR(6), 10,000 MBB replications.

The FEVDs indicate that each shock overwhelmingly explains variation in its own narrative index. The sentiment shock accounts for the large majority of the variance of aggregate oil narrative sentiment throughout the horizon, declining from above 90%

on impact to roughly the mid-70% range at longer horizons. Symmetrically, the uncertainty shock explains the vast majority of the variance of aggregate oil narrative uncertainty, remaining around 80% or more even at long horizons. Cross-contributions are present but limited: the sentiment shock explains roughly 10–15% of uncertainty variance, while the uncertainty shock explains roughly 7–10% of sentiment variance. This pattern confirms that the two identified shocks capture distinct but not completely orthogonal dimensions of oil-market narratives.

For real oil prices, the sentiment shock is clearly the dominant narrative disturbance. Its contribution is about 7–8% at very short horizons and remains economically meaningful even at long horizons, settling around 3–4%. By contrast, the uncertainty shock explains less than 1% of oil-price forecast error variance at all horizons. This difference is consistent with the impulse responses: directional narrative sentiment is the main narrative force moving oil prices, whereas uncertainty has only a limited direct role in explaining price fluctuations.

The pattern is very different for world industrial production. Here the uncertainty shock dominates at all horizons, accounting for roughly 4–6.5% of forecast error variance, while the sentiment shock contributes much less, falling below 1% after the initial months and only recovering gradually to around 1.5% at longer horizons. Thus, among the two narrative dimensions, uncertainty is quantitatively more important for explaining fluctuations in global real activity.

For world oil production, both shocks matter, but their relative importance changes with the horizon. The sentiment shock explains around 2–4% of forecast error variance and is somewhat larger in the short run. The uncertainty shock starts lower, around 1%, but rises steadily over time, overtaking sentiment after roughly one year and reaching about 6% at longer horizons. This horizon dependence is informative: sentiment matters more for short-run production adjustments associated with stronger expected market conditions, whereas uncertainty becomes increasingly important for longer-run production decisions, consistent with a real-options channel in which uncertainty delays irreversible drilling decisions, as confirmed directly by the rig-count evidence in Section 5.4.

Consumer prices are again dominated by the sentiment shock. The sentiment contribution rises quickly from about 5% on impact to roughly 8–9% at business-cycle horizons and remains above 6% even at the end of the horizon. The uncertainty shock contributes essentially nothing to CPI variance throughout. This pattern reinforces the interpretation that narrative sentiment passes through to broader prices, while narrative uncertainty mainly affects real-side outcomes rather than aggregate inflation.

Overall, the FEVD results reinforce the two-channel interpretation. Narrative sentiment is the quantitatively dominant narrative force for real oil prices and consumer prices, whereas narrative uncertainty is more important for real activity, especially

world industrial production and, at longer horizons, world oil production. The FEVD evidence therefore complements the impulse responses: sentiment primarily operates through a directional information channel linked to prices and inflation, while uncertainty operates through a real-options channel linked to production and broader activity.

6 Robustness

6.1 Placebo Tests

To confirm that the first-stage results reflect genuine narrative content rather than spurious correlation, I conduct two placebo exercises. Figure 12 reports the results.

First, I randomly permute each instrument across time 1,000 times, thereby destroying its temporal relationship with the VAR residuals while preserving its marginal distribution. The resulting placebo F -statistics are concentrated at low values in both systems, indicating that strong first-stage relationships do not arise mechanically from the distribution of the instruments alone. By contrast, the actual first-stage statistics lie far in the right tail of the placebo distributions: $F = 16.7$ in the sentiment system and $F = 27.0$ in the uncertainty system.




Figure 12: Placebo tests: permutation and cross-system instrument swap. Grey histograms: distribution of first-stage F -statistics from 1,000 random permutations of the instruments. Solid vertical line: actual F -statistic from the correctly assigned instruments. Dashed vertical line: F -statistic from swapping the instrument sets across systems (uncertainty instruments applied to the sentiment system, and vice versa). Sample: 1987:02–2022:12.

Second, I swap the instrument sets across systems, using the four uncertainty instruments to instrument sentiment and the four sentiment instruments to instrument uncertainty. The resulting cross-system first-stage statistics are small: $F = 3.1$ when the uncertainty instruments are applied to the sentiment system, and $F = 1.4$ when

the sentiment instruments are applied to the uncertainty system. Both are far below the correctly assigned first-stage values and lie much closer to the placebo region than to the actual estimates.

The placebo tests indicate that the identifying content of the instruments is not driven by generic oil-news variation. Sentiment instruments predict sentiment but have little power for uncertainty, while uncertainty instruments predict uncertainty but not sentiment. This supports the view that the two identified shocks draw on genuinely distinct dimensions of oil-market narratives rather than on a single common news factor.

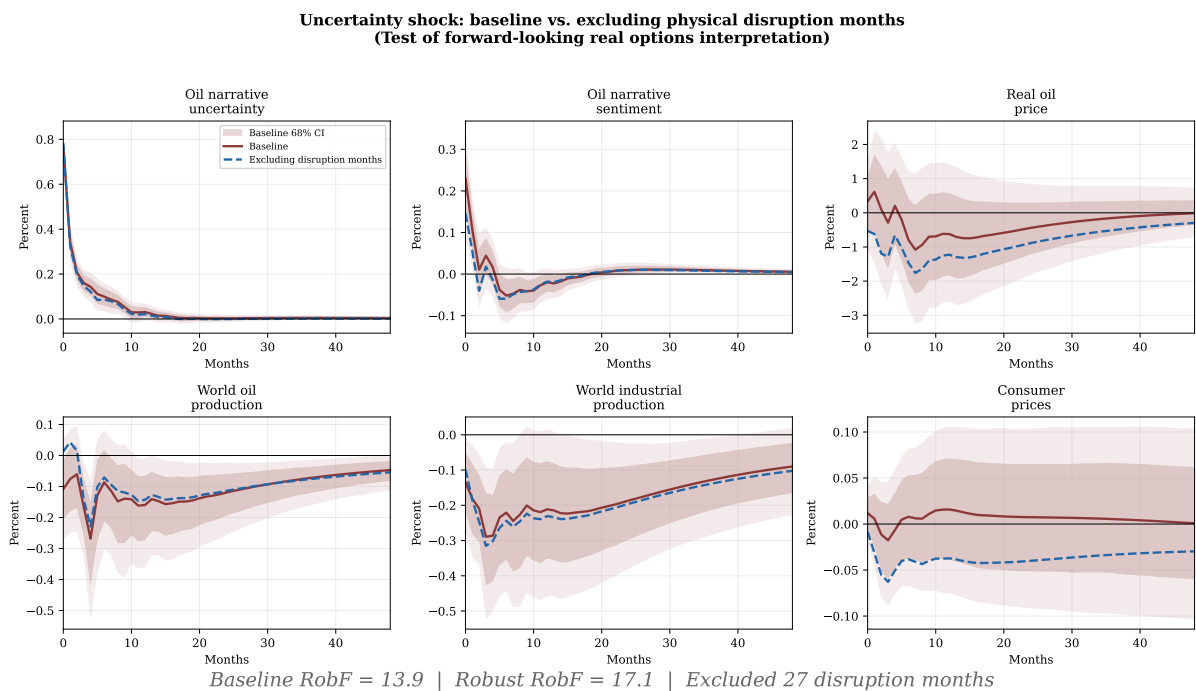


Figure 13: Uncertainty shock: baseline responses versus estimates excluding major physical disruption months. Solid line: baseline response. Dashed line: response estimated after excluding 27 disruption months. Shaded region: 68% confidence band for the baseline response. Sample: 1987:02–2022:12. Baseline robust first-stage $F = 13.9$; restricted-sample robust first-stage $F = 17.1$.

6.2 Physical Disruption Episodes

The contemporaneous decline in world oil production following an uncertainty shock could, in principle, reflect realized physical disruptions during major supply episodes rather than forward-looking adjustment. To address this concern, I re-estimate the uncertainty Proxy-SVAR after excluding 27 months associated with major physical oil supply disruptions: the Gulf War (1990:08–1991:03), Hurricanes Katrina and Rita (2005:08–2005:10), the Libyan civil war (2011:02–2011:09), Hurricanes Harvey and Irma (2017:08–2017:10), and the Russia–Ukraine invasion (2022:02–2022:06).

Figure 13 compares the baseline impulse responses with those obtained in this restricted sample. The main results are largely unchanged. Most importantly, world oil production still falls sharply on impact, remains below baseline through the medium run, and only gradually recovers thereafter following an uncertainty shock. World industrial production also continues to decline persistently, with the restricted-sample responses, if anything, somewhat more negative than in the baseline specification.

The price responses are particularly informative. If the baseline uncertainty shock primarily reflected contemporaneous realized supply outages, excluding the major disruption episodes should attenuate the decline in oil production and, if anything, strengthen the positive price response associated with reduced supply. Instead, the opposite occurs. Once these months are excluded, the real oil price response becomes more negative rather than less, while the CPI response turns mildly negative rather than remaining near zero. This pattern is difficult to reconcile with a purely realized-disruption interpretation. Instead, it suggests that the dominant effect of narrative uncertainty is to depress expected activity and delay production and investment decisions, so that contractionary demand effects outweigh any upward price pressure arising from reduced supply.

The first stage also remains strong in the restricted sample. The robust first-stage statistic rises from 13.9 in the baseline specification to 17.1 after excluding disruption months, indicating that the identifying variation in the uncertainty instruments is not driven by a small number of extreme physical supply episodes.

These results support interpreting the uncertainty shock as a forward-looking real-options shock rather than as a proxy for realized physical disruptions. The negative production response is not an artifact of contemporaneous wars, hurricanes, or other realized supply disruptions. Instead, narrative uncertainty appears to operate by raising the option value of waiting, thereby delaying drilling and production decisions and generating persistent contractions in oil-sector and broader real activity.

6.3 Additional Robustness Checks

I report seven additional specification checks in Appendix E.

Proxy-SVAR versus local projections. Appendix Figure 16 compares the baseline Proxy-SVAR impulse responses with LP-IV estimates constructed from the same external instruments following [Stock and Watson \(2018\)](#). For both the sentiment and uncertainty shocks, the LP-IV responses are close to the Proxy-SVAR estimates, with larger differences appearing only at longer horizons where the LP-IV estimates become less precise. The main patterns are unchanged: the sentiment shock raises oil prices and real activity, whereas the uncertainty shock lowers production and activity

while leaving the price response comparatively muted. This similarity suggests that the baseline results are not driven by the VAR dynamics imposed by the Proxy-SVAR.

Alternative lag order. The baseline specification uses $p = 6$. Re-estimating the model with $p = 3$ and $p = 12$ yields the same qualitative responses for both shocks, with only modest differences in the width of the confidence bands at longer horizons (Appendix Figures 17).

Cholesky identification. As a benchmark, I also estimate the VAR under recursive (Cholesky) identification using the ordering (SENT, UNC, ROP, WOP, WIP, CPI). For the sentiment shock, the resulting impulse responses are broadly similar to the Proxy-SVAR estimates. For the uncertainty shock, however, the recursive responses are less precise and harder to interpret, which is not surprising given the stronger identifying content of the external instruments in that case. I therefore prefer the Proxy-SVAR specification because it does not rely on a recursive ordering and instead uses external instrument variation for identification (Appendix Figure 18).

Leave-one-out instruments. Re-estimating each Proxy-SVAR four times, each time dropping one of the four instruments, produces point estimates that remain close to the baseline and within the 90% confidence bands in every case. In the sentiment system, all four leave-one-out specifications maintain joint F -statistics above 10, ranging from 12.9 (dropping transition) to 20.7 (dropping supply). In the uncertainty system, three of four specifications maintain robust F above 10; dropping geopolitical reduces the robust F to 9.2, indicating that geopolitical uncertainty is the most individually important instrument for the uncertainty identification. No single instrument drives the qualitative results in either system (Appendix Figure 19).

Including OPEC. Re-estimating both systems with OPEC included as a fifth instrument produces qualitatively similar impulse responses. The sentiment system overid test becomes borderline (system $J = 30.20$, $p = 0.067$), driven by the real-oil-price equation ($J = 14.31$, $p = 0.006$), consistent with the contemporaneous-feedback argument in Section 4.3. The uncertainty system passes overid cleanly with OPEC included ($p = 0.671$). The structural conclusions are unchanged: sentiment raises oil prices and activity, uncertainty depresses production and activity (Appendix Figure 20).

Excluding COVID-19. The COVID-19 pandemic produced the largest negative sentiment reading in the sample (-3.57σ) and extreme movements in oil prices, production, and activity. To verify that the two-channel results are not driven by this single episode, I re-estimate both Proxy-SVARs after excluding the 18-month window

from January 2020 through June 2021. The first-stage statistics remain strong in the restricted sample (sentiment robust $F = 17.2$; uncertainty robust $F = 14.7$). The impulse responses are qualitatively unchanged: the sentiment shock continues to raise oil prices, production, and activity, while the uncertainty shock continues to depress production and activity with a muted price response. The two-channel transmission mechanism therefore operates throughout the non-pandemic portion of the sample (Appendix Figure 21).

Alternative aggregate measures. To assess whether the two-channel result depends on the specific text-scoring methodology, I re-estimate both Proxy-SVARs replacing the aggregate narrative indices with independently constructed external measures, retaining the same four type-specific instruments. For uncertainty, I substitute the [Abiad and Qureshi \(2023\)](#) Oil Price Uncertainty index, which correlates with the baseline narrative uncertainty index at $\rho = +0.47$. The sentiment instruments have no predictive power for the alternative uncertainty measure ($F = 0.8$), confirming channel specificity, while the uncertainty instruments predict the OPU index with $F = 7.4$. The resulting impulse responses reproduce the supply-side contraction pattern: declining real oil prices, falling world industrial production, and a muted CPI response. The qualitative agreement between the baseline and alternative uncertainty specifications indicates that the real-options channel is a feature of the underlying economic environment, not of the specific uncertainty scoring procedure (Appendix Figure 22).

7 Conclusion

This paper asks whether oil-market narratives have macroeconomic effects, and whether those effects differ depending on whether narratives convey directional information or ambiguity. The answer to both questions is yes.

The main finding is that bullish narrative sentiment and narrative uncertainty behave like different types of macroeconomic disturbances. A bullish narrative sentiment shock raises the real price of oil, increases world industrial production, raises world oil production, and gradually passes through to consumer prices. This pattern is consistent with a demand-side information channel in which bullish oil-market narratives shift expectations about future scarcity or future demand. By contrast, a narrative uncertainty shock reduces oil production and real activity while producing only a muted and weakly estimated response of the real oil price and consumer prices. This pattern is consistent with a supply-side delay channel in which heightened ambiguity raises the option value of waiting and slows the adjustment of partially irreversible production and investment decisions.

These results suggest that oil-market narratives contain macroeconomically relevant variation beyond what is captured by standard physical oil shocks alone. More broadly, they imply that existing oil-shock frameworks may miss an important distinction between oil-market news that shifts directional expectations and oil-market news that primarily raises ambiguity. Distinguishing between these two forms of narrative variation helps explain why some episodes of intense oil-market coverage generate price-output co-movement, while others primarily suppress production and activity without producing a comparably strong price response.

These results also imply that episodes in which narrative sentiment and narrative uncertainty rise together can generate a macroeconomic pattern that neither shock produces in isolation. When bullish oil-market narratives coincide with elevated ambiguity, stronger demand expectations push oil prices upward while uncertainty slows the adjustment of production and investment. The result is a stagflationary combination of rising oil prices, weak production, softer real activity, and higher inflation. This perspective helps interpret episodes of intense oil-market coverage in which prices surge but production and activity respond only weakly.

The broader implication is that narratives matter not only because they summarize oil-market events, but because they shape how those events are processed by firms, investors, and policymakers. Directional tone and ambiguity are not the same object, and they do not have the same macroeconomic consequences. Recognizing that distinction helps clarify how oil news transmits to the broader economy.

References

- Abiad, Abdul and Irfan A Qureshi**, "The macroeconomic effects of oil price uncertainty," *Energy Economics*, 2023, 125, 106839.
- Agerton, Mark, Peter R Hartley, Kenneth B Medlock III, and Ted Temzelides**, "Employment impacts of upstream oil and gas investment in the United States," *Energy Economics*, 2017, 62, 171–180.
- Alessandri, Piergiorgio and Andrea Gazzani**, "Natural gas and the macroeconomy: not all energy shocks are alike," *Journal of Monetary Economics*, 2025, 151, 103749.
- Angrist, Joshua D. and Alan B. Krueger**, "Does Compulsory School Attendance Affect Schooling and Earnings?," *Quarterly Journal of Economics*, 1991, 106 (4), 979–1014.
- Aruoba, S Borağan and Thomas Drechsel**, "Identifying monetary policy shocks: A natural language approach," Technical Report, National Bureau of Economic Research 2024.
- Bachmann, Ruediger, Giorgio Topa, and Wilbert van der Klaauw**, *Handbook of economic expectations*, Elsevier, 2022.
- Baker, Scott R, Nick Bloom, and Steven J Davis**, "Measuring economic policy uncertainty," *Quarterly Journal of Economics*, 2016, 131 (4), 1593–1636.
- Barsky, Robert B. and Lutz Kilian**, "Do We Really Know That Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual*, 2002, 16, 137–183.
- Baumeister, Christiane and James D Hamilton**, "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 2019, 109 (5), 1873–1910.
- , **Lutz Kilian, and Xiaoqing Zhou**, "Are Product Spreads Useful for Forecasting? An Empirical Evaluation of the Verleger Hypothesis," *Macroeconomic Dynamics*, 2018, 22 (3), 562–580.
- Beaudry, Paul and Franck Portier**, "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 2006, 96 (4), 1293–1307.
- Bernanke, Ben S**, "Irreversibility, uncertainty, and cyclical investment," *Quarterly Journal of Economics*, 1983, 98 (1), 85–106.
- Blanchard, Olivier J. and Jordi Galí**, "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?," 2007, pp. 373–421.

- Bloom, Nick**, “The impact of uncertainty shocks,” *Econometrica*, 2009, 77 (3), 623–685.
- Bybee, Leland, Bryan T. Kelly, Asaf Manela, and Dacheng Xiu**, “Business News and Business Cycles,” *The Journal of Finance*, 2024, 79 (5), 3105–3147.
- Caldara, Dario and Matteo Iacoviello**, “Measuring geopolitical risk,” *American Economic Review*, 2022, 112 (4), 1194–1225.
- Dixit, Avinash K. and Robert S. Pindyck**, *Investment under Uncertainty*, Princeton University Press, 1994.
- Elder, John and Apostolos Serletis**, “Oil Price Uncertainty,” *Journal of Money, Credit and Banking*, 2010, 42 (6), 1137–1159.
- Ellsberg, Daniel**, “Risk, Ambiguity, and the Savage Axioms,” *Quarterly Journal of Economics*, 1961, 75 (4), 643–669.
- Gately, Dermot**, “Lessons from the 1986 Oil Price Collapse,” *Brookings Papers on Economic Activity*, 1986, 1986 (2), 237–284.
- Gavriilidis, Konstantinos, Diego R. K'anzig, Ramya Raghavan, and James H. Stock**, “The Macroeconomic Effects of Climate Policy Uncertainty,” NBER Working Paper 34762, National Bureau of Economic Research 2026.
- Gazzani, Andrea, Fabrizio Venditti, and Giovanni Veronese**, “Oil price shocks in real time,” *Journal of Monetary Economics*, 2024, 144, 103547.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy**, “Text as data,” *Journal of Economic Literature*, 2019, 57 (3), 535–574.
- Gertler, Mark and Peter Karadi**, “Monetary policy surprises, credit costs, and economic activity,” *American Economic Journal: Macroeconomics*, 2015, 7 (1), 44–76.
- Gilboa, Itzhak and David Schmeidler**, “Maxmin Expected Utility with Non-unique Prior,” *Journal of Mathematical Economics*, 1989, 18 (2), 141–153.
- Hamilton, James D.**, “What Is an Oil Shock?,” *Journal of Econometrics*, 2003, 113 (2), 363–398.
- , “Causes and Consequences of the Oil Shock of 2007–08,” *Brookings Papers on Economic Activity*, 2009, 2009 (1), 215–261.
- , “Historical Oil Shocks,” NBER Working Paper 16790, National Bureau of Economic Research 2011. Working-paper version of the later Routledge handbook chapter.

- , “Historical Oil Shocks,” in Randall E. Parker and Robert Whaples, eds., *Routledge Handbook of Major Events in Economic History*, Routledge, 2013, pp. 239–265.
- Hansen, Lars Peter**, “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 1982, 50 (4), 1029–1054.
- Jentsch, Carsten and Kurt G Lunsford**, “Asymptotically valid bootstrap inference for proxy SVARs,” *Journal of Business & Economic Statistics*, 2022, 40 (4), 1876–1891.
- Jo, Soojin**, “The Effects of Oil Price Uncertainty on Global Real Economic Activity,” *Journal of Money, Credit and Banking*, 2014, 46 (6), 1113–1135.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng**, “Measuring Uncertainty,” *American Economic Review*, 2015, 105 (3), 1177–1216.
- Känzig, Diego R**, “The macroeconomic effects of oil supply news: Evidence from OPEC announcements,” *American Economic Review*, 2021, 111 (4), 1092–1125.
- Kellogg, Ryan**, “The effect of uncertainty on investment: Evidence from Texas oil drilling,” *American Economic Review*, 2014, 104 (6), 1698–1734.
- Kilian, Lutz**, “Exogenous oil supply shocks: how big are they and how much do they matter for the US economy?,” *The review of economics and statistics*, 2008, 90 (2), 216–240.
- , “Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market,” *American Economic Review*, 2009, 99 (3), 1053–1069.
- **and Daniel P Murphy**, “The role of inventories and speculative trading in the global market for crude oil,” *Journal of Applied Econometrics*, 2014, 29 (3), 454–478.
- **and Helmut Lütkepohl**, *Structural Vector Autoregressive Analysis*, Cambridge University Press, 2017.
- Larsen, Vegard H. and Leif Anders Thorsrud**, “The Value of News for Economic Developments,” *Journal of Econometrics*, 2019, 210 (1), 203–218.
- Loughran, Tim and Bill McDonald**, “When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks,” *Journal of Finance*, 2011, 66 (1), 35–65.
- , – , **and Ioannis Pravidis**, “Assimilation of oil news into prices,” *International Review of Financial Analysis*, 2019, 63, 105–118.
- Mankiw, N. Gregory and Ricardo Reis**, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 2002, 117 (4), 1295–1328.

- McCracken, Michael W and Serena Ng**, “FRED-MD: A monthly database for macroeconomic research,” *Journal of Business & Economic Statistics*, 2016, 34 (4), 574–589.
- Mertens, Karel and Morten O Ravn**, “The dynamic effects of personal and corporate income tax changes in the United States,” *American Economic Review*, 2013, 103 (4), 1212–1247.
- Miranda-Agrippino, Silvia and Giovanni Ricco**, “The transmission of monetary policy shocks,” *American Economic Journal: Macroeconomics*, 2021, 13 (3), 74–107.
- Olea, José Luis Montiel and Carolin E Pflueger**, “A robust test for weak instruments,” *Journal of Business & Economic Statistics*, 2013, 31 (3), 358–369.
- , **James H. Stock, and Mark W. Watson**, “Inference in Structural Vector Autoregressions Identified with an External Instrument,” *Journal of Econometrics*, 2021, 225 (1), 74–87.
- Ramey, Valerie A**, “Can government purchases stimulate the economy?,” *Journal of Economic Literature*, 2011, 49 (3), 673–685.
- Romer, Christina D and David H Romer**, “A new measure of monetary shocks: Derivation and implications,” *American Economic Review*, 2004, 94 (4), 1055–1084.
- Sargan, J. Denis**, “The Estimation of Economic Relationships Using Instrumental Variables,” *Econometrica*, 1958, 26 (3), 393–415.
- Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson**, “Measuring news sentiment,” *Journal of Econometrics*, 2022, 228 (2), 221–243.
- Shiller, Robert J**, “Narrative economics,” *American Economic Review*, 2017, 107 (4), 967–1004.
- Sims, Christopher A.**, “Implications of Rational Inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.
- Stock, James H. and Mark W. Watson**, “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *The Economic Journal*, 2018, 128 (610), 917–948.
- **and Motohiro Yogo**, “Testing for Weak Instruments in Linear IV Regression,” in Donald W. K. Andrews and James H. Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 2005, pp. 80–108.
- Tetlock, Paul C.**, “Giving Content to Investor Sentiment: The Role of Media in the Stock Market,” *Journal of Finance*, 2007, 62 (3), 1139–1168.

Online Appendix

Not All Oil Narrative Shocks Are Alike: The Macroeconomic Effects of News-Driven Sentiment and Uncertainty Shocks

Gabriel A. Abeyie

Cameron School of Business
University of North Carolina Wilmington

April 21, 2026

This appendix provides supplementary material for the main paper. The contents are organized as follows.

Section	Contents
A Classification Dictionaries	Complete keyword lists for the five narrative type sub-dictionaries
B LLM Validation	Independent validation of the oil-relevance filter and narrative classification using a large language model
C Instrument Diagnostics	Granger non-causality tests and Ljung-Box serial correlation tests for raw vs. purged instruments
D Incremental Identification	Correlations between the identified narrative shocks and twelve existing structural shock measures
E Additional Robustness Checks	LP-IV comparison, alternative lag orders, Cholesky identification, leave-one-out instruments, OPEC inclusion, COVID-19 exclusion, and alternative aggregate uncertainty measure

A Classification Dictionaries

Table 7 reports the complete keyword lists for each of the five narrative type sub-dictionaries used to classify oil-relevant articles. An article is assigned to a narrative type if it contains at least two keyword matches (with title matches weighted three times body-text matches). Articles that do not meet this threshold for any type are classified as *general oil*.

Table 7: Narrative Classification Dictionaries

Narrative type	Keywords
Geopolitical conflict (52 terms)	invasion, invade, military strike, air strike, airstrike, missile, troops, armed forces, military intervention, cease-fire, ceasefire, warfare, war zone, civil war, proxy war, oil sanctions, economic sanctions, trade sanctions, sanctions on Iran, sanctions on Russia, oil embargo, trade embargo, export ban, import ban, Strait of Hormuz, Persian Gulf crisis, Gulf War, Iraq War, Iraq invasion, Kuwait invasion, Libyan conflict, Libya crisis, Syrian conflict, Yemen conflict, Iran nuclear, Iranian nuclear, Crimea, Ukraine conflict, Ukraine war, Ukraine invasion, Russian invasion, Russian aggression, Middle East conflict, Middle East crisis, Middle East tension, Arab Spring, regime change, terrorist attack, terrorism, geopolitical risk, geopolitical tension, geopolitical crisis, geopolitical instability, political instability, territorial dispute
Supply disruption (42 terms)	supply disruption, production disruption, output disruption, supply outage, production outage, unplanned outage, pipeline disruption, pipeline explosion, pipeline leak, pipeline shut, pipeline closure, refinery outage, refinery explosion, refinery fire, refinery shutdown, refinery shut down, platform shutdown, platform evacuat, shut-in production, shut in production, force majeure, hurricane, tropical storm, typhoon, cyclone, flooding, earthquake, natural disaster, Gulf of Mexico storm, Gulf of Mexico hurricane, oil spill, oil leak, tanker accident, Suez Canal, Bab el-Mandeb, chokepoint, infrastructure damage, sabotage, supply shortage, oil shortage, gasoline shortage, diesel shortage, heating oil shortage, inventory drawdown, stockpile decline, strategic petroleum reserve, SPR release
OPEC coordination (40 terms)	OPEC meeting, OPEC summit, OPEC decision, OPEC agreement, OPEC deal, OPEC accord, OPEC output, OPEC production, OPEC quota, OPEC compliance, OPEC cut, production cut, output cut, OPEC increase, production increase, output increase, OPEC freeze, production freeze, output freeze, OPEC plus, OPEC+, OPEC strategy, OPEC policy, OPEC target, OPEC market share, OPEC price war, OPEC spare capacity, spare capacity, Saudi production, Saudi output, Saudi oil, Saudi Arabia oil, Saudi Aramco, oil minister, energy minister, cartel, cartel discipline, quota cheat, overproduction, production ceiling, voluntary cut, coordinated cut, supply management, market balance
Speculation/ Finan. (52 terms)	oil futures, crude futures, futures contract, futures market, futures trading, futures price, commodity futures, energy futures, oil options, crude options, options trading, oil derivative, energy derivative, oil swap, commodity swap, oil ETF, commodity ETF, energy ETF, speculator, speculat, speculative position, speculative demand, speculative buying, speculative bubble, oil bubble, commodity speculation, excessive speculation, position limits, financialization, financialisation, index fund, index investor, commodity index, passive investment, commodity investment, hedge fund, commodity trading advisor, managed money, non-commercial, commodity trading, algorithmic trading, contango, backwardation, roll yield, margin call, margin requirement, open interest, trading volume, long position, short position, net long, net short, CFTC, commitment of traders
Energy transition (64 terms)	climate change, global warming, greenhouse gas, carbon emission, carbon dioxide, CO2 emission, carbon tax, carbon price, carbon trading, emission trading, cap and trade, cap-and-trade, Paris Agreement, Paris Accord, climate accord, Kyoto Protocol, COP26, COP27, COP28, COP29, net zero, net-zero, carbon neutral, decarbonization, decarbonisation, climate policy, climate regulation, ESG, environmental social governance, renewable energy, solar energy, solar power, wind energy, wind power, wind farm, electric vehicle, EV adoption, battery electric, hydrogen fuel, green hydrogen, biofuel, ethanol mandate, clean energy, green energy, energy transition, low carbon, peak oil demand, oil demand peak, stranded asset, fossil fuel divestment, keep it in the ground, oil demand destruction, demand destruction, Inflation Reduction Act, European Green Deal, EU Green Deal, carbon border adjustment, methane emission, methane emissions, methane regulation, carbon capture, direct air capture, offshore wind, sustainable aviation fuel

B LLM Validation of Dictionary Classification

To assess the accuracy of the dictionary-based classification, I compare it against independent classifications from a large language model (Claude Sonnet 4).

B.1 Oil-Relevance Filter

Panel A of Table 8 reports the oil-relevance validation. A stratified random sample of 500 articles (250 dictionary-flagged as oil-relevant, 250 not) is independently classified by the LLM. Each article’s text is truncated to the first 2,000 words. The dictionary achieves 87% accuracy, 81.2% precision, and 91.9% recall against the LLM benchmark, with a false positive rate of 18.8%. These results are consistent with rates typically observed for text-based indices in this literature (Baker et al., 2016; Caldara and Iacoviello, 2022; Gavriilidis et al., 2026).

The high recall (91.9%) confirms that the dictionary rarely misses articles the LLM considers oil-relevant. The 18.8% false positive rate reflects the dictionary’s broader net: it flags articles that mention oil terms anywhere in the full text, while the LLM evaluates a truncated version. For the paper’s purposes, high recall is more important than high precision, since type-specific sentiment and uncertainty scores are averaged across all flagged articles within each month, and occasional marginal inclusions are diluted by the large volume of clearly oil-relevant coverage.

B.2 LLM Classification Validation

Narrative type classification. Panel B reports type-specific validation. For each of 500 stratified oil-relevant articles (approximately 80 per specific type plus 100 general), the LLM is asked five separate binary questions corresponding to the five narrative categories. All five types achieve alignment rates above 65%, with energy transition highest at 92.5%, followed by OPEC at 75%, geopolitical at 72.5%, speculation at 70%, and supply at 67.5%. The variation across types reflects differences in how distinctive each category’s vocabulary is: transition and OPEC narratives use highly specific terminology that both dictionary and LLM identify consistently, while speculation and supply narratives overlap more with general oil-market reporting.

Figure 14 displays the cross-contamination matrix. The off-diagonal patterns are economically interpretable: OPEC articles are frequently also flagged as geopolitical, geopolitical articles as supply-related, and supply articles as geopolitical. These overlaps reflect the multi-dimensional nature of oil-market narratives. An OPEC production cut announced during a geopolitical crisis contains both coordination and geopolitical content; a hurricane that disrupts Gulf Coast refineries is both a supply event and a geopolitical concern. The dictionary assigns a single dominant type based on relative keyword counts, while the LLM detects all relevant dimensions through contextual reading. This multi-dimensionality does not threaten the identification strategy. The Proxy-SVAR exploits the fact that the type-specific indices capture economically distinct sources of narrative variation, and the leave-one-out analysis in Appendix E.4 confirms that no single instrument drives the results.

Table 8: LLM Validation of Dictionary-Based Classification*Panel A: Oil-Relevance Filter*

Metric	Value
Sample size	500
Accuracy	87.0%
Precision	81.2%
Recall	91.9%
F_1 score	86.2%
False positive rate	18.8%

Panel B: Narrative Type Classification (Binary LLM Prompts)

Type	N (dictionary)	LLM agrees	Alignment	FP rate
Geopolitical	80	58	72.5%	27.5%
Opec	80	60	75.0%	25.0%
Speculation	80	56	70.0%	30.0%
Supply	80	54	67.5%	32.5%
Transition	80	74	92.5%	7.5%

Notes: Panel A reports agreement between the dictionary-based oil-relevance filter and Claude Haiku 4.5 on a stratified sample of 500 articles (250 oil-relevant, 250 not). Panel B reports type-specific validation using binary LLM prompts. For each of 500 stratified oil-relevant articles (approximately 80 per specific type), the LLM is asked five separate binary questions corresponding to the five narrative types. Alignment is the share of dictionary-classified articles where the LLM agrees. False positive (FP) rate is the share where the LLM disagrees.




Figure 14: Cross-contamination matrix: dictionary vs. LLM classification. Each cell shows the share of articles classified as the row type by the dictionary that the LLM flags as the column type. Diagonal entries correspond to alignment rates in Table 8.

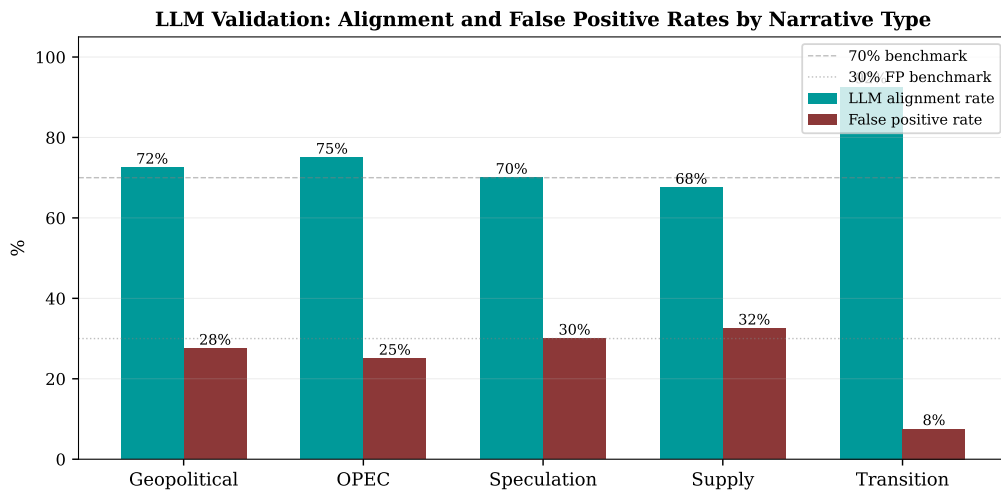


Figure 15: Alignment and false positive rates by narrative type. Teal: alignment rate. Dark red: false positive rate.

C Instrument Diagnostics

C.1 Granger Non-Causality Test

Table 9 reports Granger non-causality tests for the raw and purged instruments. The tests assess whether lagged macroeconomic state variables predict the instruments. Before purging, several raw instruments exhibit significant forecastability, as expected. After purging, no instrument is significantly predicted by lagged state variables at the 5% level, confirming that the purging procedure successfully removes predictable feedback from past oil-market and macroeconomic conditions.

Table 9: Granger Non-Causality Tests: Raw vs. Purged Instruments

Instrument	Raw (unpurged)		Purged	
	<i>F</i> -stat	<i>p</i> -value	<i>F</i> -stat	<i>p</i> -value
<i>Sentiment instruments</i>				
Geopolitical	1.08	0.367	0.00	1.000
Speculation	2.12	0.002	0.01	1.000
Supply	1.59	0.040	0.02	1.000
Transition	0.72	0.832	0.01	1.000
<i>Uncertainty instruments</i>				
Geopolitical	0.60	0.933	0.01	1.000
Speculation	1.58	0.041	0.00	1.000
Supply	0.87	0.642	0.02	1.000
Transition	1.13	0.307	0.20	1.000
<i>Joint tests (Fisher χ^2)</i>				
Sentiment (4 instruments)	21.52	0.006	0.00	1.000
Uncertainty (4 instruments)	9.79	0.280	0.00	1.000
All (8 instruments)	31.31	0.012	0.00	1.000

Notes: *F*-statistics from regressions of each instrument on 6 lags of $\Delta \log \text{WIP}$, $\Delta \log \text{CPI}$, $\Delta \log \text{ROP}$, and $\Delta \log \text{WOP}$. Bold *p*-values indicate rejection at the 5% level. Purging regresses each instrument on 12 own lags and 12 lags of the four state variables; the residual replaces the raw instrument. Joint tests use Fisher's method ($-2 \sum \log p_i \sim \chi^2(2k)$). Sample: 1987:02–2022:12.

C.2 Serial Correlation

Table 10 reports Ljung-Box Q tests for serial correlation in the raw and purged instruments at lags 6, 12, 18, and 24. The purged instruments show substantially reduced serial correlation relative to the raw series, consistent with the removal of own-lag predictability in the purging step.

Table 10: Autocorrelation Diagnostics: Ljung-Box Q Tests

Instrument	Q(6)		Q(12)		Q(18)		Q(24)	
	Raw	Purged	Raw	Purged	Raw	Purged	Raw	Purged
<i>Sentiment instruments</i>								
Geopolitical	0.003	1.000	0.002	1.000	0.010	0.997	0.009	0.990
Speculation	0.000	1.000	0.000	1.000	0.000	1.000	0.000	0.704
Supply	0.020	0.995	0.000	0.999	0.000	0.995	0.000	0.999
Transition	0.000	0.989	0.000	0.999	0.000	0.959	0.000	0.977
<i>Uncertainty instruments</i>								
Geopolitical	0.007	1.000	0.023	1.000	0.076	0.995	0.115	0.983
Speculation	0.000	0.982	0.000	0.911	0.000	0.133	0.000	0.048
Supply	0.000	0.998	0.000	1.000	0.000	0.995	0.000	0.991
Transition	0.000	0.993	0.000	1.000	0.000	0.982	0.000	0.957

Notes: p -values from the Ljung-Box Q test for serial correlation at lags 6, 12, 18, and 24. "Raw" refers to the unpurged type-specific instruments; "Purged" refers to residuals from regressing each instrument on 12 own lags and 12 lags of the four macro state variables. Bold values indicate rejection at the 5% level. Sample: 1987:02–2022:12.

C.3 Over-Identification Test (with OPEC Instrument)

Table 11 reports the over-identification test when OPEC is included as a fifth instrument. The sentiment system rejects at the 10% level (system $J = 30.20$, $p = 0.067$), with the rejection driven entirely by the real-oil-price equation ($J = 14.31$, $p = 0.006$). The C-test confirms that OPEC is the sole source of this heterogeneity ($C = 13.43$, $p < 0.001$). Despite this diagnostic tension, the structural impulse responses are qualitatively unchanged when OPEC is included (Appendix Figure 20).

Table 11: Over-Identification Tests (OPEC Included)

<i>Panel A: Narrative Sentiment Shock</i>			
Equation	J -statistic	p -value	
Other index	1.526	0.822	
Real oil price	14.314	0.006	**
World oil prod	3.975	0.409	
World IP	6.106	0.191	
CPI	4.279	0.370	
System	30.201	0.067	

<i>Panel B: Narrative Uncertainty Shock</i>			
Equation	J -statistic	p -value	
Other index	3.406	0.492	
Real oil price	5.970	0.201	
World oil prod	3.101	0.541	
World IP	3.019	0.555	
CPI	1.217	0.875	
System	16.714	0.671	

<i>Panel C: Instrument Exclusion (C-test), Sentiment Shock</i>			
Dropped instrument	F (remaining 4)	C -statistic	p -value
Geopolitical	18.5	3.114	0.078
OPEC	16.7	13.428	0.000 **
Speculation	21.3	3.003	0.083
Supply	24.7	0.001	0.981
Transition	18.1	3.225	0.073

Notes: This table reports the 5-instrument specification (OPEC included) as a robustness check. See Table 6 for the headline 4-instrument results. Panels A and B report Sargan J -statistics from 2SLS regressions of each reduced-form VAR residual $u_{j,t}$ on the 5 purged type-specific instruments, using the fitted value $\hat{u}_{1,t}$ from the first-stage regression. Under H_0 (all instruments valid), $J = T \times R^2 \sim \chi^2(4)$, where 4 is the number of over-identifying restrictions ($k - 1 = 5 - 1$). The system statistic sums the equation-level J -statistics; under H_0 it is distributed $\chi^2(20)$. Panel C reports the C -statistic (difference-in- J) for the sentiment system, testing the marginal validity of each instrument conditional on the remaining 4 being valid; $C \sim \chi^2(1)$ under H_0 . * $p < 0.05$; ** $p < 0.01$.

D Incremental Identification

To verify that the narrative shocks contain variation distinct from existing structural shock measures, Table 12 reports Pearson correlations between the identified narrative sentiment and uncertainty shock series and twelve structural shocks from the oil, gas, and monetary policy literatures.

Table 12: Correlation of Narrative Shocks with Existing Structural Shock Measures

Shock	Sentiment shock	Uncertainty shock	<i>n</i>
	ρ	ρ	
Kilian (2008) oil supply shock	-0.178	-0.219*	68
Kilian (2009) oil supply shock	-0.082	-0.129**	245
Kilian (2009) global demand shock	+0.085	+0.001	245
Kilian (2009) oil-specific demand shock	+0.115*	+0.072	245
Bachmann et al. (2022) oil price shock	+0.002	+0.046	425
Alessandri and Gazzani (2025) gas supply shock	+0.039	+0.045	263
Romer and Romer (2004) monetary policy shock	+0.020	+0.031	245
Aruoba and Drechsel (2024) monetary policy shock	+0.036	+0.045	255
Gazzani et al. (2024) oil demand shock	+0.172**	-0.069	187
Gazzani et al. (2024) oil supply/geopolitical risk shock	+0.191***	+0.060	187
Ramey (2011) news shock	+0.103	+0.180*	105
Känzig (2021) OPEC supply news shock	+0.322***	+0.112**	425

Notes: Pearson correlations between the identified narrative sentiment and uncertainty structural shock series and existing structural shock measures from the literature. Small correlations indicate that the narrative shocks contain variation largely orthogonal to existing identified shocks, supporting the claim that the two-channel framework identifies novel structural variation. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

The narrative *uncertainty* shock is essentially orthogonal to every existing measure. All correlations are below 0.22 in absolute value, and only three reach marginal significance: the Kilian (2008) quarterly oil supply shock ($\rho = -0.22$, $n = 68$), the Kilian (2009) monthly oil supply shock ($\rho = -0.13$), and the Känzig (2021) OPEC supply news shock ($\rho = +0.11$). The negative correlations with oil supply shocks are economically interpretable: realized supply disruptions tend to coincide with elevated narrative uncertainty. But the magnitudes are small, with less than 5% shared variance in each case. The uncertainty shock is uncorrelated with monetary policy shocks ($|\rho| < 0.05$), the Bachmann et al. (2022) oil price shock ($\rho = +0.05$), and the Alessandri and Gazzani (2025) gas supply shock ($\rho = +0.05$). This confirms that the supply-side real options channel identified in Section 5.2 captures structural variation that existing oil-market frameworks do not isolate.

The narrative *sentiment* shock is also largely orthogonal to existing measures, with most correlations below 0.12 in absolute value. Three correlations are larger and economically interpretable. The strongest is with the Känzig (2021) OPEC supply news shock ($\rho = +0.32$), reflecting partial overlap between narrative-driven demand ex-

pectations and OPEC-announcement-driven supply news. Even here, the shared variance is roughly 10%, leaving 90% of the narrative sentiment shock as variation not captured by Känzig’s identification. The sentiment shock also correlates with the [Gazzani et al. \(2024\)](#) oil demand shock ($\rho = +0.17$) and oil supply/geopolitical risk shock ($\rho = +0.19$), consistent with the fact that bullish oil-market narratives respond to both demand-side developments and geopolitical supply risks. The sentiment shock is essentially uncorrelated with the [Bachmann et al. \(2022\)](#) oil price shock ($\rho = +0.00$), both monetary policy shocks ($|\rho| < 0.04$), and the [Alessandri and Gazzani \(2025\)](#) gas supply shock ($\rho = +0.04$), confirming that the narrative sentiment channel is distinct from these established identification strategies.

Taken together, these correlations confirm that the narrative shocks capture novel structural variation not summarized by existing shock decompositions. The uncertainty shock is nearly fully orthogonal to all twelve comparison measures, and the sentiment shock retains substantial independent variation even after accounting for modest and economically interpretable overlaps with the [Känzig \(2021\)](#) OPEC supply news shock and the [Gazzani et al. \(2024\)](#) oil demand and geopolitical risk shocks.

E Additional Robustness Checks

This appendix reports robustness checks for the baseline Proxy-SVAR results. In each case, the qualitative two-channel pattern is preserved: the sentiment shock raises oil prices and activity, while the uncertainty shock depresses production and activity with a muted price response.

Figure 16 compares the baseline Proxy-SVAR impulse responses with local projection estimates using the same external instruments (LP-IV) following [Stock and Watson \(2018\)](#). The LP-IV point estimates track the Proxy-SVAR closely for both shocks, with wider confidence bands at longer horizons. The agreement confirms that the structural results are not artifacts of the VAR dynamics.

Figure 17 re-estimates both systems with $p = 3$ and $p = 12$. The impulse responses are qualitatively unchanged, with minor differences in confidence band width at longer horizons.

Figure 18 compares the Proxy-SVAR impulse responses with recursive (Cholesky) identification using the ordering (SENT, UNC, ROP, WOP, WIP, CPI). The Cholesky responses are qualitatively similar for the sentiment shock but less precisely estimated for the uncertainty shock, where the recursive ordering assumption is harder to justify.

Figure 19 re-estimates each system four times, dropping one instrument at a time. All leave-one-out specifications produce point estimates close to the baseline and within the 90% confidence bands. In the sentiment system, joint F -statistics range from 12.9

(dropping transition) to 20.7 (dropping supply). In the uncertainty system, robust F -statistics range from 9.2 (dropping geopolitical) to 18.3 (dropping transition).

Figure 20 re-estimates both systems with OPEC included as a fifth instrument. The impulse responses are qualitatively similar to the headline specification. The sentiment system overid test becomes borderline (system $J = 30.20$, $p = 0.067$), driven by the real-oil-price equation ($J = 14.31$, $p = 0.006$). The uncertainty system passes overid cleanly ($p = 0.671$). See Section 4.3 for the economic justification for excluding OPEC from the headline specification.

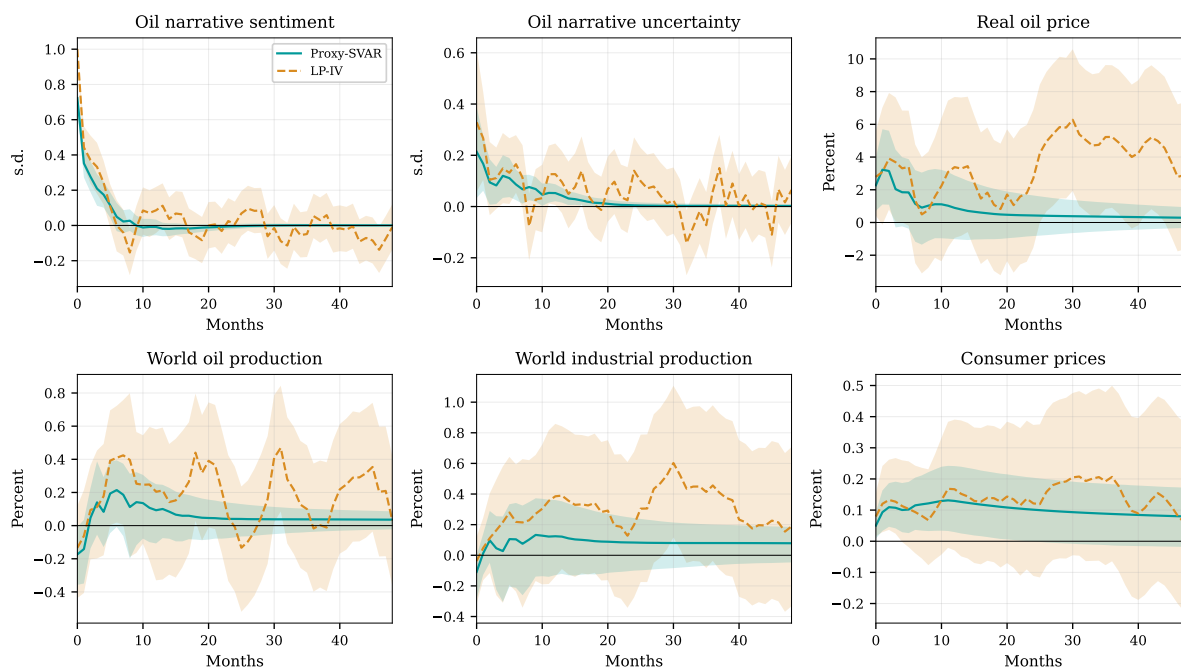
Figure 21 re-estimates both systems after excluding the 18-month COVID-19 window (January 2020 through June 2021). The sentiment shock responses are essentially unchanged. The uncertainty shock preserves the key qualitative pattern: world industrial production contracts persistently. The real oil price response, which was weakly negative in the baseline, becomes weakly positive when COVID is excluded, consistent with the theoretical prediction that the net price effect of an uncertainty shock is ambiguous. First-stage statistics strengthen slightly in the restricted sample (sentiment robust $F = 17.2$; uncertainty robust $F = 14.7$).

Figure 22 re-estimates the uncertainty Proxy-SVAR replacing the baseline narrative uncertainty index with the [Abiad and Qureshi \(2023\)](#) Oil Price Uncertainty index, retaining the same four type-specific instruments. The uncertainty instruments predict the OPU index with $F = 7.4$ (robust $F = 5.3$). The impulse responses reproduce the supply-side contraction pattern: declining world industrial production, a muted CPI response, and qualitatively similar real oil price and production dynamics. The sentiment instruments have no predictive power for the alternative measure ($F = 0.8$), confirming channel specificity.¹²

¹²The alternative sentiment measure uses a generic financial dictionary that classifies geopolitical supply-tightening events as negative rather than bullish, which is inconsistent with the [Loughran et al. \(2019\)](#) oil-specific scoring used in the baseline. See Section 3.3.

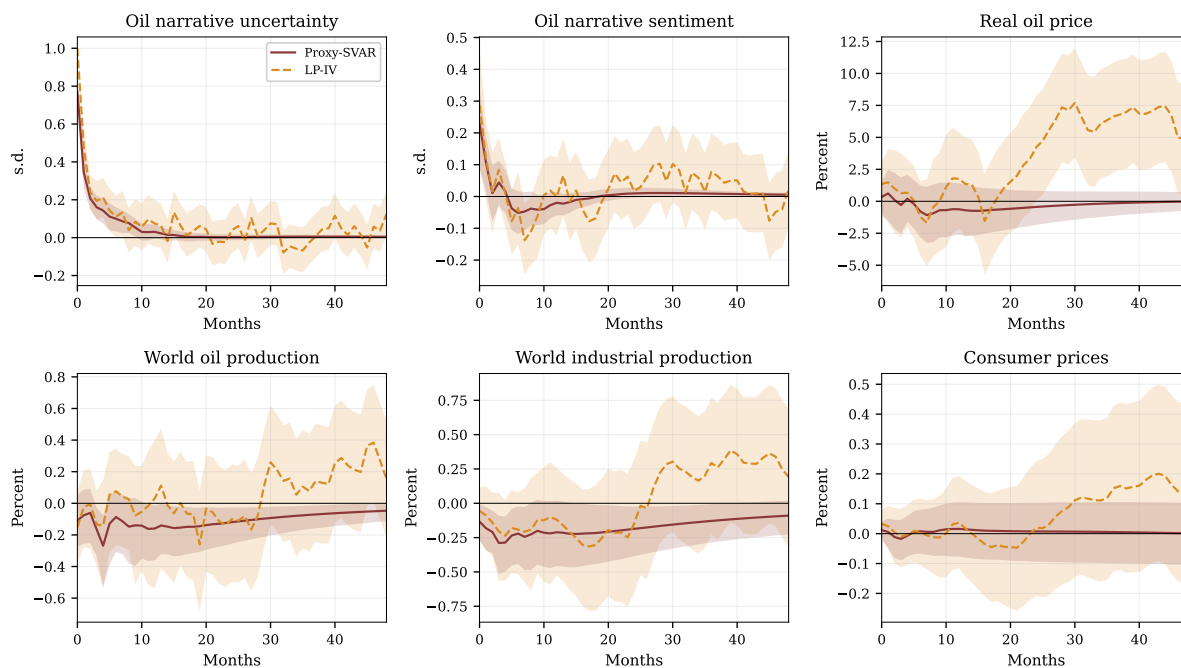
E.1 LP-IV vs. Proxy-SVAR

Sentiment shock: Proxy-SVAR vs LP-IV



(a) Narrative sentiment shock

Uncertainty shock: Proxy-SVAR vs LP-IV

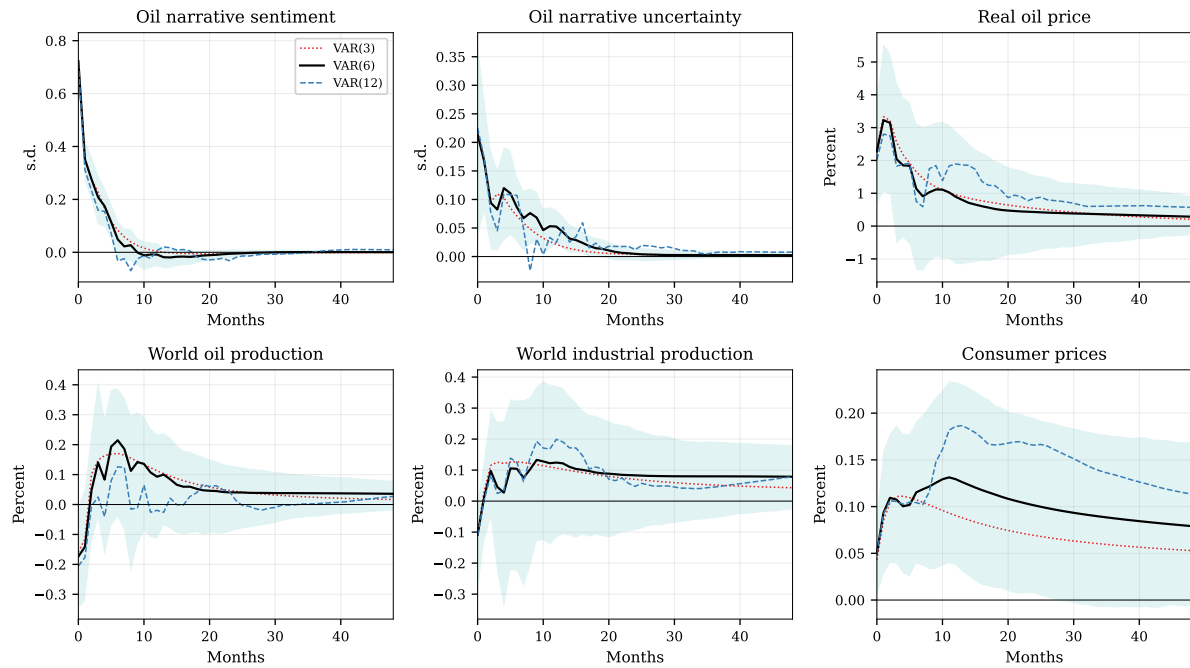


(b) Narrative uncertainty shock

Figure 16: Proxy-SVAR vs. LP-IV impulse responses. Solid line: Proxy-SVAR with 90% MBB bands (shaded). Dashed: LP-IV with 90% MBB bands. Sample: 1987:02–2022:12, VAR(6).

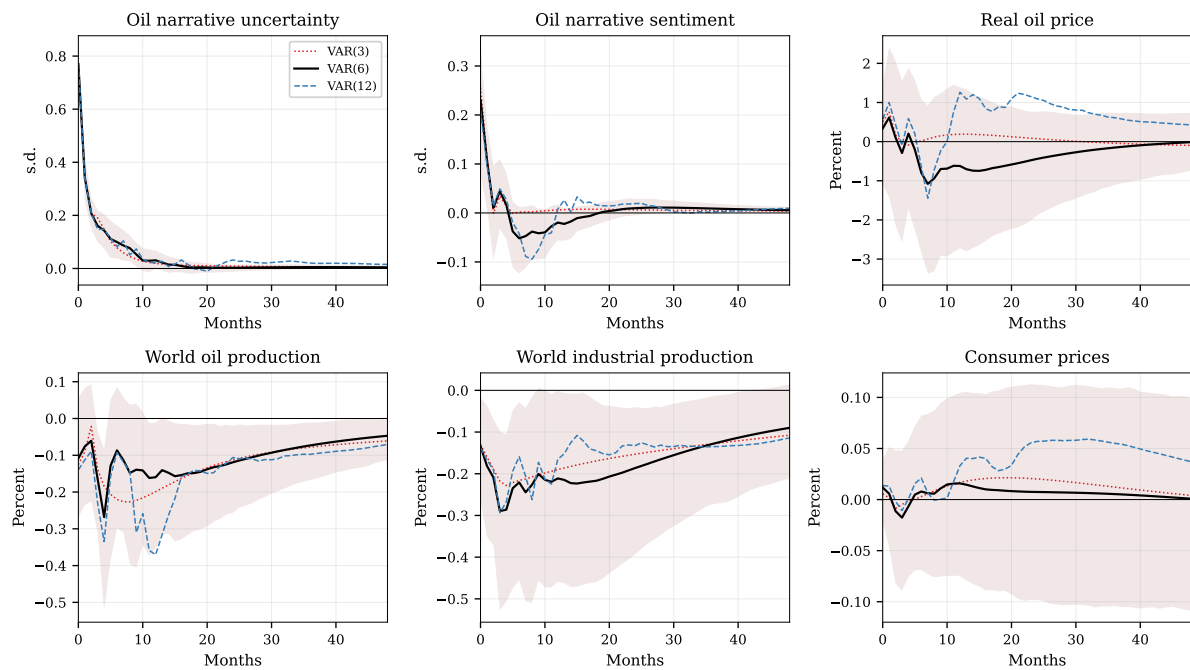
E.2 Alternative Lag Order

Sentiment shock: Alternative lag orders



(a) Narrative sentiment shock

Uncertainty shock: Alternative lag orders

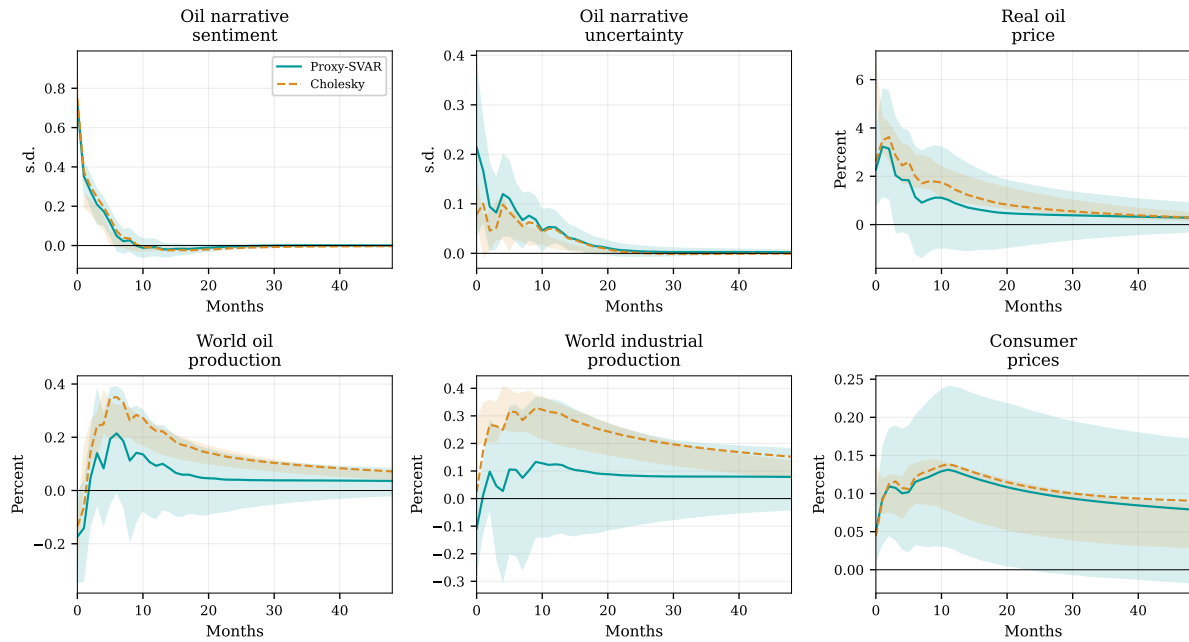


(b) Narrative uncertainty shock

Figure 17: Impulse responses under alternative lag orders. Baseline: VAR(6). Alternatives: VAR(3) and VAR(12). Sample: 1987:02–2022:12.

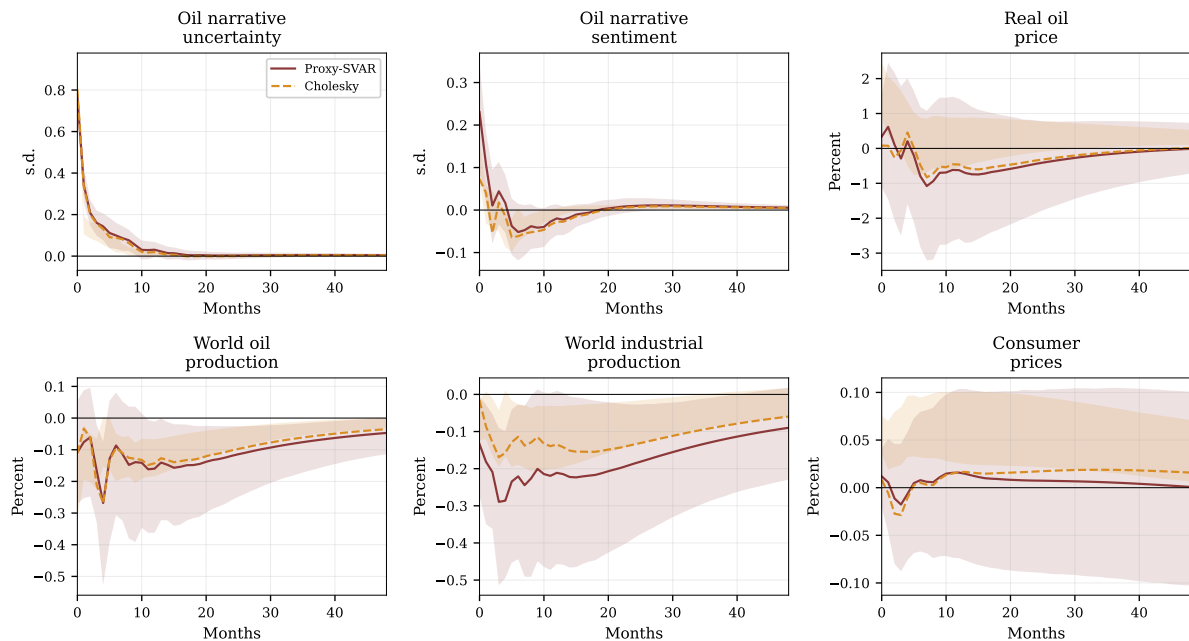
E.3 Cholesky Identification

Sentiment shock: Proxy-SVAR vs Cholesky identification
Ordering: (SENT, UNC, ROP, WOP, WIP, CPI), VAR(6)



(a) Narrative sentiment shock

Uncertainty shock: Proxy-SVAR vs Cholesky identification
Ordering: (UNC, SENT, ROP, WOP, WIP, CPI), VAR(6)

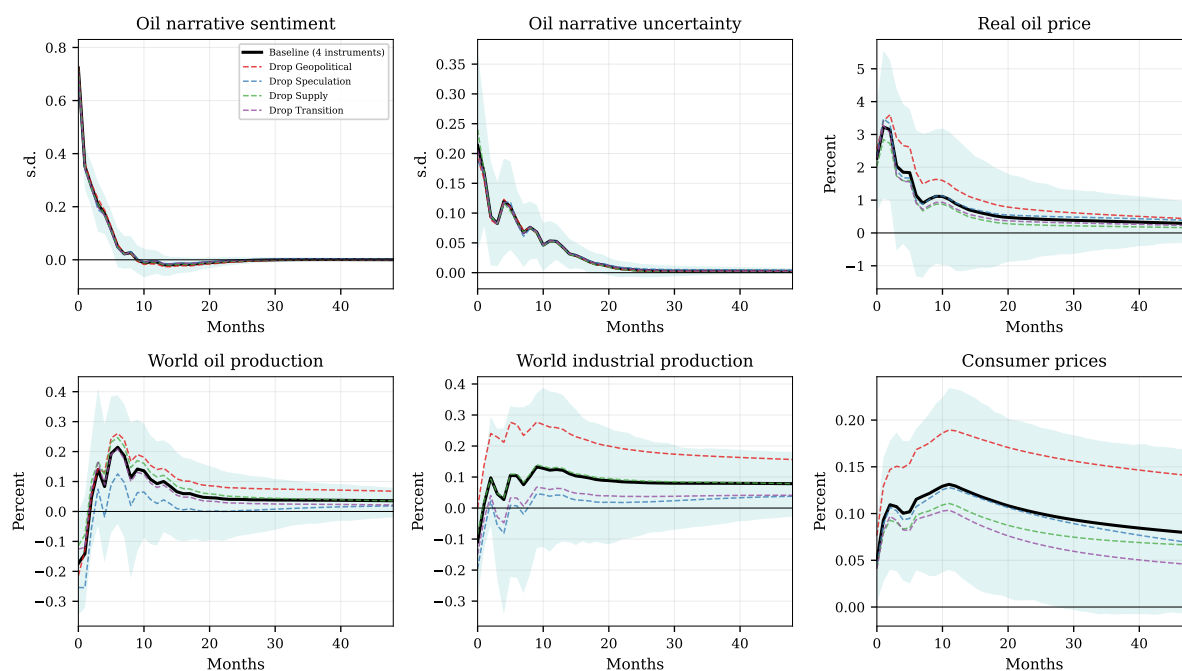


(b) Narrative uncertainty shock

Figure 18: Proxy-SVAR vs. Cholesky impulse responses. Solid: Proxy-SVAR with 68% bands. Dashed: Cholesky. Sample: 1987:02–2022:12, VAR(6).

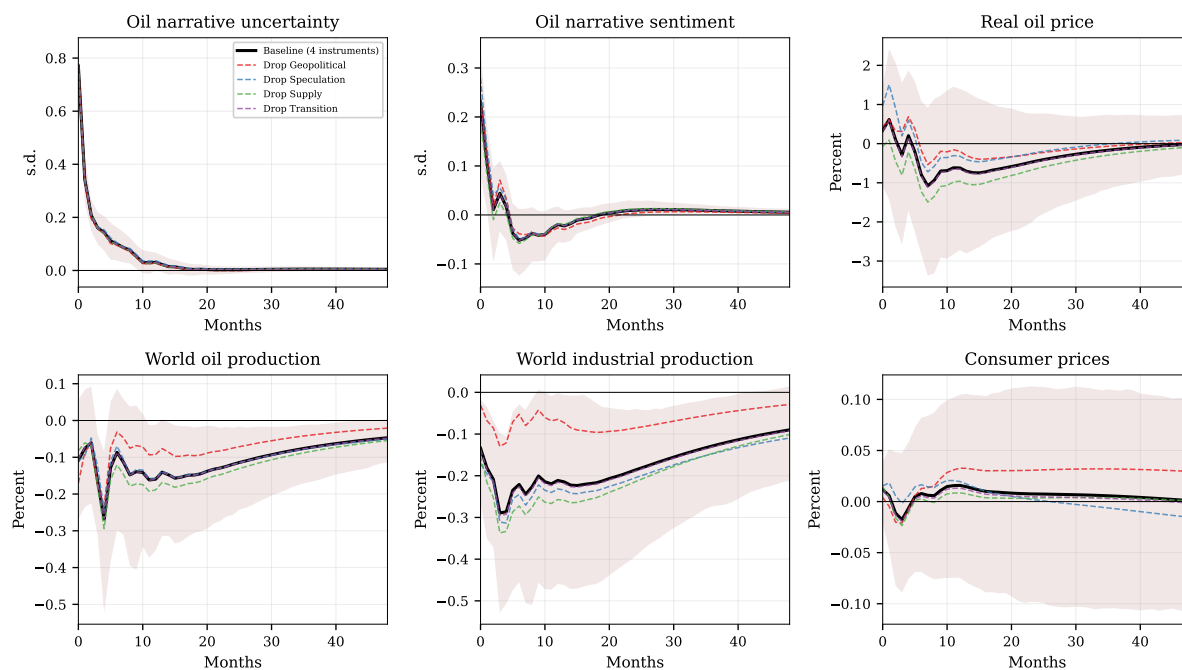
E.4 Leave-One-Out Instrument Analysis

Sentiment shock: Leave-one-out instrument robustness



(a) Narrative sentiment shock

Uncertainty shock: Leave-one-out instrument robustness

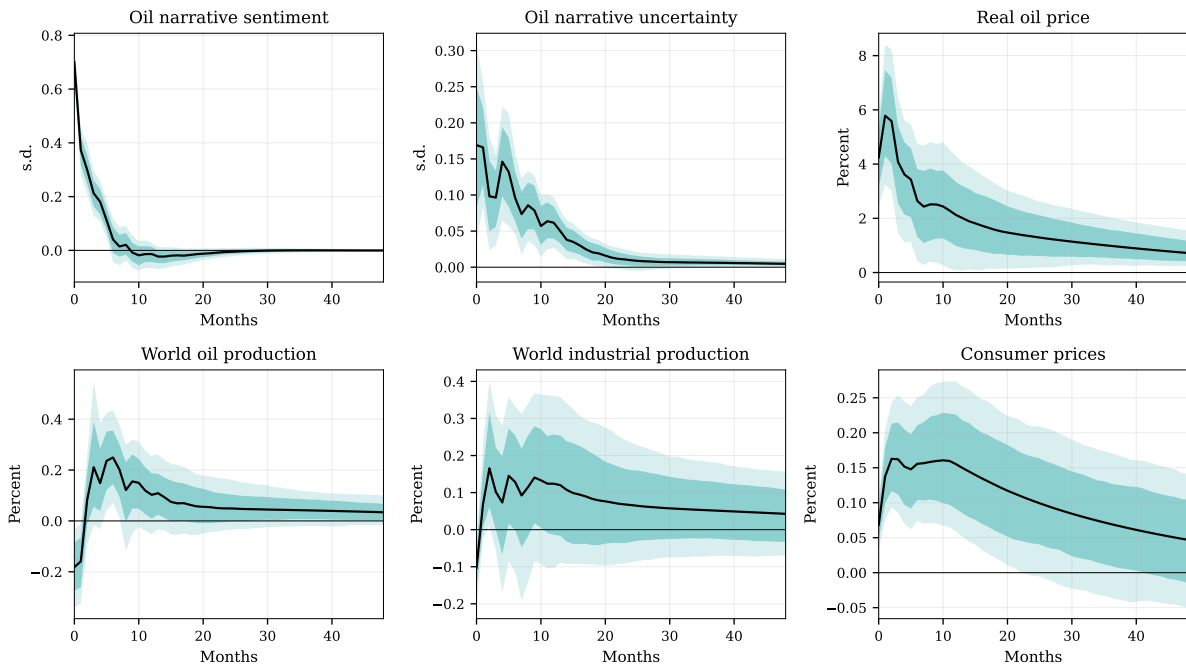


(b) Narrative uncertainty shock

Figure 19: Black solid: baseline with 90% bands. Colored dashed: each leave-one-out specification. Sample: 1987:02–2022:12, VAR(6).

E.5 Including OPEC Instrument

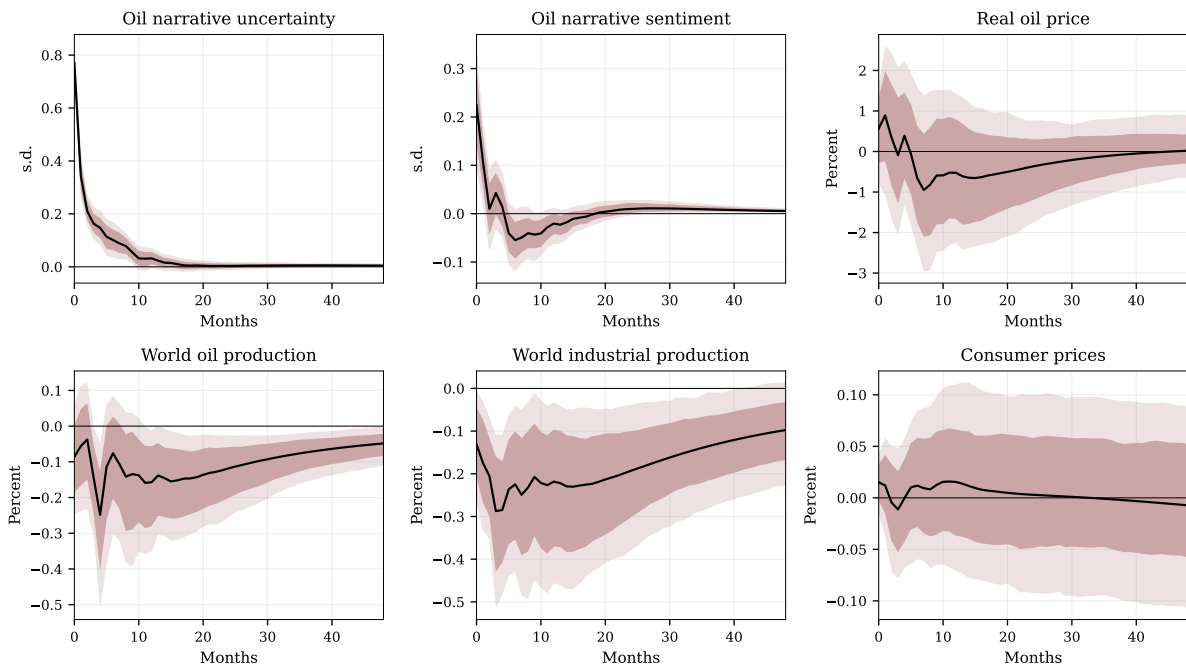
Sentiment shock (5 instruments, OPEC included)



First-stage: $F = 20.3$ (robust 21.4), Eff. $F = 20.5$, $R^2 = 19.5\%$, Adj. $R^2 = 18.5\%$

(a) Narrative sentiment shock

Uncertainty shock (5 instruments, OPEC included)



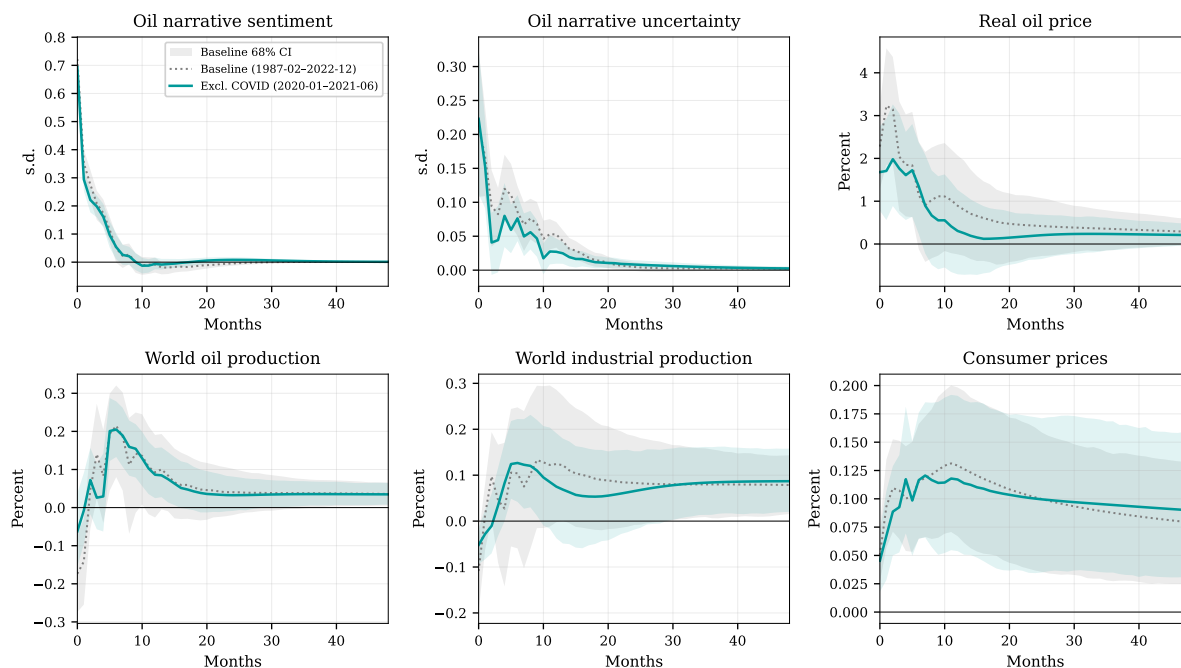
First-stage: $F = 22.2$ (robust 11.7), Eff. $F = 10.4$, $R^2 = 20.9\%$, Adj. $R^2 = 20.0\%$

(b) Narrative uncertainty shock

Figure 20: Impulse responses with OPEC included (5 instruments). Sample: 1987:02–2022:12, VAR(6).

E.6 Excluding COVID-19

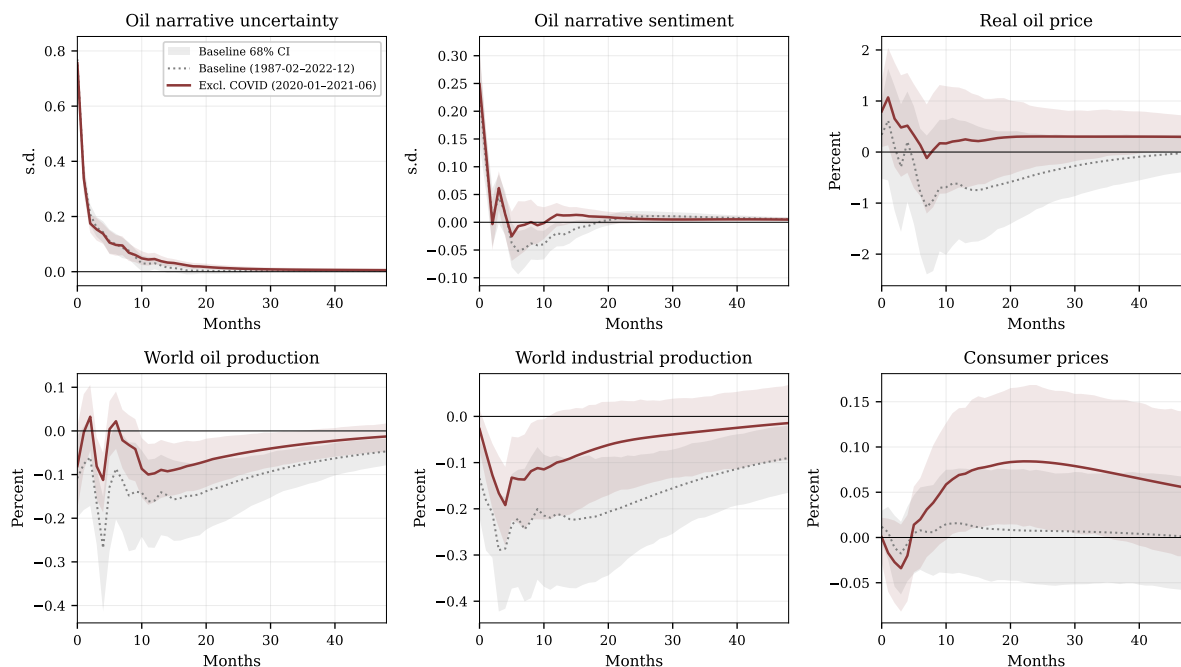
Sentiment shock: Baseline vs excluding COVID-19 (2020-01-2021-06)



Robust F: Baseline=16.8, Excl. COVID=17.2

(a) Narrative sentiment shock

Uncertainty shock: Baseline vs excluding COVID-19 (2020-01-2021-06)



Robust F: Baseline=13.9, Excl. COVID=14.7

(b) Narrative uncertainty shock

Figure 21: Baseline vs. excluding COVID-19 (2020:01–2021:06). Sample: 1987:02–2022:12 (excluding 18 months).

E.7 Alternative Aggregate Uncertainty Measure

Uncertainty shock: Baseline vs alternative aggregate measure

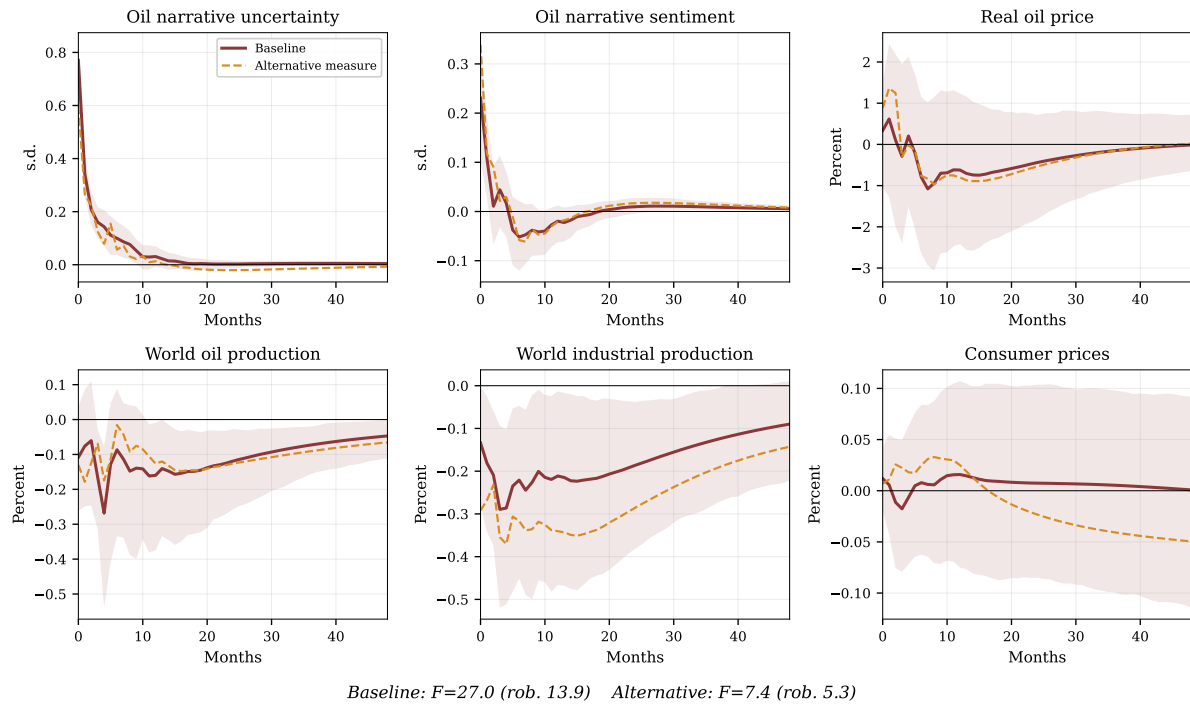


Figure 22: Uncertainty shock: baseline narrative uncertainty index vs. [Abiad and Qureshi \(2023\)](#) Oil Price Uncertainty index. Same four type-specific instruments. Sample: 1987:02–2022:12, VAR(6).