UNPACKING COMMODITY PRICE FLUCTUATIONS: READING THE NEWS TO UNDERSTAND INFLATION

DIMITRIS MALLIAROPULOS, EVGENIA PASSARI, FILIPPOS PETROULAKIS

ABSTRACT. We show that text-based indicators of supply and demand disturbances in commodity markets provide distinct information about future inflation movements, relative to existing predictors and various measures of inflation expectations derived from financial prices or from surveys of professional forecasters and consumers. Specifically, we document that demand disturbances play a significantly larger role in prediction than supply disturbances because they typically lead to more correlated increases in quantities and prices of goods across the consumer basket, resulting in a clear and positive relationship between commodity prices and overall inflation. Local projections confirm that demand disturbances generate substantially larger inflation responses in the short- to medium-term. In terms of magnitudes, commodity-specific indicators reduce out-of-sample inflation forecast errors at the one-year horizon by up to 30 percent.

JEL classification codes: C19, E31, E37, F40, Q02.

Keywords: inflation, commodities, news, textual analysis, demand, supply.

D. Malliaropulos: University of Piraeus and Bank of Greece, 21 El. Venizelou, 10250 Athens, Greece (e-mail: dmalliaropulos@bankofgreece.gr). E. Passari: Université Paris-Dauphine, PSL and CEPR. DRM Finance, Université Paris Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris, France (e-mail: evgenia.passari@dauphine.psl.eu). F. Petroulakis: Bank of Greece, 21 El. Venizelou, 10250 Athens, Greece (e-mail: fpetroulakis@bankofgreece.gr). The authors thank George-Marios Angeletos, Marta Banbura, Gadi Barlevy, Francesco Bianchi, Ryan Charhour, Ludmila Fadejeva, Raffaella Giacomini, Yuriy Gorodnichenko, François Gourio, Refet Gürkaynak, Matteo Iacoviello, Michele Lenza, Andrei Levchenko, Francesca Monti, Dimitris Papanikolaou, Pascal Paul, Giorgio Primiceri, Ricardo Reis, Hélène Rey, Alexander Richter, and the participants of the Sciences Po - OFCE - Banque de France - CEPR Workshop on Empirical Monetary Economics 2023, the 2024 AEA session "New Developments in the Study of Inflation Origins and Inflation Dynamics", the 31st CEPR European Summer Symposium in International Macroeconomics (ESSIM), the PSE Macro Days 2024, the ECB "New Methods for Sentiment Analysis with Artificial Intelligence in Central Banks" 2025 workshop, the 2nd LBS Alumni Conference, and the EC-CEPR-JIE Conference on "Global Shocks, Macroeconomic Spillovers and Geopolitical Risks" for comments and discussions. Evgenia Passari gratefully acknowledges financial support from the Bank of Greece, the Europlace Institute of Finance (EIF) & the Labex Louis Bachelier and the FinTech & Digital Finance Chair at Université Paris Dauphine - PSL, as well as hospitality from the Kellogg School of Management, Northwestern University, the Federal Reserve Bank of Chicago and the Kiel Institute of the World Economy where this research was partly conducted. The views expressed in this paper are those of the authors and should under no circumstances be interpreted as reflecting those of the Bank of Greece or the European System of Central Banks.

1. Introduction

Monetary authorities recognize that fluctuations in commodity prices are significant determinants of inflation dynamics (Bernanke, 2008). Commodity prices serve as leading indicators of inflation since they are sensitive to changing economic conditions and can influence aggregate prices through distribution channels and second-round effects. Understanding the intricate relationship between commodity price changes and inflation is vital for effective monetary policy, especially in the face of evolving global challenges that can have large effects on commodity markets. How exactly commodity prices affect inflation, however, likely depends on the drivers of commodity price fluctuations. Such fluctuations can arise from demand shifts associated with economic booms and recessions, hedging operations and speculative trading behaviors, supply constraints precipitating precautionary demand pressures, or geopolitical events, including trade agreements and wars. Each of these factors exerts heterogeneous influences on aggregate price indices and macroeconomic outcomes, both in terms of magnitudes and persistence. Our aim is to disentangle these forces using text-based indicators of demand and supply disturbances in commodity markets.[1]

We show that commodity prices provide distinct information about future inflation movements relative to existing predictors, inflation expectations and survey forecasts. Importantly, however, this predictive power is obscured by the fact that commodity prices are driven by both supply and demand shocks, and these shocks have different passthrough to inflation. Our starting point is that under the presence of aggregate demand shocks, which typically affect most sectors of the economy in a homogeneous manner, the relationship between commodity prices and inflation tends to be direct and unambiguous, reflecting substantial pass-through effects. In line with this intuition, we find that positive aggregate demand shocks are generally associated with widespread increases in both prices and quantities across the consumer basket, thus establishing

[1]Several studies have documented that specific commodities (for instance, oil) are informative about inflation, but this relation is not stable across different countries and time periods. Commodity prices were effective predictors of inflation until the early 1980s, but their predictive power has diminished since the mid-1980s (Stock and Watson, 2003). The predictive power of commodity prices for aggregate price indices is weak, despite the widespread view that commodity futures capture substantial information regarding underlying supply and demand conditions aggregated by futures markets (Bernanke, 2008). An additional complication, recently emphasized by Shapiro (2020), arises from the aggregation of categorical inflation measures at a broad-index level which can obscure inflation dynamics due to the varying sensitivities of different sectors to distinct underlying factors.

a stronger link to overall inflation dynamics.[2] Conversely, supply shocks that are idiosyncratic to specific commodities can generate varying effects on the general price level, depending on the interplay between substitute and complementary goods. For instance, a negative supply shock in the cocoa market may elevate cocoa prices while exerting downward pressure on the price of sugar, given their complementary nature in consumption. Meanwhile, commodities such as coffee, which are imperfect substitutes for cocoa, may see limited demand effects, thereby remaining relatively insulated from price changes. Consequently, although a commodity-specific supply shock might elevate the price of the affected good, its broader impact on inflation may be attenuated due to weak pass-through to the overall consumption basket, a prediction that is consistent with our findings.

Identifying supply and demand shocks across different commodity markets entails significant challenges. In theory, techniques that employ traditional data and econometric methods that have been used extensively for the analysis of the oil market could also be adopted to separately disentangle supply from demand for different commodities. In practice, these more traditional methods have been reserved, until recently, almost exclusively for the study of the oil market. This has been the result of data availability, intuitive modeling assumptions, conventional sign-restrictions choices, as well as a heightened interest in the oil market versus the markets of other commodities. Although some more recent papers, including Güntner, Reif, and Wolters (2024) and Winne and Peersman (2016), provide supply and demand shocks using structural techniques for the natural gas and food markets, respectively, we are the first to apply a unified framework for the study of a large number of commodities, spanning all sectors, to the study of inflation.

We employ the automated, narrative approach of Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025), that combines textual analysis with human auditing to construct disaggregated supply- and demand-side disturbances for a wide array of commodities from business news. This framework has a number of advantages. The most important one is the capacity to assimilate extensive informational sets across different markets. Additionally, the textual methodology necessitates fewer a priori parametric restrictions and circumvents a number of identification challenges inherent in the structural form estimation of understudied commodity markets, such as metals and

[2]More broadly, shocks to aggregate demand tend to predict future inflation more accurately compared to idiosyncratic supply shocks, particularly when the latter are not accommodated by monetary authorities (Boughton and Branson, 1991).

livestock. Sourcing supply and demand disturbances from a rich panel of commodities is crucial for the mechanism of the paper, since it allows us to isolate common from idiosyncratic components in commodity price movements; the common component carries important information for the prediction of aggregate inflation. Moreover, this framework helps us distinguish between different drivers of commodity price developments. For instance, commodity price fluctuations might reflect changes in global demand that are common across commodities, commodity-specific demand developments, or supply-related events such natural disasters and wars; each of these may have differential effects on inflation.

In addition to these composite indicators which span the full spectrum of traded commodities, we also use text-based measures of key risk factors that affect commodity markets, such as recessionary regimes, natural disasters, or geopolitical risk. Employing the narrative-based, commodity supply and demand indicators allows us to cut through the sources of confusion that arise when one examines the informational content of commodity price developments alone. We then proceed to explore whether these measures that separately identify supply- and demand-side disturbances across commodity markets are able to provide additional information regarding the future trajectory of inflation over and above standard predictors employed in the literature.

Using these indicators, we report several new findings. First, narrative-sourced supply and demand disturbances help predict future inflation, over and beyond common predictors such as past inflation, the policy rate, the unemployment gap, yield curve measures, and the log growth of a trade-weighted dollar index. For example, textual indicators help reduce the out-of-sample forecast error of inflation by 20% to 30%, depending on the horizon and the exact specification. The fact that textual indicators provide incremental information on top of commodity returns is particularly encouraging and illustrates the benefit of separately examining the supply and demand drivers of commodity prices, given that their importance and persistence may vary. For instance, we find that, on average, demand-side disturbances contribute more to the out-of-sample predictability of different inflation baskets. A notable benefit of our indices is that they are news-based, and so are available at a timely fashion, unlike models that use production data, which are published with delays.

We further assess the validity of the textual commodity supply and demand indicators by looking at their performance across different inflation regimes. We follow the approach of Joseph et al. (2021) and split the sample into periods when inflation is increasing, falling or remains stable. Our framework provides clear improvements

over the baseline model for headline inflation at the one-year horizon, especially for episodes of falling and stable headline inflation. For episodes of increasing inflation, we show that the narrative measures provide substantial and consistent improvements across several specifications. For shorter horizons (six months), the addition of text-based commodity measures again improves the performance of specifications that incorporate additional predictors such as tradable commodity returns.

After establishing that textual measures for oil compare favorably with the structurally identified oil shocks of Baumeister and Hamilton (2019) and Känzig (2021),[3] we examine the effects of our proxies for different inflation baskets using local projections (Jordà, 2005). These local projections confirm the quantitatively more important role of demand relative to supply disturbances - on average. Supply disturbances matter only for the recent period that includes the COVID-19 pandemic. We further show that our indices contain valuable information for future inflation that is not reflected in various measures of inflation expectations, such as measures derived from financial market prices and survey-based measures of professional forecasters and consumers. Hence, our indices are not merely picking up already-available information but insights that are not apparent to either market participants, professional forecasters, or households.

What drives the ability of our text-based indicators to forecast inflation relative to existing models? We argue that the forecasting improvement over the benchmark autoregressive specification, which is notoriously hard to beat (Faust and Wright, 2013; Bańbura, Lenza, and Paredes, 2024), comes from disentangling supply from demand dynamics in commodity prices. This decomposition is possible due to the use of textual data on news articles and, most importantly, the classification into supply- or demand-driven news. In particular, the source of the price movement matters, not just the size: we show that isolating supply- from demand-driven price developments is crucial for the improved performance in inflation predictability even when both supply and demand disturbances result in commodity price movements of the same direction. This argument is confirmed by auxiliary forecasts using real-time data, which yield very similar results. If forecasting performance were worse with real-time versus revised inflation measurements, then the timeliness of our indices, rather than the textual information allowing the disentangling of supply and demand, would explain the forecasting gains.

[3]The impulse responses of the main variables to our measure of oil supply are reasonably close and qualitatively similar to the impulse responses to the shocks of Baumeister and Hamilton (2019) and Känzig (2021) (either to both or to one of them).

A further reason for the performance of our composite indicators is the implied weighting scheme. In particular, the construction of the composite indices that implicitly uses news-weights for individual commodities is informative as it reflects patterns in news coverage. This is because the individual commodity weights are time-varying and driven by the news cycle. Media coverage patterns provide valuable information about which commodity movements matter most for broader inflation dynamics. An alternative scheme where the construction of the indices uses production weights leads to slightly worse predictions. This performance differential, though non-trivial, is not major, nor statistically significant, indicating that the primary benefit of our methodology comes from separately identifying supply and demand disturbances, as these propagate differently through the consumer price basket. In addition, incorporating the ex-ante defined commodity risk factors that (Lumbanraja et al., 2025) source from news articles seems to lead to some improvement in forecasting performance, especially during the COVID-19 shock. [4]

Our contribution to the literature is threefold. Our primary contribution is the application of textually-derived supply and demand measures that span the broad commodity market for inflation prediction on the backdrop of the disconnect between commodity prices and inflation documented by the literature.[5] The universality of business news, along with their high frequency and extensive coverage of the highly heterogeneous commodity markets, make the employed textual measures a useful tool for the analysis of the transmission of global shocks and their implications for inflation. This is particularly important in the wake of the recent policy challenges following the COVID-19 pandemic and the global geopolitical tensions which have generated a lot of interest and some disagreement regarding the relative role of supply and demand factors for inflation (Blanchard and Bernanke, 2023; Bańbura, Bobeica, and Martínez Hernández, 2023; Giannone and Primiceri, 2024; Bergholt and Ulvedal, 2024). In this context, the use of natural language processing methods can deepen our understanding of inflation behavior. Our forecasting exercise demonstrates that the textual supply and demand indices yield considerable forecasting improvements compared to conventional autoregressive benchmarks. Although we do not contribute to the debate about inflation decomposition in the recent episodes directly, the real-time

[4] The improvement arises from the inclusion of the natural disaster factor in particular.

[5] See the works of Hooker (2002), Stock and Watson (2003). A notable exception is the work of Gospodinov and Ng (2013a), which demonstrates that the principal components of convenience yields retain predictive power for aggregate price indices and commodity prices.

applicability of the employed framework allows for its integration into existing forecasting methodologies employed by monetary authorities.

Second, our work contributes to the growing literature that uses textual analysis for the study of various economic and financial outcomes. Recent contributions include the works of Tetlock (2007), Gentzkow and Shapiro (2010), Hoberg and Phillips (2010), Boudoukh et al. (2013), Alexopoulos and Cohen (2015), Baker, Bloom, and Davis (2016), Allcott and Gentzkow (2017), Hassan et al. (2019), Angelico et al. (2022), Hassan et al. (2023a), and Hassan et al. (2023b). The study of Angelico et al. (2022)employs Italian tweets to construct a real-time measure of consumers' inflation expectations. The authors show that Twitter-based indicators, built using Latent Dirichlet Allocation (LDA) models, combined with a dictionary-based approach, are highly correlated with monthly survey-based and daily market-based inflation expectations. Relative to this body of work, our focus is on inflation prediction. Another paper that leverages the use of computer linguistics for inflation prediction is the study of Hong, Jiang, Meng, and Xue (2024) that breaks down economic stories from over 880,000 Wall Street Journal articles into different topics, using LDA to forecast inflation.

Unlike the works of Angelico et al. (2022) and Hong et al. (2024), we do not systematically exploit narratives from the news corpus on top of the sourcing of supply and demand disturbances.[6] For this reason, we employ a textual measure that exploits information in a supervised and semi-supervised textual framework which, differently to LDA methods, allows us to define the source of commodity-price fluctuations ex-ante, consistent with our proposed economic mechanism: forecast gains mainly arise from disentangling supply from demand dynamics. The idea is that demand disturbances tend to lead to correlated changes in quantities and prices of commodities across the consumer basket, resulting in a clear and positive relationship between commodity prices and overall inflation. By contrast, commodity supply disturbances tend to be more idiosyncratic, which, together with substitution patterns, tends to lead to weaker aggregate effects.

Third, our work contributes to the extensive literature that studies the impact of oil supply and demand shocks on inflation (and other macroeconomic variables). Several of these works follow structural approaches (Kilian, 2008; Kilian, Rebucci, and Spatafora, 2009; Kilian, 2009; Kilian and Vigfusson, 2017; Känzig, 2021). Additionally,

[6] In the robustness section we augment the forecasting model with the predefined risk factors built by Lumbanraja et al. (2025). Although the addition of the Natural Disaster factor leads to some improvement in forecasting performance, especially during the COVID-19 shock, predictability throughout the analysis is largely driven from disentangling supply from demand forces.

recent papers that construct narrative measures of shock identification in the spirit of Romer and Romer (2010), include the works of Wu and Cavallo (2012), Caldara, Cavallo, and Iacoviello (2019), Loughran, McDonald, and Pragidis (2019) and the study of Datta and Dias (2019). Our approach differs in that we employ proxies of supply and demand shocks for the universe of tradable commodities from Lumbanraja et al. (2025), instead of focusing on the oil market. This is central to the economic mechanism of the paper because it allows us to account for correlated changes in quantities and prices across commodity markets via the inclusion of the composite demand indicator in the forecasting model.

The remainder of the paper is organized as follows. Section 2 introduces our methodology and the text-based commodity indices. Section 3 highlights the distinct properties of supply and demand indices and sets the ground for the empirical analysis. Section 4 presents a forecasting exercise in which we study the differential impact of the proposed text-based commodity indices on a number of broad and disaggregated inflation measures, across different sample periods. Section 5 compares the automated, narrative supply and demand measures to shocks affecting oil prices identified by the literature using structural vector autoregressions (SVAR). Section 6, studies the responses of prices to commodity-market events picked by the text-based measures. Section 7 concludes.

2. Textual Analysis and Narrative Identification of Commodity Price Developments

A Text-based Framework for Disentangling Commodity Supply and Demand. We first present the general framework of Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025) that uses a narrative approach for the construction of supply and demand developments in commodity markets. Commodity-related information is collected from more than a million articles from Reuters between January 2000 and June 2023. Indices of supply and demand are constructed first at the market-wide level, which summarizes the dynamics across the universe of commodities, and subsequently the analysis is tailored to span a large number of commodity categories, including energy, industrial and precious metals, agricultural commodities, and livestock. Ultimately, individual commodities such as oil, natural gas, copper, iron, gold, wheat, corn, sugar, and hogs are also individually analyzed and mapped.

For the construction of the indices, news reading is simulated in three steps. The first step is content analysis. Words and word combinations are identified within sentences that can be attributed to supply and demand factors. During the second step of refinement, the process is overlayed with additional algorithms that cater to negations, as well as to a number of exceptions that vary across commodity categories. The third step involves an extensive human auditing exercise.

A crucial step in content analysis is the use of supervised and semi-supervised methods for learning; the algorithm identifies words and word combinations that can be attributed to supply and demand factors. The starting point is the search for the most popular words over the universe of business news articles. The words that appear more frequently are identified and classified into supply and demand lists. The process is complemented with a large number of human checks. This forms the basis of the creation of the supply and demand dictionaries. Additionally, standard dictionaries of "increase" and "decrease" words that are typically employed in the literature that uses textual analysis for the measurement of economic outcomes are employed. The "supply" and "demand" words are then combined with the "increase" or "decrease" words to form pairs that we count as signals. The following equations detail the construction of suppply increase, supply decrease, demand increase and demand decrease indicators.

$$SI_t = \frac{1}{N_t}\sum_{w}^{W_t}\{1[w = Supply] \times 1[|w - pos_{Increase}| \leq i]\} \tag{1}$$

$$SD_t = \frac{1}{N_t}\sum_{w}^{W_t}\{1[w = Supply] \times 1[|w - pos_{Decrease}| \leq i]\} \tag{2}$$

$$DI_t = \frac{1}{N_t}\sum_{w}^{W_t}\{1[w = Demand] \times 1[|w - pos_{Increase}| \leq i]\} \tag{3}$$

$$DD_t = \frac{1}{N_t}\sum_{w}^{W_t}\{1[w = Demand] \times 1[|w - pos_{Decrease}| \leq i]\} \tag{4}$$

where $w = 0, 1, ..., W_t$ are the words contained in articles, with N_t the number of articles published at date t, and $pos_{Decrease}$ ($pos_{Increase}$) the position of the nearest synonym of decrease (increase). To measure changes in demand or supply, we thus count only mentions of supply or demand that occur within a i-word window of a synonym of increase or decrease. We set $i = 4$ but the validity of the indices is robust to the choice

of different word-window sizes.

Net supply is defined as the difference between supply decrease and supply increase and net demand is defined as the difference between demand increase and demand decrease. Consequently, an increase in net supply and an increase in net demand should both have a positive effect on inflation. To quantify the net effects, we thus generate:

$$NetDemand_t = DI_t - DD_t \tag{5}$$

$$NetSupply_t = SD_t - SI_t \tag{6}$$

A standardized version of the indices divides the net supply and demand indicators with the total number of articles the algorithm "reads" per day, ensuring that they are not artificially inflated by changes in the publication policy of news articles.

The Drivers of Supply and Demand Developments. The origins of supply and demand developments can have different implications for inflation. For instance, natural disasters are thought to have an impact mostly on short- to medium-term inflation while environmental regulation could have a lasting effect on inflation by permanently affecting relative prices through expectations. This is highly relevant for policy because it implies that central banks can look through transitory shocks, but may have to respond to persistent shocks. Lumbanraja et al. (2025) retrieve four main drivers of commodity prices according to information collected from IMF commodity market reports; namely, a business cycle /recessionary factor, a geopolitical risk factor, a natural disasters factor, and a climate change factor. The commodity drivers' dictionaries are derived using a semi-supervised approach, employing a rich list of authoritative texts for each driver category, according to the approach of Engle et al. (2020). In particular, the authoritative texts for the business cycle driver are composed of the full history of the Business Cycle Dating Committee Announcements of the NBER (1979-2020).

The employed narrative commodity indices help shed light on the complicated relationship between commodity prices and inflation dynamics by distinguishing among aggregate and commodity-specific demand and supply developments. The proposed framework complements concurrent structural measures for the identification of supply and demand developments that focus on the oil market because it spans most individual tradable commodities. This allows us to study the differential impact of supply and demand developments for different inflation measures and assess their suitability.

Measures of Supply and Demand Developments for a Broad Commodity Index. For the construction of supply and demand indicators for the global commodity index Lumbanraja et al. (2025) retain all articles that provide information about commodity markets that reference any commodity that forms a constituent of the Composite Spot Commodity Index from the Standard & Poors, Goldman Sachs Commodity Index (GSCI) spot price series.[7]

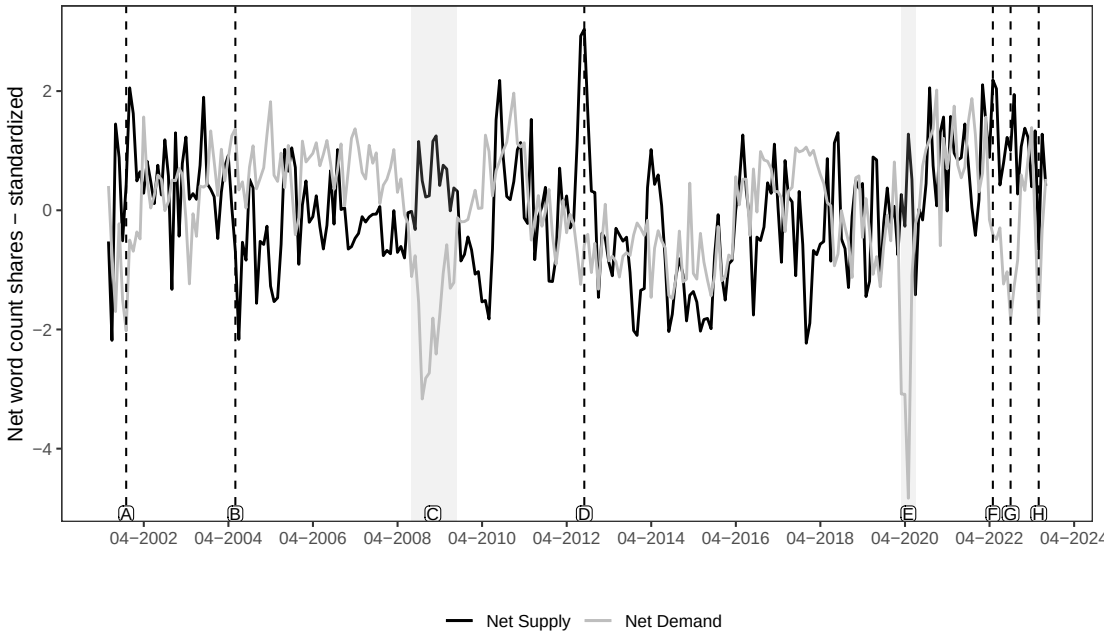
Figure 1 from Lumbanraja et al. (2025) plots the monthly standardized net supply and demand indicators for the full sample period, after adjusting the net supply and demand indicators with the total number of articles per day. The peaks and troughs of the standardized net supply and net demand indices map to well-known commodity-wide developments. Large spikes correspond to major events such the global financial crisis, the trade war, the COVID-19 epidemic, the Russo-Ukrainian War and the U.S. debt-ceiling crisis. Sizable net supply peaks correspond to important oil production cuts by OPEC and to natural disasters.

Measures of Supply and Demand Developments for Individual Commodities. The narrative approach is subsequently tailored to cover commodity categories (energy, industrial metals, precious metals, agricultural commodities and livestock) and individual commodities; oil, natural gas, heating oil, copper, aluminium, zinc, nickel, wheat, corn, soybean, sugar, cocoa, coffee, cotton, cattle and hogs. Methodologically, this is achieved by augmenting the dictionaries of the broad commodity indicators with words that capture commodity-specific supply and demand elements.

Figures A1 and A2 in the appendix plot the standardized net supply and demand indicators for crude oil and wheat for the 2001-2023 period. Interestingly, while the crude-oil narrative indices map to well-known oil developments, such as OPEC meetings and other notable business cycle-related events, including the weakening global demand following the U.S. recessions of 2001 and 2008, and the dramatic increase in US production between 2014 and 2016, wheat is marked by events of a different nature. The wheat index captures a number a supply disruptions that coincide with extreme weather events and natural disasters that affect production. The peaks in demand, which appears to be rather inelastic, coincide with trade deals between U.S. and China and stockpiling around the pandemic, which is indicative of precautionary demand.

[7] This index tracks the prices of major physical commodities for which there are active, liquid futures markets.

FIGURE 1. Standardized Net supply and Standardized Net Demand across All Commodity-Related News Articles



Note: This figure plots the standardized net supply and demand indicators for the period between 2001 and 2020. The bars map a number of well-known commodity-wide developments. These events are: [A] US recession and events of 9/11 weighed on commodity markets (118th Extraordinary Meeting of the OPEC Conference – production cut), [B] OPEC agreed to raise output quotas (131st Extraordinary Meeting), [C] Global Financial Crisis, [D] Worst drought in more than 50 years in the U.S. severely cut corn and soybeans productions, sparking all-time high prices, [E] COVID-19 dented demand and disrupted supply chains, [F] The Russo-Ukrainian War raised concerns about supply shortages in grains, oil, and metals, [G] Demand concerns from rising global interest rates, China's COVID-19 lockdown extensions, [H] 2023 United States debt-ceiling crisis.
Source: Lumbanraja et al. (2025)

3. PROPERTIES OF SUPPLY AND DEMAND MEASURES

To better understand the behavior and properties of the supply and demand indices from Lumbanraja et al. (2025) in the context of our analysis, we carry out a number of tests. Our premise is that when aggregate demand shocks occur, such as during an economic expansion or increased consumer confidence, they exert upward pressure on prices across a broad spectrum of commodities. This phenomenon is known to be more pronounced in procyclical goods, such as oil or industrial metals but is also present in other commodity categories. For example, an increase in consumer spending often correlates with rising oil prices, which, in turn, influences transportation and production costs across most sectors. As these costs rise, the higher prices tend to get passed through to consumers, leading to sustained inflationary pressures. This

relationship between aggregate demand and commodity prices highlights a significant and homogeneous component that links the behavior of individual commodities to their contribution to the overall price level.

Empirically, the textual composite demand measure captures an aggregate component that reflects a broad-based change in overall economic activity. By encompassing various demand signals across different commodities, this measure allows us to discern shifts (e.g. in consumer behavior or investment patterns) that signal fluctuations in economic performance. Its importance lies in its ability to provide a comprehensive view of demand trends rather than focus solely on isolated sectors. Descriptive evidence highlights this property, as the composite demand measure displays a high correlation with a textual measure of the business cycle from Lumbanraja et al. (2025). This correlation suggests that as demand across commodities shifts, it coincides with changes in important business cycle components. For instance, a rise in consumer confidence leads to increased spending, which, in turn, bolsters demand for goods and services across multiple sectors, including manufacturing and retail. This interconnectedness reinforces the composite measure's role as a valuable predictor of economic trends. Additionally, we can measure aggregate demand through the calculation of the first principal component of all commodity-specific demand indicators. The fact that this first principal component correlates with a textual measure of the business cycle corroborates this point. Figure 2 plots the 1st principal component of individual commodities' net demand indices, the net demand composite index, and a measure of the business cycle, illustrating the high correlation across these measures which is in line with the mechanism of the paper: aggregate demand shocks tend to be correlated across different commodities, leading to a strong pass-through effect on the consumption basket.

In contrast, supply shocks tend to exhibit idiosyncratic characteristics that can lead to diverse outcomes depending on market conditions and the interrelationships between different goods. Taking the cocoa market as an example, a negative supply shock - perhaps a disease that affects cocoa crops - can result in a sharp increase in cocoa prices. However, since cocoa and sugar are typically used for the production of chocolate, the reduced supply of cocoa can depress the demand for sugar. In such cases, the complementarity between these two commodities results in divergent pricing movements, where sugar prices may decrease even as cocoa prices rise. Thus, while cocoa experiences significant inflationary tension, the overall impact on the consumer price index could be mitigated, showcasing the nuanced effects of supply shocks. The behavior of

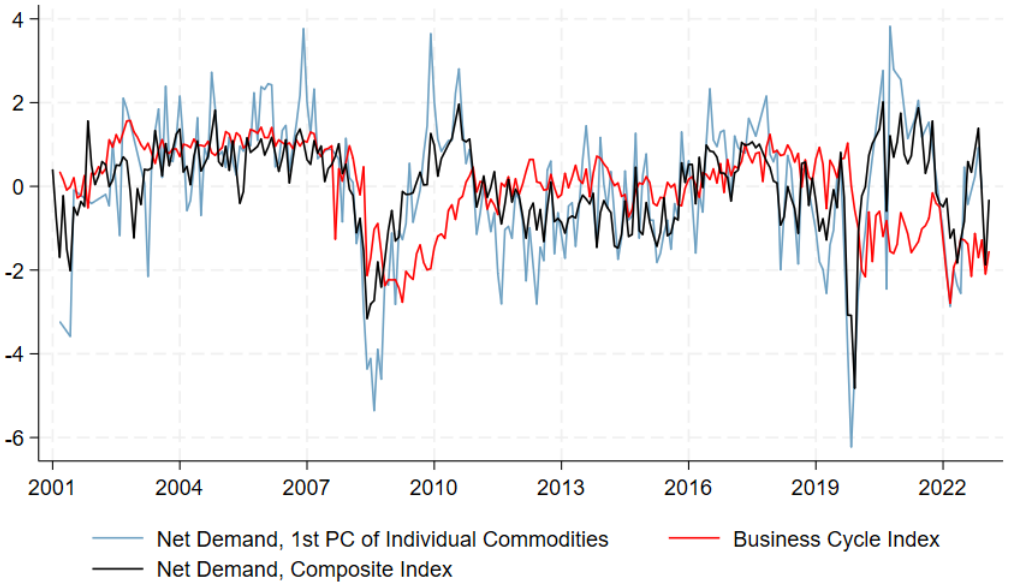
commodities such as coffee, which act as imperfect substitutes for cocoa further complicates the aggregate effects on prices baskets. A supply disruption affecting cocoa may not lead to a proportional increase in demand for coffee and subsequently coffee prices, but instead may see consumers shifting preferences to tea or other substitutes. This limited demand response within the context of substitute goods demonstrates that supply shocks have unique properties, but also interact with consumer behavior in ways that can dampen their broader inflationary impact. Hence, the interconnectedness of commodities through substitutability and complementarity could lead to varying inflation outcomes even when specific commodities face supply constraints.

Figure 3 plots the eigenvalues of the principal components of textually-extracted individual commodity net demand and supply developments and provides some evidence that illustrates this point. The scree plot of the principal component analysis of individual commodity net supply indicators does not suggest a natural break between high and low eigenvalues - however, the first eigenvalue from the principal component analysis of the individual commodity net demand indicators is notably larger, showcasing the importance of an aggregate demand component.

Furthermore, Figure 4 presents the pairwise correlations of the supply and demand indices of individual commodities with those of oil. We focus on this relationship due to the unique importance of oil in the production stages across economic sectors, but also due to its prominence in a literature that has historically focused solely on oil behavior to proxy for global supply chain disruptions. With the exception of wheat and sugar (both essential food products with relatively inelastic consumer demand) pairwise correlations of net supply for individual commodities with oil are notably lower than the corresponding correlations of net demand, consistent with our argument.

Finally, it is important to recognize that different supply dynamics across commodity sectors further complicate the inflation landscape. While demand dynamics tend to exhibit significant commonality across diverse commodities, supply dynamics are more fragmented. For example, agricultural commodities like grains may respond strongly to extreme weather events, while energy and metals markets often display low correlations in their supply shocks. Therefore, the idiosyncratic nature of supply dynamics (shaped by seasonal conditions, regional disruptions, and specific industry characteristics) is central in commodity price transmission. In particular, the presence of complement and substitute goods becomes pivotal as they influence how supply shocks propagate through the economy, reinforcing the notion that while commodity-specific

FIGURE 2. Net Demand Indices and the Business Cycle



Notes: The figure plots the 1st Principal Component of net demand indices of individual commodities, the net demand composite index, and the (negative) of the business cycle index. Commodities included are aluminium, cattle, cocoa, coffee, copper, corn, cotton, gasoline, hogs, natural gas, oil, soybean, sugar, wheat and zinc, spanning the period May 2001-June 2023.

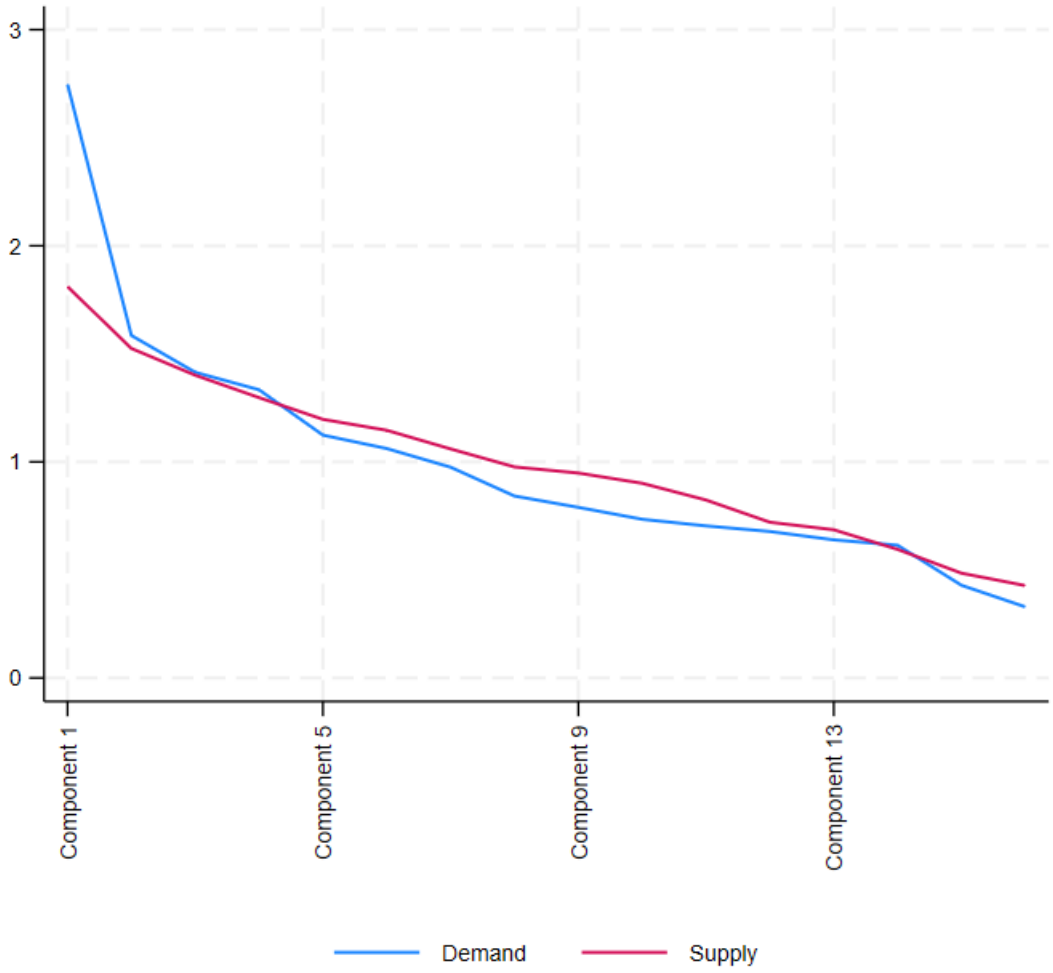
disruptions can elevate prices, their ultimate effect on inflation may be significantly muted due to a weak pass-through effect on the overall consumption basket.

4. DO THE NARRATIVE COMMODITY INDICES IMPROVE INFLATION FORECASTS?

The analysis presented in the previous sections suggests that textual measures of commodity supply and demand, derived from business news, may contain valuable informational embedded in equilibrium aggregate commodity prices. This finding prompts further investigation into whether differentiating between supply- and demand-side dynamics may provide distinct information about future inflation relative to existing predictors, potentially leading to more accurate inflation forecasts.

In this section, we evaluate the proposition outlined above. The existing literature remains divided on the utility of commodity prices for inflation forecasting, particularly during the Great Moderation and the post-Great Financial Crisis period. Rather than taking a definitive position on the general forecasting value of commodity prices, we focus on whether our textual measures of supply and demand dynamics can enhance inflation predictions relative to baseline models, both with and without the inclusion of commodity prices. The analysis is conducted for both headline inflation and core inflation, the latter excluding volatile components such as food and energy prices. By

FIGURE 3. Scree Plots from Principal Component Analysis

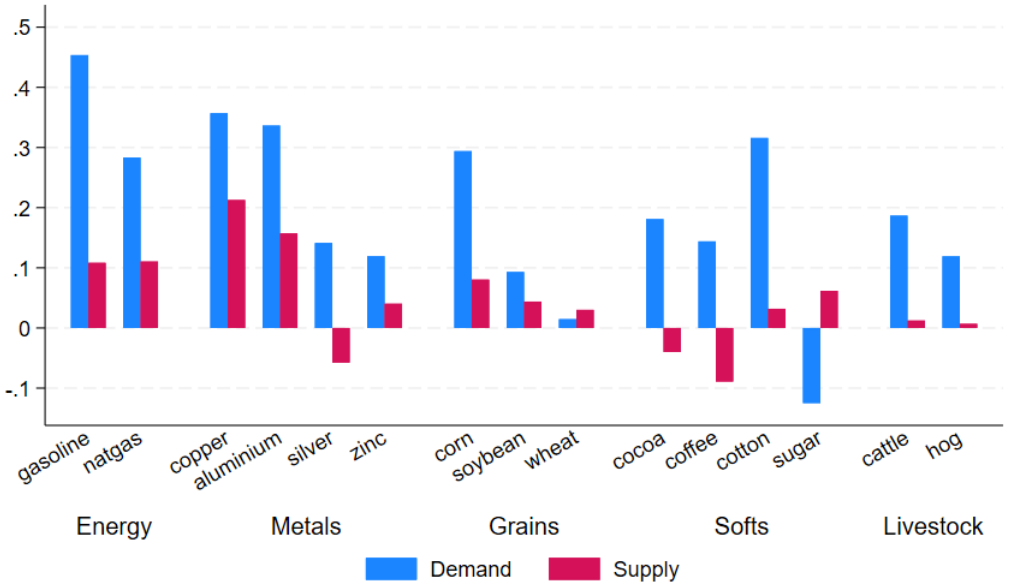


Note: Scree plots of eigenvalues after principal component analysis conducted on net supply and net demand indices of individual commodities including aluminium, cattle, cocoa, coffee, copper, corn, cotton, gasoline, hogs, natural gas, oil, soybean, sugar, wheat and zinc, spanning the period May 2001-June 2023.

doing so, we aim to assess the incremental value that text-based indicators provide over traditional forecasting models in capturing inflation dynamics, especially in the context of heterogeneous drivers across different inflation measures.

We conduct two broad sets of empirical exercises. In the first, we incorporate our textual indices into a baseline model, without the inclusion of tradable commodity prices or other control variables. This allows us to assess the standalone predictive value of our measures. In the second set of exercises, we introduce commodity prices along with additional macroeconomic and financial controls, evaluating whether our

FIGURE 4. Correlations of commodity indices with oil index



Notes: The figure shows correlations of net demand and net supply for each individual commodity against oil, spanning the period May 2001-June 2023.

indices offer incremental informational content beyond the effects captured by commodity price movements and other macroeconomic and financial variables. The analysis employs a simple linear model, adhering to parsimonious specifications commonly used in the literature, particularly following the frameworks of Gospodinov and Ng (2013a) and Stock and Watson (2003). Additionally, existing literature has shown that simple, autoregressive specifications are not easy to improve on (Faust and Wright, 2013; Bańbura, Lenza, and Paredes, 2024). Specifically, we estimate linear regressions of the following form:

$$\pi^h_{t+h} = \alpha_0 + \beta_0(L)\pi_t + \beta_1(L)index_t + \beta_2(L)x_t + \epsilon_{t+h}. \tag{7}$$

In the above expression, π^h_{t+h} represents inflation over a horizon of h months from time t, where $\pi^h_{t+h} = (1200/h) \times ln(p_{t+h}/p_t)$, with p_t denoting the price level at time t. Therefore, $\pi_t = 1200 \times ln(p_t/p_{t-1})$ is the one-month annualized inflation rate at time t. The coefficients $\beta_k(L)$ are lag polynomials, $index_t$ denotes our textual measure of commodity demand or supply, and x_t control variables.[8]

For the choice of control variables x_t, we closely follow the literature (Gospodinov and Ng, 2013b; Garratt and Petrella, 2022), and include the policy rate (effective federal

[8]Although these models are not strictly autoregressive, we characterize them as autoregressive for convenience. As the dependent variable is inflation h periods ahead, and the regressor is a polynomial of one-period inflation, these models are only truly autoregressive if $h = 1$.

funds rate), the unemployment gap (to capture Phillips Curve effects), a yield curve measure (the 10 year minus the 2 year Treasury yield), and the log growth of a trade-weighted dollar index. In each exercise we also separately add the one-month log growth in aggregate commodity prices, in order to show the performance of our indicators with and without commodity prices, as it is not clear from the literature whether these do improve prediction. We employ rolling window forecasts, using 120-month windows as is standard in the literature. A potential concern with such long windows is that crises are frequently present in the sample.[9] The $\beta_k(L)$ terms are of order 1.[10]

The richness of our indices allows for a variety of specifications, which can tease out different patterns in the data. For simplicity, we focus on our composite indices in this section. In line with the parsimony requirement, we examine several combinations of these indices. Specifically, we consider models incorporating: i) net supply and net demand together; ii) only net supply; iii) only net demand; iv) demand increase and decrease; v) supply increase and decrease; vi) supply and demand increase and decrease. Initially, we limit our analysis to the pre-COVID period to establish baseline results. Subsequently, we address the econometric challenges unique to the COVID-19 period, to enhance the robustness of our results in the face of the unprecedented disruptions associated with the pandemic. These distinct specifications allow us to capture the nuanced effects of supply and demand shocks across different time-frames and economic conditions, thereby enhancing the reliability of our inflation forecasting models.

We analyze both Consumer Price Index (CPI) and Personal Consumption Expenditures (PCE) inflation, acknowledging the significant differences between these measures and their potential impact on our findings. CPI is a Laspeyres index derived from consumer survey data, while PCE represents a Fisher-ideal index based on business surveys. Notably, CPI captures only out-of-pocket expenditures, whereas PCE encompasses indirect expenditures, such as those covered by employer-provided health insurance. The PCE is the Federal Reserve's preferred inflation metric; in contrast, the European Central Bank (ECB) focuses on the Harmonized Index of Consumer Prices (HICP), which is also a Laspeyres index sourced from consumer surveys. Although

[9] In addition to the simple AR models, as specified in (7), we also considered models incorporating moving average terms, which have been shown to improve inflation forecasts. While these indeed provided prediction improvements until 2019, their performance drastically deteriorates during the high inflation episode of 2021-2023, with baseline errors 30-40% worse than the AR models.

[10] We also test AR models with two lags, yielding similar results. Lag selection criteria were not employed, as they produce different optimal structures depending on the covariates, and we aimed to avoid confounding our results with varying lag structures.

our analysis centers on the U.S. context, the universal applicability of our measures ensures that our examination of CPI remains relevant and valuable in broader economic discussions.

4.1. **Accounting for the COVID shock.** The unprecedented shock induced by the pandemic has complicated standard time series estimation methods, primarily due to the presence of extreme outliers in the data. This phenomenon has prompted the literature to explore various strategies for addressing these challenges. Researchers have investigated a range of potential solutions, seeking to mitigate the distortions caused by these outliers and enhance the robustness of econometric estimates in the context of such peculiar economic conditions.

In our analysis, we encounter two principal challenges. The first is the non-stationarity introduced by the COVID-19 pandemic, which has significant implications for parameter estimation and the reliability of our models. Second, even if we achieve accurate parameter estimates, our demand indices experience a dramatic collapse in magnitude during the initial three months of the shock. This pronounced decline complicates the forecasting process for the relevant time horizons, as the stability and predictive power of our indices are undermined by these extraordinary conditions.

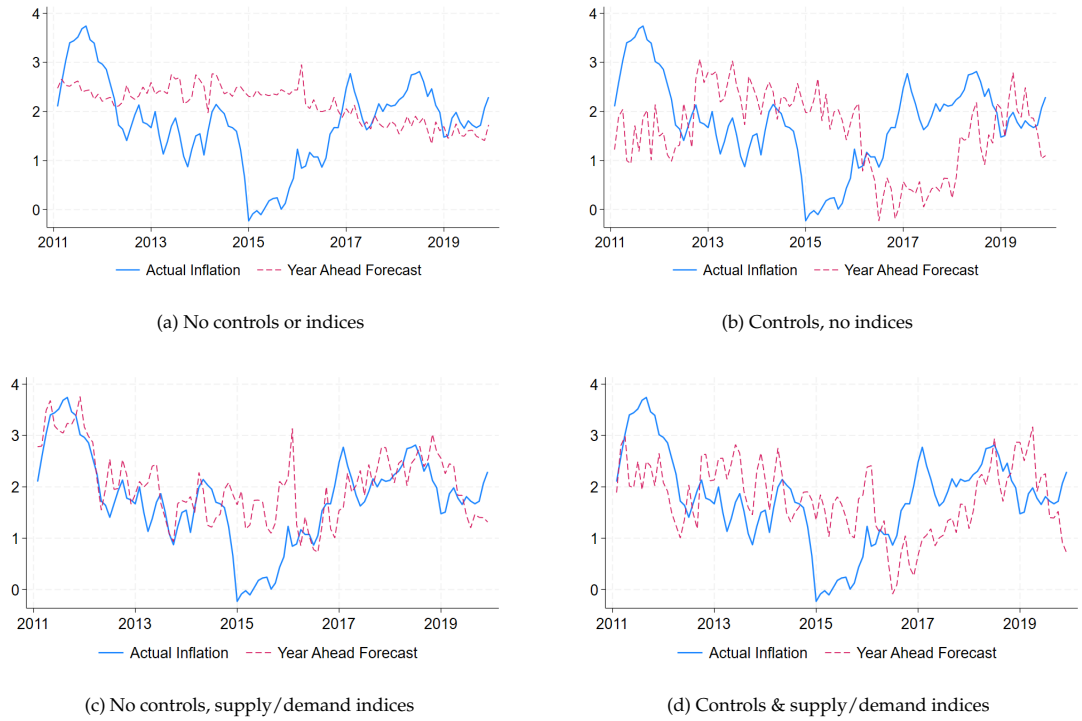
To address this challenge, we adopt the approach proposed by Lenza and Primiceri (2022), who contend that while excluding extreme outliers from the COVID episode may be suitable for parameter estimation, it is not advisable for forecasting, as this practice risks underestimating uncertainty. They recommend incorporating stochastic volatility within a Bayesian VAR framework, which effectively scales down the variables of interest during the months impacted by the shock. To preserve the parsimony of our model, we implement a simplified version of their approach by replacing the values of the demand indices for the initial three months of the shock with their mean from 2019. Additionally, we include a dummy variable equal to one for the COVID era in our parameter estimation, as suggested by Ng (2021).[11]

4.2. **Forecasts.**

We present our results separately for the pre-COVID period and the entire sample, to evaluate the performance of our framework across different economic regimes. We begin with the forecasting results for Consumer Price Index (CPI) inflation during the pre-COVID period, as detailed in Figure 5 and Table 1.

[11] Ng (2021) also advocates for incorporating the log growth of hospitalizations or infections to account for the shock. While this method does enhance prediction accuracy, it offers a smaller improvement compared to the scaling-down approach.

FIGURE 5. Inflation Forecasts versus Actual Inflation, Pre-COVID



(a) No controls or indices

(b) Controls, no indices

(c) No controls, supply/demand indices

(d) Controls & supply/demand indices

Note: Inflation forecasts (12-months ahead) plotted together with actual inflation spanning the period February 2011-December 2019. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar monthly log change, the unemployment gap, and log monthly changes in commodity prices. When included, indices are the net supply and demand composite indices.

Figure 5 plots 12-month ahead inflation forecasts along with actual inflation for the period between December 2010 and December 2019. The top panels display the behavior of forecasts derived from models that only contain autoregressive components (Panel A), or a full set of controls including log monthly changes in commodity prices (Panel B). The bottom panels further introduce the textual indicators of net demand to the aforementioned models. The improvement in the forecasting performance is evident. Including the full set of controls leads to an overall worsening of the forecast performance, which is not uncommon in inflation forecasting, where autoregressive models often dominate. Still, the specification with net demand (plus controls) clearly gives a better prediction.[12]

[12]Note that the first few forecasts have significantly lower sample size than 120, because our indices start in May 2001. This is not in general a problem, except for the first two periods of the out-of-sample forecast, which, when including controls, give a structural break. As such, we omit the first two periods in the chart, but do include them in the tables. Note that, this small sample issue is not a concern for specifications without additional controls.

The upper part of Panel A of Table 1 presents the full results for models that do not incorporate any control variables (i.e., the set x is null). The Root Mean Square Error (RMSE) for the baseline model across all specifications is reported in the second column for each forecast horizon. We assess the statistical significance of the forecasts relative to the benchmark using the Diebold-Mariano test. Focusing on headline inflation at the h=12 horizon, most Autoregressive (AR) models demonstrate a clear improvement over the baseline, particularly those in columns 1 and 3, which feature a combination of net demand and net supply and net demand alone, respectively. Secondary enhancements are observed in columns 4 and 6, which consider demand increase/decrease and supply/demand increase/decrease. This indicates that demand serves as a considerably stronger predictor of future price growth. We also evaluate specifications that include a set of control variables: the effective federal funds rate, the spread between 10-year and 2-year U.S. Treasury yields, trade-weighted U.S. dollar monthly log returns, and the unemployment gap. Results are reported in the lower part of Panel A of Table 1. Our framework consistently yields significant improvements for the h=12 horizon even against this benchmark, though the performance of the baseline model is worse with the inclusion of the controls.

Panel B presents the results for a baseline model that incorporates commodity price returns as a predictor. Compared to Panel A, the Root Mean Square Error for the autoregressive model is generally lower (however, these differences are not statistically significant). For the h=12 horizon, the models identified as successful in Panel A continue to provide substantial improvements in predictive accuracy in Panel B, while additionally the model in column 3 now also yields highly significant improvement for h=6 (though the magnitude improvement is in fact smaller). This finding indicates that, despite the inclusion of commodity price information in the baseline model, our indices enhance forecasting performance. When controls are included (and the overall predictive power is lower across the board), the predictive gains for the autoregressive model persist only for the h=12 horizon.

It is important to note that the baseline model presented here, which incorporates commodity inflation, exhibits a lower Root Mean Square Error (RMSE) than that observed in Panel A. Consequently, our predictive gains are understated when compared to a simple autoregressive model. Nevertheless, given that our primary objective is to demonstrate that our indices offer additional informational value irrespective of the inclusion of commodity price inflation, we concentrate our analysis on within-class comparisons.

TABLE 1. Forecast performance (RMSE), Headline CPI inflation, Pre-COVID

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: without commodities								
Autoregressive								
	h=12	1.034	0.675*	1.019	0.696**	0.821***	1.096	0.756**
	h=6	1.411	0.848	1.059	0.831	0.910	1.118	0.996
	h=3	1.729	1.035	1.129	0.959	1.035	1.208	1.182
Controls, Autoregressive								
	h=12	1.308	0.776**	1.061	0.732***	0.821*	1.088	0.790**
	h=6	1.600	0.998	1.022	0.992	1.003	1.098	1.076
	h=3	1.988	1.025	1.023	1.007	1.011	1.075	1.082
Panel B: with commodities								
Autoregressive								
	h=12	0.979	0.722*	1.020	0.734**	0.890	1.122	0.837**
	h=6	1.418	0.898	1.067	0.864**	0.930	1.085	1.007
	h=3	1.695	1.052	1.124	0.958	1.013	1.164	1.168
Controls, Autoregressive								
	h=12	1.281	0.816**	1.046	0.767***	0.854*	1.083	0.836**
	h=6	1.584	0.985	1.039	0.976	0.991	1.110	1.079
	h=3	1.837	1.026	1.034	0.995	1.006	1.082	1.097
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease						X	X	
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of headline CPI inflation across different monthly horizons, using the composite commodity indices. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Figure 6 plots inflation forecasts together with actual inflation for the period between February 2011 and June 2024. As before, the top panels display the behavior of forecasts derived from models that only contain autoregressive components (Panel A), or also a full set of controls including log monthly changes in commodity prices (Panel B). The bottom panels further introduce the textual indicators of net demand to the aforementioned models. The inability of all models to timely forecast the inflation surge following the outbreak of the pandemic is obvious and expected. However, there is still some improvement in the forecasting performance of the latter period of the sample when one employs the textual indicators, over the benchmark case or the specification including controls and log monthly changes in commodity prices. Once again, the inclusion of further controls leads to worse performance.

TABLE 2. Forecast performance (RMSE), Headline CPI inflation

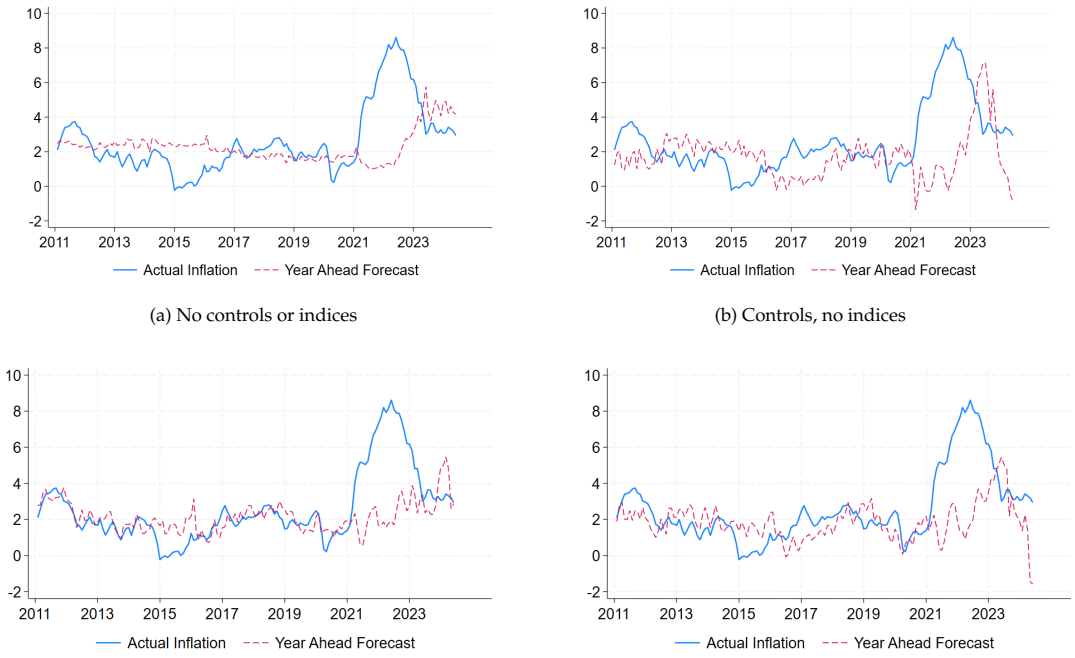
	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: without commodities								
Autoregressive								
	h=12	2.177	0.872*	0.994	0.884*	0.912	1.004	0.885
	h=6	2.235	0.894*	1.008	0.894**	0.934	1.029	0.961
	h=3	2.560	1.006	1.039	0.985	1.030	1.064	1.076
Controls, Autoregressive								
	h=12	2.562	0.815*	1.009	0.812**	0.845*	1.013	0.819**
	h=6	2.542	0.897	1.004	0.898	0.923	1.027	0.943
	h=3	2.791	0.996	1.002	0.992	1.021	1.022	1.027
Panel B: with commodities								
Autoregressive								
	h=12	2.149	0.895*	0.998	0.904*	0.931	1.016	0.909
	h=6	2.213	0.924*	1.018	0.913**	0.945	1.027	0.976
	h=3	2.375	1.029	1.050	0.989	1.024	1.063	1.091
Controls, Autoregressive								
	h=12	2.567	0.815*	1.007	0.811*	0.842*	1.014	0.815*
	h=6	2.503	0.909	1.015	0.904*	0.927	1.037	0.956
	h=3	2.534	1.023	1.015	1.006	1.032	1.036	1.063
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease						X	X	
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of headline CPI inflation across different monthly horizons, using the composite commodity indices. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

Table 2 shows results for the full sample, which includes the COVID period. As expected, the baseline performance is substantially worse than before, due to the large inflation shock. Despite these shifts in baseline performance, the relative effectiveness of our models compared to the baseline remains, even compared to models which include commodity prices. The inclusion of controls reduces performance markedly, especially for the h=12, but in fact the performance of our preferred models is affected little (i.e., while the autoregressive benchmark RMSE rises a lot with the inclusion of controls, the benchmark RMSE with our indices stays essentially the same). Thus, within each class of models examined, our indices yield substantial predictive gains, highlighting their relevance and utility across varying inflationary regimes.

Overall, our commodity demand and supply indices appear to contain highly informative content for forecasting aggregate price indices. Demand-driven indices yield significant improvements in predicting headline inflation, particularly over longer horizons, when no additional controls are included. Even when controlling for commodity

FIGURE 6. Inflation Forecasts versus Actual Inflation



(a) No controls or indices

(b) Controls, no indices

(c) No controls, supply/demand indices

(d) Controls & supply/demand indices

Note: Inflation forecasts plotted together with actual inflation spanning the period February 2011-June 2024. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar monthly log change, the unemployment gap, and log monthly changes in commodity prices. When included, indices are the net supply and demand composite indices.

price inflation and macro and financial variables, these indices continue to provide notable forecasting gains, particularly for longer horizons.

Tables 3 and 4 present the results for headline PCE inflation.[13] Given the key differences between CPI and PCE measures, the results are not necessarily expected to align. For PCE, the baseline RMSE is generally lower than for CPI. In all, the performance of our indices in improving forecasts relative to the baseline is similar with CPI.

We provide a summary graphical measure of the strength of our indices in inflation forecasting in Figure 7, in the spirit of Hong et al. (2025), who also use news data. We plot the cumulative sum of squared prediction errors (SEE) for our preferred specification (with net demand and commodity prices) relative to the cumulative SSE for the autoregressive benchmark (also with commodity prices). We do so for both the CPI

[13]For completeness, we report, in tables C1 and C2 in the appendix, results for core CPI and PCE inflation, which remove food and energy items from the index. Core inflation presents a much more challenging forecasting benchmark, as it includes far fewer commodity-driven components relative to headline, and is substantially less volatile. A notable difference between Core CPI and PCE is the latter includes food services (e.g. restaurant meals). As such, we do see some forecast improvements for Core PCE, though few specifications are statistically significant).

TABLE 3. Forecast performance (RMSE), Headline PCE inflation, Pre-COVID

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: without commodities								
Autoregressive								
	h=12	0.834	0.699	1.005	0.719*	0.841**	1.040	0.763*
	h=6	1.034	0.811	1.044	0.806	0.856	1.086	0.938
	h=3	1.242	0.993	1.082	0.954	1.010	1.162	1.132
Controls, Autoregressive								
	h=12	1.016	0.842	1.095	0.773***	0.882	1.095	0.814
	h=6	1.165	0.992	1.019	0.993	0.998	1.063	1.051
	h=3	1.410	1.033	1.014	1.011	1.016	1.051	1.076
Panel B: with commodities								
Autoregressive								
	h=12	0.796	0.745	1.012	0.756*	0.902	1.063	0.831*
	h=6	1.030	0.865	1.056	0.848*	0.886	1.051	0.950
	h=3	1.173	1.004	1.094	0.950	0.989	1.124	1.115
Controls, Autoregressive								
	h=12	0.990	0.905	1.086	0.805***	0.919	1.095	0.872
	h=6	1.113	0.993	1.031	1.001	1.005	1.081	1.072
	h=3	1.240	1.009	1.015	0.999	1.007	1.057	1.073
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease						X	X	
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of headline PCE inflation across different monthly horizons, using the composite commodity indices. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

and PCE, for the 12- and 6-month horizons. The figure shows a more or less consistent increase in cumulative SSE for the full sample, indicating that the narrative indices outperform the autoregressive benchmark essentially all the time (as noted by Hong et al. 2025). The time variation is quite similar to that of Hong et al. (2025), who also show a gradual increase in relative cumulative SSE from 2010-2015, and then a plateau until COVID hit, when the improvement became even larger. Their earlier sample also shows outsize gains during the global financial crisis; taken together, the results in both papers show the important role that news can play in the timely analysis of inflation.

4.3. **Forecasting performance across different inflation regimes.**

Another way to assess the performance of our measures is to see how well they perform across different inflation regimes. Good average performance may mask bad

TABLE 4. Forecast performance (RMSE), Headline PCE inflation

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: without commodities								
Autoregressive								
	h=12	1.741	0.899**	0.998	0.906**	0.928	0.996	0.904
	h=6	1.637	0.896*	1.006	0.899**	0.936	1.022	0.965
	h=3	1.744	0.993	1.024	0.985	1.028	1.052	1.069
Controls, Autoregressive								
	h=12	1.998	0.847	1.015	0.840*	0.875*	1.016	0.847*
	h=6	1.808	0.901	1.003	0.907	0.932	1.019	0.948
	h=3	1.873	0.999	0.998	0.997	1.025	1.013	1.033
Panel B: with commodities								
Autoregressive								
	h=12	1.739	0.910**	1.001	0.915**	0.934	1.005	0.912
	h=6	1.622	0.918**	1.015	0.913***	0.942	1.018	0.971
	h=3	1.621	1.001	1.034	0.981	1.014	1.044	1.068
Controls, Autoregressive								
	h=12	2.006	0.842	1.014	0.829*	0.863	1.017	0.833*
	h=6	1.765	0.905*	1.011	0.909	0.931	1.030	0.953
	h=3	1.693	1.003	1.004	1.000	1.023	1.020	1.044
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease						X	X	
Demand Increase/Decrease						X		X

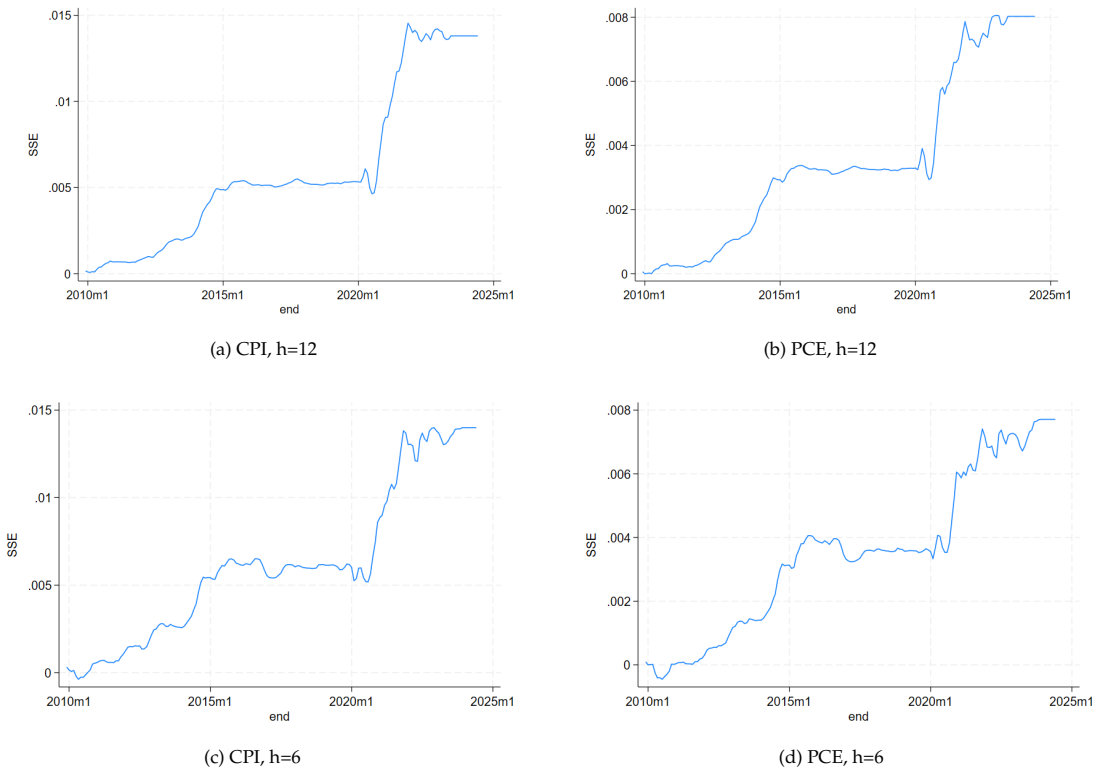
Note: The table shows results from forecasts of headline PCE inflation across different monthly horizons, using the composite commodity indices. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

performance at time of falling or rising inflation, rendering the framework problematic when good forecasts are most important.

We follow Joseph et al. (2021) and split the sample in periods when inflation is increasing, falling or is stable. We define an episode as increasing when the 3-month moving average change in inflation is positive for at least 5 consecutive months; falling when the same measure registers a negative change for at least 5 consecutive months; and stable when it is neither falling nor rising.

Results are shown in Appendix Tables D1 for CPI and D2 for PCE, for h=12 and h=6 horizons. Overall our framework provides clear improvements over the baseline for headline inflation for h=12 for models with demand indices, especially for episodes of falling and stable headline inflation, and to a somewhat lesser extent for episodes with increasing inflation. For h=6 for CPI, similar patterns hold, though improvements

FIGURE 7. Difference in Cumulative SSE



(a) CPI, h=12

(b) PCE, h=12

(c) CPI, h=6

(d) PCE, h=6

Note: The figures plot the difference between the cumulative sums of squared errors (SSE) between the autoregressive (AR) model with commodity prices, and the model which includes the net demand indices.

for episodes of increasing inflation are generally weak.[14] For PCE, models with demand indices provide improvements for all episodes for h=12 and for both falling and increasing inflation episodes for h=6.

4.4. **Robustness Checks.**

4.4.1. *Alternative specifications.* In the preceding section, we used specifications standard in the literature to assess the performance of our indices. Here, we discuss a number of robustness checks and alternative specifications.

First, we examine whether the reliance on the MSE criterion may affect our conclusions. A popular alternative is the mean absolute error (MAE), given by the mean of the absolute differences between two forecasts. The literature frequently considers both MSE and MAE in tandem; we opted to focus on MSE for brevity, but MAE provides an important cross-check, particularly since it is less sensitive to outliers. Results for this alternative criterion are shown in Appendix tables E1 and E2. Results in favor

[14] Note that because the sample can be quite small for some cases, we do not perform Diebold-Mariano tests.

of our indices are substantially stronger now. Our preferred specifications are even more statistically significantly better than the baseline across all specifications for the 12-month horizon, and in some cases also for the 6-month horizon (full sample with commodities).

We then consider the role of the timeliness of data releases, and whether using real-time data has a bearing on our results. Macro time series often undergo revisions and currently published series may differ, often substantially, from what was available to forecasters in real time. This is not a major concern in our case, since most of our variables are either not revised (e.g. the financial data), or, in the case of price indices, are revised lightly. Still, it is useful to make sure that the performance of our indices is preserved with real-time data. We obtain real-time series from the Philadephia Fed Real-Time Data Set for Macroeconomists, for CPI, unemployment and the NAIRU. Results for CPI inflation are reported in Appendix Table E3. As expected, baseline forecasts (not including covariates) are barely affected. The largest discrepancies arise only in the specifications with the unemployment gap. In all these cases, the relative performance of our indices is unaffected.

Another alternative involves expanding the set of possible controls. Given that our time series span is limited, we are restricted in the number of regressors we can add, otherwise we risk overfitting and a reduced forecasting performance. Moreover, results are sensitive to the choice of covariates and the sample used (with and without the high inflation period after COVID). As such, in order to better gauge the predictive power of different regressors across samples, we adopt the following approach. We take the baseline autoregressive model and two models where we add our two preferred sets of indices to the autoregressive model, and add 10 different predictors, one at a time. The two preferred sets of indices are the one where we only include Net Demand (ND), and the one where we have Net Demand + Net Supply (NDNS).

The predictors are the four we included before (EFFR, yield curve, the log monthly growth of the trade-weighted dollar index and the unemployment gap), plus the log monthly growth of industrial production (IP) and the SP500 index (SPX), the CBOE Volatility Index (VIX), together with three measures of inflation expectations: the Survey of Professional Forecasters (SPF), the options-implied market expectations from the Cleveland Fed (CLEV) and the household expectations from the Michigan Survey (MICH). Results are shown in Tables E4 and E5 in the Appendix, for CPI and PCE,

respectively, for forecasts with a 12-month horizon. Each line shows the relative performance of adding further predictors to benchmark (either the autoregressive model, the ND or the NDNS).

Two main features of the results stand out. First, no predictor manages to improve on the predictive ability of our forecasts. This is the case either when the alternative predictors are included to the autoregressive baseline on their own, or together with our indices. This is the case either before COVID or in the full sample, in specifications with and without commodity price growth, and either for CPI or PCE. Second, the gains from including the additional predictors are uneven across samples. This is most clearly shown for USD and VIX. Including USD together with commodities in the pre-COVID sample leads to forecasting gains against all three baseline models. It is in fact the only case where the gains are statistically significant. By contrast, including USD in the full sample specification leads to a slight worsening of the forecast. We get the opposite result for VIX, which leads to improvements over the full sample, but substantial worsening for the pre-COVID sample. Including any of the expectations measures does not improve any of the baselines, and sometimes leads to substantial worsening.

Furthermore, in tables E6 and E7 we show results where we include a third lag for the autoregressive component, our indices, and the commodity price growth. This is not an optimal choice, either in terms of standard lag selection criteria or predictive performance. The predictive errors for almost all specifications are higher in this case, but the relative performance of our indices remains intact. Adding either further lags to the other predictors or to our indices and the autoregressive component leads to substantially worse forecasts for all models.

4.4.2. *Alternative composite index weighting scheme.* Next, we consider the role of the weighting scheme used in the construction of the composite indices. In our inflation forecasts, we face a methodological constraint: including separate supply and demand variables for each commodity would push against few degrees of freedom and produce poor forecasts. Instead, we employ composite indices that allow commodities to be naturally weighted by their prevalence in news coverage. This implies a choice of how to weigh the different commodities in the construction of the composite. The composite indices used in forecasting are (implicitly) news-aggregated, in the sense that the full sample of news articles in the dataset is used to extract the indices, without distinguishing among commodities. Commodities with more news coverage will

naturally feature more prominently in the indices. Our view is that this method is useful precisely because it captures news bias, and more salient commodities or events (which have larger aggregate price effects) are picked more in the news. Note that this implicit weighting is, by definition, time-varying: commodities have a higher weight in the indices precisely when they experience newsworthy supply or demand shocks. The net demand index particularly benefits from this approach, effectively capturing underlying demand factors through higher and common news coverage patterns.

An alternative approach would be to use traditional fixed weights based on global production (analogous to the GSCI Composite Index methodology), using five-year averages of production and trading data between 2019 and 2024. Table E8 in the Appendix presents results from forecasting exercises (with commodity price growth but no other controls), using the fixed production-weights weighting schemes for the calculation of the composite indices. We only focus on our preferred combination of textual indices (the one that includes net demand only), but include results for both CPI and PCE, and also show, for convenience, both the baseline prediction errors and the relative performance of models with the news-weighted indices.

The results indicate that the news-weighted scheme outperforms the production-weighted approach in forecasting both CPI and PCE inflation. This suggests media coverage patterns provide valuable information about which commodity movements matter most for broader inflation dynamics. However, the performance differential, though non-trivial, is not major, nor statistically significant, indicating that the primary benefit of our methodology comes from separately identifying supply and demand disturbances, as these propagate differently through the consumer price basket.

4.4.3. *Exploiting additional information from textual data.* An alternative approach to using the composite indices for supply and demand would be to use indices designed to capture explicit topics identified in the corpus of the data. We consider forecasts which include three commodity-related risk factors that Lumbanraja et al. (2025) derive from the "Commodity Special Feature" of the IMF's World Economic Outlook and from the "Commodity Markets Outlook" of the World Bank. These include a business cycle (BC) factor, a geopolitical risk (GR) factor, and a natural disasters (NDis) factor. These factors represent the major narrative themes that drive commodity price movements, according to expert texts. Appendix table E9 includes forecasts evaluations for these factors.

For comparison, we also include the relative performance of two specifications of our indices (Net Demand+Net Supply and Net Demand only). We show results for

CPI and PCE, for models with and without commodity price growth, both pre-COVID and for the full sample, using both the MSE and MAE criteria. We see that the BC and GR factors do not provide any improvement to the baseline, and in fact worsen the forecast in most cases.

The NDis factor, however, does lead to sizable improvements in the full sample estimation results. This is not very surprising, given that it spikes during the pandemic period. If we combine the NDis factor with our composite indices, the gain is greater, in fact greater than including the composite indices alone. These gains are not always statistically significant under the MSE criterion, but appear significant when we employ the MAE criterion, possibly indicating less stable forecasting gains resulting from the NDis factor.

5. Oil Price News Shocks and Oil Price Shocks

In order to assess the forecasting power of the textual indicators we do not need to make a statement about the exogeneity of commodity-supply and demand signals that are sourced from the news. This is important because the endogenous nature of commodity prices, which are influenced by macroeconomic conditions, complicates the estimation of their causal effects. To address this challenge, the literature has traditionally employed structural vector autoregressions (SVAR) to identify distinct shocks affecting oil prices. These have been identified using (i) zero restrictions (Kilian, 2009), (ii) sign restrictions (Kilian and Murphy, 2012; Lippi and Nobili, 2012; Baumeister and Peersman, 2013; Baumeister and Hamilton, 2019), (iii) narrative information (Antolin-Diaz and Rubio-Ramirez, 2018; Caldara, Cavallo, and Iacoviello, 2019; Zhou, 2019), and more recently (iv) high-frequency variation in oil prices around OPEC announcements (Känzig, 2021). Despite the extensive development of methodologies for modeling oil price movements, there is a notable scarcity of analyses on other tradable commodities, such as natural gas and wheat, which have gained significance during the recent period of turbulence.

The purpose of this section is to show that the Oil Net Supply and Net Demand Indices contain valuable information for the understanding and future trajectory of macroeconomic fundamentals, similar to structurally-estimated shocks. To do so, we perform an exercise similar to Lumbanraja et al. (2025), who take the "surprise" in the Oil Net Supply series, given by the residuals of an AR(1) regression, and compare the impulse responses of a number of relevant variables to this surprise, relative to two

well-known frameworks. The first is the oil supply shock of Baumeister and Hamilton (2019), identified using sign restrictions with informative priors drawn from the literature. The second is the news shock about future oil supply of Känzig (2021), who uses a high-frequency identification scheme exploiting reactions of futures prices to OPEC announcements, and uses them as an external instrument in a standard oil VAR.[15] These two shock series are estimated using two different structural methods and represent the frontier for oil-price modeling. Establishing this benchmark is important against the backdrop of a lack of proxies for the drivers of other commodities (besides oil), which is exactly the gap that this paper aims to fill. Different from Lumbanraja et al. (2025), we clean from demand in the local projection exercise and plot responses to the residuals from an autoregressive process for the supply-unigram indicator.

We use the six variables used in the exercise of Känzig (2021): the real WTI spot price, global oil production, inventories (proxied by OECD inventories), global industrial production, US industrial production and US CPI. Following Känzig (2021), the shocks are normalized to a shock that raises the real oil price by 10% on impact. Throughout this set of exercises (as well as in the following section), we plot, as is standard, the coefficients of the Local Projections across horizons, together with 90% and 95% confidence intervals. However, as is well known, the coefficients in Local Projections are highly serially correlated, which can result in separately insignificant coefficients that are jointly significant. We employ the fix of Jordà (2023), and plot the inverted statistic of the test of joint significance at each step h of total horizon H, which is given by $\pm\sigma_h\sqrt{d(H,\alpha)/H}$, where $d(H,\alpha)$ is the critical value of the joint null for significance level α (with an asymptotic χ^2 distribution), and σ_h is the standard deviation of the estimated coefficient at horizon h. These significance bands are useful because while each coefficient may not be statistically significant, they can still be jointly significant, which is the economic effect of interest (i.e. whether a treatment had an effect on the outcome.

Figure 8 and Table 5 present compelling evidence that the textual indices correlate and offer comparable predictions to those of structural methods popular in the literature. In particular, Table 5 presents the correlations of the Net Supply and Net Demand Oil Indices of Lumbanraja et al. (2025) with those of Baumeister and Hamilton (2019) and Känzig (2021). The results indicate that a surprise of the simplest Net Supply indicator for oil (computed as the residual of an autoregressive process of order one) has

[15]The authors of these papers provide updated versions of their shocks on their respective websites. We use shocks up to March 2023 from both papers (and not the vintages provided in the replication packages).

a correlation of 31% with the oil shock series of Baumeister and Hamilton (2019) and 35% with Känzig (2021). Their correlation is 37% for the same sample. The surprise of the Net Demand Oil Index has a correlation of 41% with the sum of the three demand components of Baumeister and Hamilton (2019). All correlations are highly significant (over 99.9% confidence level).[16]

The study of the impulse responses of the local projections (Figure 8) further reveals that a negative surprise of the Net Supply Index of Lumbanraja et al. (2025) predicts a substantial and immediate increase in oil prices. Global oil production decreases over the subsequent 24 months. There is also a small but statistically significant short-term inventory depletion. In addition, a negative supply shock leads to a marked and sustained decline in US and global industrial production, highlighting global vulnerabilities to higher oil prices. U.S. consumer prices rise significantly, on impact, and continue to increase for approximately one year.

Interestingly, despite the fact that the textual methodology does not include any data on prices or quantities, the textual indices are still able to capture similar dynamics to methods that do incorporate this information. The results across the three methods are broadly consistent, both qualitatively and in terms of statistical significance.

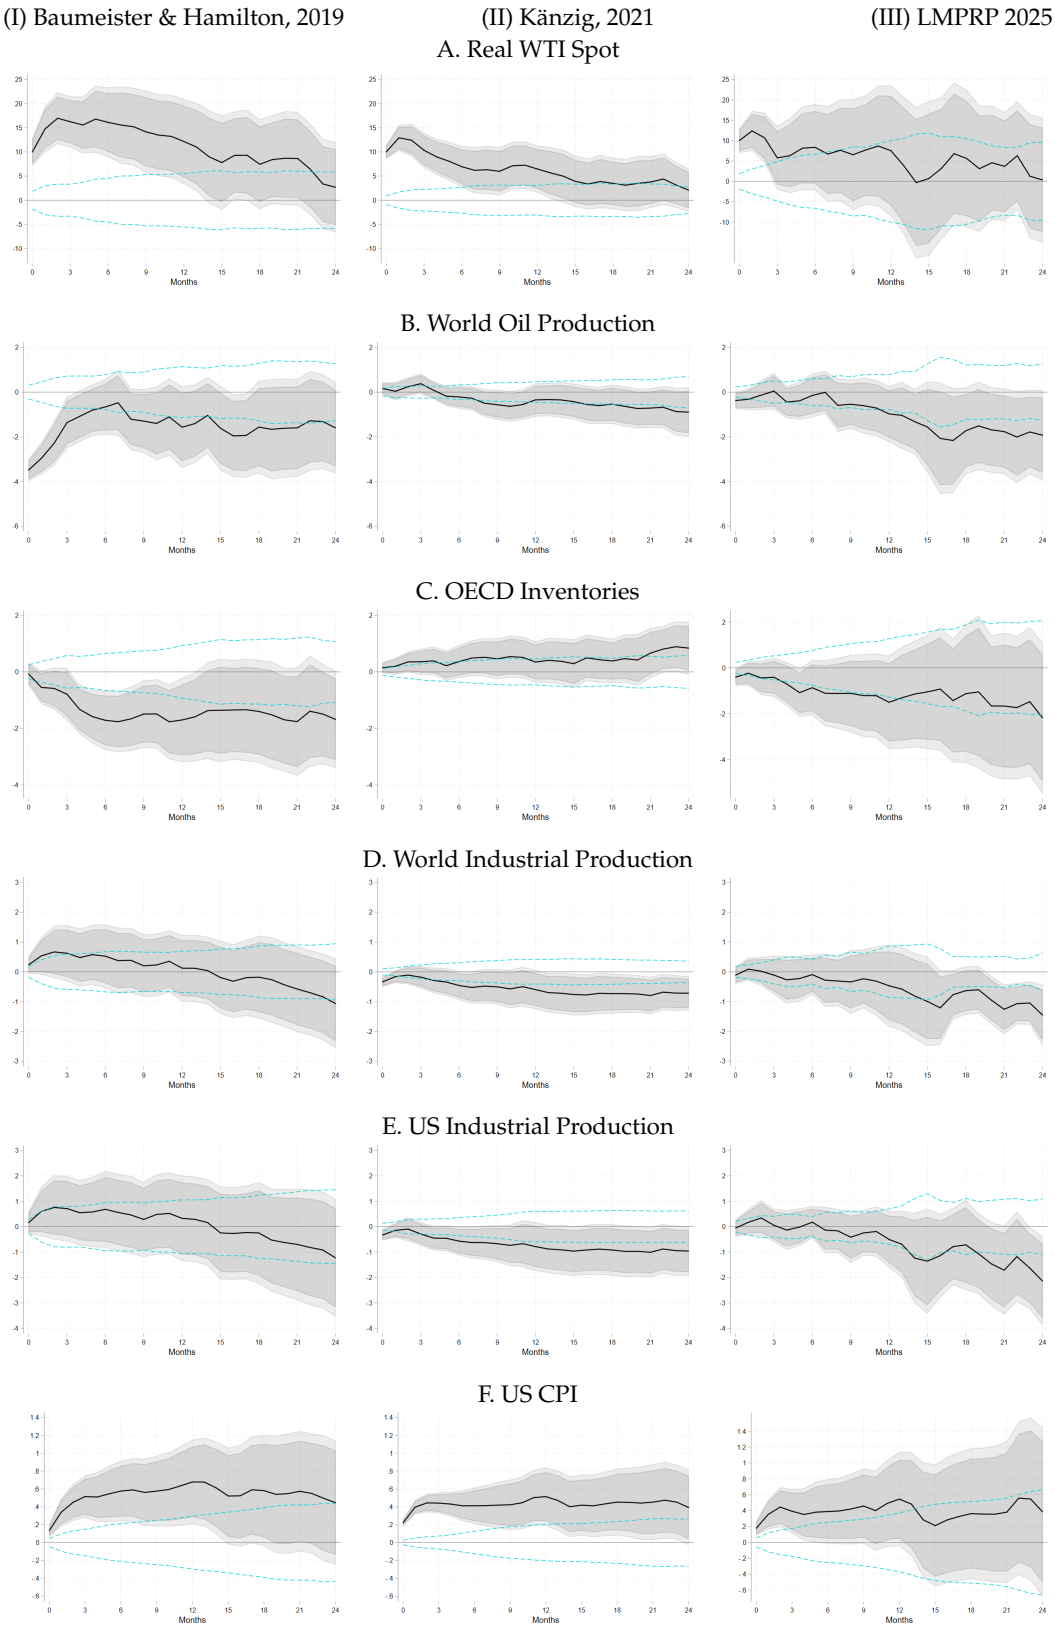
TABLE 5. Correlations

	Supply			Demand	
	LMPRP 2025	Kanzig 2021	BH 2019	LMPRP 2025	BH 2019
LMPRP 2025	1.0000			1.0000	
Kanzig 2021	0.3534	1.0000			
BH 2019	0.3138	0.3735	1.0000	0.4109	1.0000

Note: The table shows correlations of the structural shocks of Baumeister and Hamilton (2019) (BH 2019) and Känzig (2021) with a surprise series in the indices of Lumbanraja et al. (2025) (LMPRP 2025), given by the residuals of a regression of the index on its own lag. We use the Net Oil Supply and Demand indices, respectively.

[16]Each local projection includes 24 own-dependent variable lags. Figure A3 in the Appendix repeats the analysis with local projections from specifications that further control for the other variables. In particular, following Känzig (2021), we control for 12 lags of all variables under consideration for US IP and CPI; for the oil market variables and world IP, we control only for the oil market and for world IP.

FIGURE 8. Local Projections Impulse Responses



Note: This figure from Lumbanraja et al. (2025) shows the impulse responses to an oil supply shock (mean and 90% and 95% CI) for the Real Spot Price of Oil, World Oil Production, OECD Oil Inventories, World Industrial Production, US Industrial Production, and US CPI. The teal dashed lines show significance bands for the test that all coefficients of interest are jointly zero, as in Jordà (2023). The shocks are the oil supply shocks of Baumeister and Hamilton (2019) and Känzig (2021) and a "surprise" series, calculated from our Net Supply Crude Oil Index as the residual of an AR(1), controlling for lag Net Demand (LMPRP 2025). The shock is normalized to a shock producing an impact increase of 10% in the real price of oil.

Notably, all three exercises match quite closely the effects on the two price series. Our Net Supply surprise yields overall less precise and more erratic impulse responses in further horizons, but this is a well-known feature of local projections relative to Vector Autoregressions (VAR) used to identify the shocks in the other two papers.[17] Overall, the impulse responses of the Net Supply Oil Index compare well with other state-of-the-art measures of the literature that use structural methods.

6. The Differential Impact of Commodity Supply and Demand Measures on Inflation

In this section, we study the responses of prices to commodity-market events picked by the textual-based measures. Although the commodity news indicators are not shocks–but rather proxies of shocks–and one cannot readily separate the commodity supply channel from demand or from other inflation drivers, the leveling exercise of the previous section gives us some comfort that the constructed "commodity-news shocks" series behaves in a highly similar way to other, well-established commodity shocks series.

Armed with this observation and the hypothesis that our news-derived proxies are close trackers of real shock behavior, we employ the Local Projections framework of Jordà (2005), which involves regressing successive forward values of the dependent variable on its lags, our indices, and relevant controls. The controls are the log change of industrial production, S&P 500 log returns, the federal funds rate, the 10-year minus 2-year US treasury yield, VIX, and trade-weighted US dollar log returns. We are mainly interested in disentangling the different effects of net supply and demand developments, and we hence include the net supply and demand indices for each category of indices together. We include up to an 18-month horizon, and use Newey-West standard errors.

The impulse responses have been normalized to have the same scale. The indices we use are scale free, and hence the unit standard deviation normalisation commonly used cannot say how large the shocks are, which is not very useful for studying passthrough. The normalisation is akin to the unit effect normalisation of Stock and Watson (2018). We normalize by shocks that raise the annual GSCI growth by its median absolute value (14.3%). Ideally we would normalize by the effect on inflation, but on occasion

[17] Intuitively, as detailed by Känzig (2021) and Nakamura and Steinsson (2018), countless shocks hit the economy in longer horizons. Without the additional structure imposed by Structural VARs, estimation is bound to be less precise, also because the sample for high-frequency series is much shorter than with standard macro series.

the impact effect is negative. We could have chosen a different horizon, but in this case we would assume the same effect for both supply and demand. The dependent variable is the natural logarithm of the respective price level (headline CPI or one of its components), and so the coefficients give 100*log points change from the baseline.

A serious issue one needs to address is the non-stationarity induced by the COVID-19 shock, which is particularly severe in our framework. Lenza and Primiceri (2022) argue that, for the calculation of impulse responses, dropping the pandemic is typically sufficient. However, we also want to examine how well our indices perform during this volatile episode, especially because the pandemic coincided with a generational spike in inflation. As such, we take a simple version of the Lenza and Primiceri (2022), whereby we replace our indices with their 2019 mean for the January-April 2020 period, and include a COVID-19 dummy in the regressions, as suggested by Ng (2021). We provide further details on these methods in Section 4.

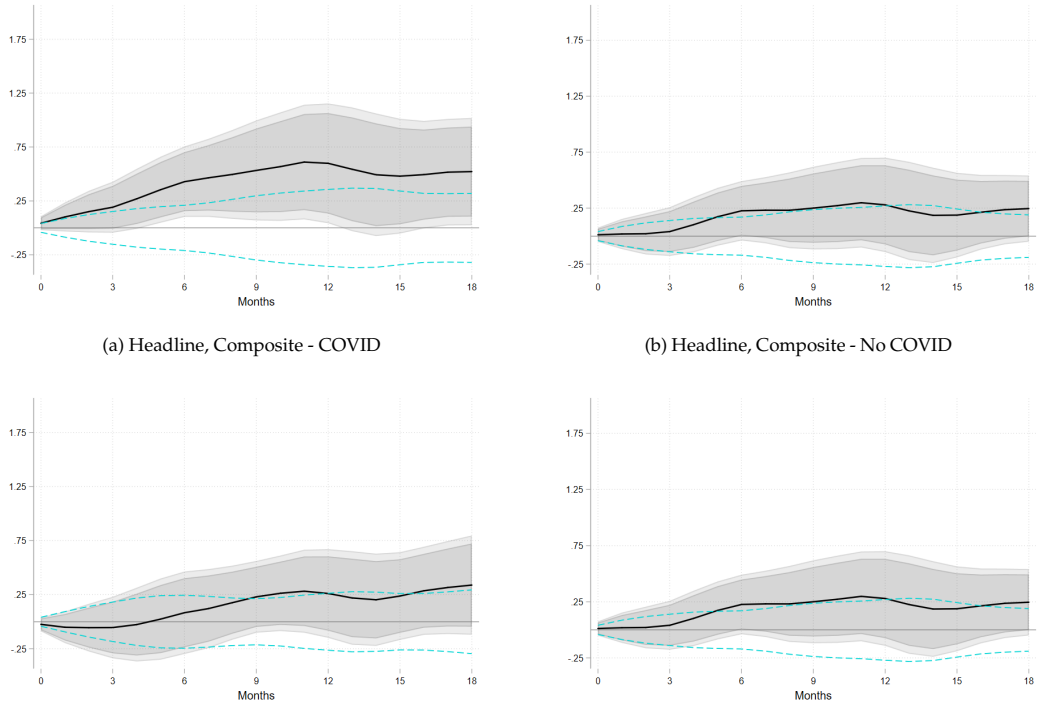
The Impact of the COVID-19 Pandemic on Inflation. We have already underlined the importance of pinning down the origins of commodity price movements for the modeling of inflation dynamics. The COVID-19 pandemic has posed additional challenges for the understanding of the commodity price-inflation link by creating large and persistent supply shocks with grave and long-lasting inflationary consequences. The inflation discussion this time moved away from crude oil considerations and rather focused on a broader commodity basket that was affected by global supply-chain disruptions.

Figure 9 shows the impact of including the data of the first year of the COVID-19 pandemic in our sample. Panel A plots the Local Projection of CPI to a standardized unit increase in our composite net supply index between 2001 and 2023. This sample includes the Ukraine war, but results are very similar even before the large increase in global energy prices starting in early 2022. The panel to the right (Panel B), plots the same response, but calculated only for the pre-COVID sample (until December 2019). The effect of supply developments is much smaller in the pre-COVID sample, and is hardly statistically significant at the 18-month horizon.

Instead, the result reflects the various supply-related problems created due to the COVID pandemic, such as lockdowns, supply chain disruptions, the rising cost of global shipping, and any other hindrance to the supply of commodities. In particular, 2021 was a year marked by persistent low inventories for a number of key metals, with copper, aluminum, nickel and zinc markets registering a deficit due to both inventory and distribution problems. This is consistent with the inspection of the same set of impulse responses coming only from the oil market (Panels C and D). Indeed,

the responses of CPI to a standardized unit increase in our net supply of oil index are similar in terms of statistical significance with or without the COVID sample.

FIGURE 9. CPI Inflation Response to Supply Developments - COVID



(a) Headline, Composite - COVID

(b) Headline, Composite - No COVID

(c) Headline, Oil - COVID

(d) Headline, Oil - No COVID

Note: Log CPI response following a one unit increase in contemporaneous composite net supply up to an 18-month horizon (90% and 95% CI). The response is normalized to a shock that raises the GSCI by 10% in 12 months. The teal dashed lines show significance bands for the test that all coefficients of interest are jointly zero, as in Jordà (2023). Controls: industrial production log change, S&P 500 log returns, FFR, 10-year minus 2-year US treasury, VIX, trade-weighted US dollar log returns. Each specification also includes a constant and 13 lags of log CPI. COVID corrections using a COVID dummy and rescaling of the indices for January-April 2020. Sample: March 2001-June 2023.

Commodity News Indices and Inflation. We now turn to the effects of commodity developments (both supply and demand) on aggregate price growth. While some of the US literature tends to focus on PCE inflation, because this is the measure targeted by the Fed, in this section, we focus only on CPI inflation. We do so because most other countries rely on CPI inflation, which is computed from consumer surveys, while PCE is measured from business surveys. Since our framework is designed to be used in any jurisdiction, CPI is the more natural choice. However, we report results for both CPI and PCE in the next section on forecasting.

The first row (Panels A and B) of Figure 10 shows the response of annual CPI inflation to supply and demand developments, using the composite indices. The response to demand is almost twice as high, and equally persistent. We argue that this is due to

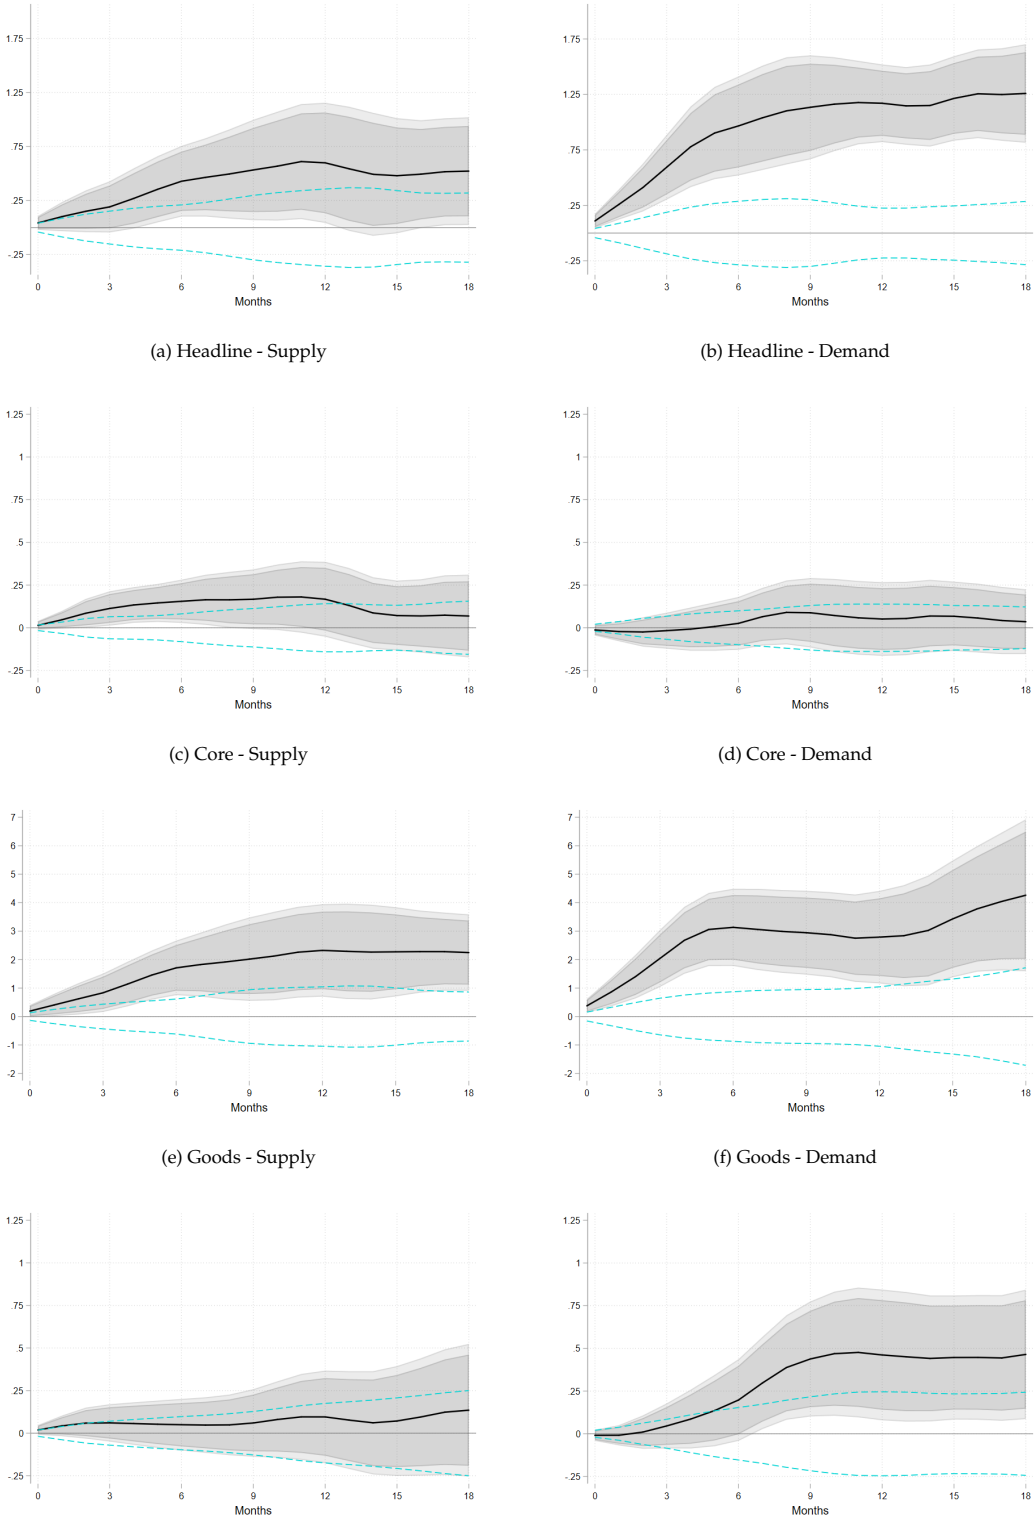
the fact that our composite demand measure captures essentially an aggregate demand component, and so signals a broad-based change in demand. This is evidenced by the fact that the demand measure is highly correlated with the business cycle driver.

The second row (Panels C and D) shows the response of core CPI inflation to the same disturbance. As expected, the response is much smaller, because most commodities have been removed from the price index. In this case, the response to supply developments is somewhat larger, and significant for a few months (according to both the standard confidence intervals and the joint test). In addition to direct effects from other commodities which are still present in the core price index and are captured in our composite indices (e.g. metals), the response also likely captures second-round effects.

Panels E and H further show the responses of goods and services inflation. Services inflation is structurally higher than goods inflation, as it benefits less by productivity growth, but is much less volatile. The response of both is significant and persistent for demand developments, but naturally more pronounced for goods. Goods inflation also responds to supply developments, about half as strongly as to demand developments, while services inflation has a very weak response to supply developments. It should be noted that energy is classified under services in CPI, and so services inflation differs from core inflation. At the same time, a larger chunk of goods CPI is composed of commodities than services CPI, which can also account for the different responses.

We then turn to considering the response of goods CPI, broken down into food and goods less food (GLF). The latter accounts for around two thirds of the total goods CPI basket, and is composed of roughly equal parts of durables and non-durables less food. GLF inflation is more volatile than food inflation, and so the magnitude of the response to demand developments is substantially larger and highly persistent (Figure 11). The response of food inflation is more gradual, peaking at 9 months, and eventually disappears. This disparate behavior can account for the response of headline CPI inflation to demand developments, which rises fast, plateaus at 8 months, but does not revert even at 18 months (Figure 10b). For both categories, there are positive effects from supply as well; for food CPI, in particular, the response to supply developments is slower than for demand but is persistent, and the overall effect deeper.

FIGURE 10. CPI Inflation Response to Composite Developments



(a) Headline - Supply

(b) Headline - Demand

(c) Core - Supply

(d) Core - Demand

(e) Goods - Supply

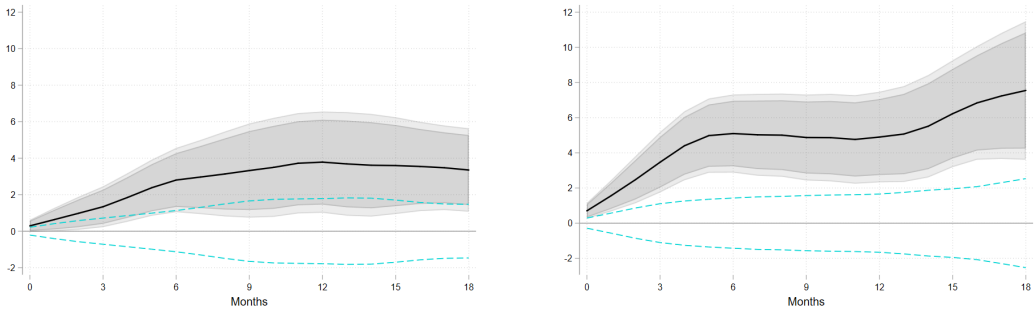
(f) Goods - Demand

(g) Services - Supply

(h) Services - Demand

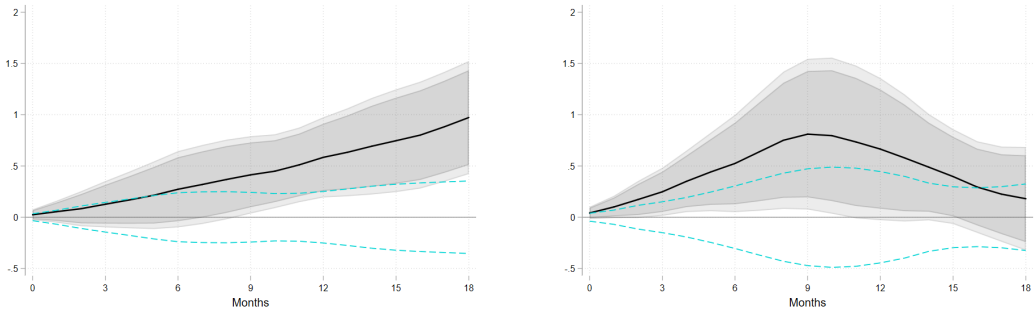
Note: Log CPI response following a one unit increase in contemporaneous composite net supply and net demand, up to an 18-month horizon (90% and 95% CI). The response is normalized to a shock that raises the GSCI by 10% in 12 months. The teal dashed lines show significance bands for the test that all coefficients of interest are jointly zero, as in Jordà (2023). Controls: industrial production log change, S&P 500 log returns, FFR, 10-year minus 2-year US treasury, VIX, trade-weighted US dollar log returns. Each specification also includes a constant and 13 lags of log CPI. COVID corrections using a COVID dummy and rescaling of the indices for January-April 2020. Sample: March 2001-June 2023.

FIGURE 11. CPI Goods Inflation Response to Commodity Developments



(a) Goods less food - Supply

(b) Goods less food - Demand



(c) Food - Supply

(d) Food - Demand

Note: Log CPI response following a one unit increase in contemporaneous composite net supply and net demand, up to an 18-month horizon (90% and 95% CI). The response is normalized to a shock that raises the GSCI by 10% in 12 months. The teal dashed lines show significance bands for the test that all coefficients of interest are jointly zero, as in Jordà (2023). Controls: industrial production log change, S&P 500 log returns, FFR, 10-year minus 2-year US treasury, VIX, trade-weighted US dollar log returns. Each specification also includes a constant and 13 lags of log CPI. COVID corrections using a COVID dummy and rescaling of the indices for January-April 2020. Sample: March 2001-June 2023.

Commodity News Indices and Inflation Expectations. Overall our textual indices provide intuitively reasonable results, showing that our news-reading method manages to extract information not otherwise available in standard macro controls. However, it is possible that we are just picking up information already available to economic agents, which is, however, not embodied yet in time t controls. One way to test this proposition is to see whether our indices can also "predict" agents' expectations; if they do, then we are simply picking up knowledge incorporated into people's forecasts. As such, we repeat the local projection exercise, but instead of regressing future prices on the textual indices, we instead regress various measures of inflation expectations on the textual indices.

We use three different measures of expectations. The first is a market-based measure from inflation swaps, as computed by Haubrich, Pennacchi, and Ritchken (2012),

and published by the Cleveland Fed (CLEV). It incorporates expectations for CPI inflation implicit in bond market valuations. The second comes from the Survey of Professional Forecasters (SPF), a long survey conducted by the Philadelphia Fed, and reflects the views of experts. SPF data is given at a quarterly basis, and we use interpolated monthly measures of annual growth rates, for both CPI and PCE inflation. The third measure comes from the University of Michigan Survey of Consumers (MICH), and is a household measure of inflation expectations. As such, it is not tied to any specific price basket, and is considered to be a less sophisticated measure of inflation expectations.

The results of this exercise are shown in Table 6. Because expectations data are mostly available on an annual basis, we do not conduct a full impulse response analysis, and instead only consider expected inflation over the next 12 months. As such, the baseline comparison is relative to impulse responses of annual price growth rates, not levels, as in the previous exercise. These are shown in Column 1 of Panel A, for actual CPI inflation 12 months ahead, in log growth form. For this exercise, we do not normalize the coefficients, as we only compare the same variables of interest across models. As before, the response to demand is significant 12 months ahead; for supply, the coefficient has a p-value for 0.133, and, as we saw before, a test of joint-significance will reject the null for $h = 12$. Column 2 shows results for the SPF measure. Both coefficients are at least an order of magnitude smaller, and not significant. Column 3 shows results for the SPF measure 12 months ahead, i.e., t+24 inflation expectations of professional forecasters formed at time t+12. Interestingly, the response is much larger in this case, suggesting a delayed response of professional forecasters' inflation expectations to the information picked up by our indices. Columns 4 and 5 repeat these exercises for the market-based measure of inflation expectations from inflation swaps (CLEV), and the results are very similar; essentially no effect for t+12 expectations at t (column 4), and a much larger delayed response of t+24 expectations at time t+12 (column 5), though still much smaller than the baseline effect on CPI growth.

Panel B repeats this exercise for PCE inflation, which is only available in SPF, and also provides results for MICH, which is not specific to any particular price basket. The results are very much the same. For completeness, Panels C and D repeat the previous two exercises removing the high inflation sample which began in the COVID era, and where expectations were highly anchored, hence vastly underpredicting actual inflation. However, the results are very similar with the full sample. Overall, the message from Table 6 is that our indices contain information about future inflation not

TABLE 6. Commodity News Indices and Inflation Expectations

	(1)	(2)	(3)	(4)	(5)
Panel A - CPI					
	Baseline	SPF	F12 SPF	Clev	F12 Clev
NetDemand	0.885***	-0.0334	0.150**	0.0493	0.147
	(4.40)	(-0.68)	(3.18)	(0.65)	(1.96)
NetSupply	0.378	0.0115	0.0340	0.00310	0.0644
	(1.51)	(0.33)	(0.85)	(0.06)	(0.66)
N	257	266	259	266	259
Panel B - PCE and Consumers					
	Baseline	SPF	F12 SPF	Michigan	F12 MICH
NetDemand	0.314***	-0.0539	0.124**	-0.0598	0.0711
	(4.23)	(-0.94)	(3.29)	(-0.50)	(1.03)
NetSupply	0.162*	0.00796	0.0361	-0.00314	0.0563
	(2.04)	(0.19)	(0.96)	(-0.03)	(0.86)
N	257	197	202	266	257
Panel C - CPI, Until 2019					
	Baseline	SPF	F12 SPF	Clev	F12 Clev
NetDemand	0.894***	-0.0339	0.128**	0.0215	0.162
	(3.80)	(-0.65)	(2.97)	(0.30)	(1.74)
NetSupply	0.340	-0.0261	0.0112	-0.0593	0.0450
	(1.39)	(-0.82)	(0.31)	(-1.20)	(0.46)
N	224	224	224	224	224
Panel D - PCE and Consumers, Until 2019					
NetDemand	0.254**	-0.102*	0.118**	-0.0254	0.0571
	(3.19)	(-2.04)	(3.11)	(-0.16)	(0.65)
NetSupply	0.112	-0.0191	0.0123	-0.0865	0.0354
	(1.62)	(-0.46)	(0.43)	(-0.70)	(0.57)
N	224	155	167	224	224

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table reports coefficients from a one standard-deviation increase in contemporaneous composite supply and demand developments. The dependent variable in the baseline column is the annual log growth rate of CPI or PCE, 12 months ahead, as indicated. In columns 2 and 4, the dependent variable is expected inflation at t+12, as of time t, according to the indicated measure. In columns 3 and 5, the dependent variable is expected inflation at t+24, as of t+12. SPF=Survey of Professional Forecasters; Clev=market-implied expectations from Haubrich, Pennacchi, and Ritchken (2012); MICH=Michigan consumer survey. Controls: industrial production log change, S&P 500 log change, FFR, 10-year minus 2-year US treasury, VIX, trade-weighted US dollar log change. Each specification also includes a constant and 13 lags of log CPI. COVID corrections using a COVID dummy and rescaling of the indices for January-April 2020. Sample: March 2001-June 2023.

included in measures of inflation expectations derived from financial prices or inflation expectations of professional forecasters and consumers.

7. Concluding Remarks

In this paper, we employ a computer-based, narrative approach that builds on the work of Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025) for the analysis of inflation as emanating from developments in commodities. By leveraging the rich informational content of business news, we distinguish between demand- and supply-side drivers of commodity price movements and their respective impacts on future inflation, controlling for widely-used predictors. Our findings point to the dominant role of demand-side developments as inflation drivers. Nonetheless, we also document a notable increase in the significance of supply-side shocks under certain conditions, particularly in recent years. The COVID-19 pandemic, in particular, has introduced new complexities in understanding the relationship between commodity prices and inflation.

We find that our indices play a crucial role in enhancing out-of-sample forecasts of inflation by significantly reducing forecast errors. These indicators not only facilitate an understanding of the persistence of shocks generated by various commodity drivers, but they also offer substantial incremental information beyond that provided by commodity returns alone. This underscores the importance of analyzing supply and demand side factors independently, given the varying persistence they may exhibit. The text-based commodity demand and supply measures also demonstrate robust performance across different economic regimes, including periods characterized by declining, stable, and rising headline inflation. The employed framework has the additional benefit of catering to the forecasting of inflation components via the mapping of individual commodities to different inflation baskets, yielding further forecasting improvements. Furthermore, narratively dissecting the drivers of commodity supply and demand provides novel insights regarding the magnitude and persistence of the pass-through effect to inflation.

The interplay of global supply chain disruptions and substantial fiscal stimulus, compounded by the geopolitical ramifications of Russia's invasion of Ukraine, has precipitated multi-decade highs in annual inflation indicators, necessitating a rigorous analysis of the underlying inflation dynamics. There has been a lot of interest and some disagreement regarding the relative role of supply and demand factors for inflation in these recent episodes (Blanchard and Bernanke, 2023; Bańbura, Bobeica, and Martínez Hernández, 2023; Giannone and Primiceri, 2024; Bergholt and Ulvedal, 2024).

In this context, the growing literature that utilizes advancements in computational linguistics to analyze economic narratives is emerging as a promising avenue for deepening our understanding of the complexities inherent in inflationary behaviors. While additional research in this domain is essential, the implications for monetary policy could be significant. Our forecasting exercise, while intentionally straightforward and parsimonious, demonstrates that the textual indices employed herein yield considerable forecasting improvements compared to conventional autoregressive benchmarks. The timely applicability of this framework allows for its integration into existing forecasting methodologies employed by central banks, thereby enhancing the precision and responsiveness of inflation predictions.

References

Alexopoulos, Michelle and Jon Cohen. 2015. "The power of print: Uncertainty shocks, markets, and the economy." *International Review of Economics & Finance* 40 (C):8–28. URL https://ideas.repec.org/a/eee/reveco/v40y2015icp8-28.html.

Allcott, Hunt and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31 (2):211–236. URL https://ideas.repec.org/a/aea/jecper/v31y2017i2p211-36.html.

Angelico, Cristina, Juri Marcucci, Marcello Miccoli, and Filippo Quarta. 2022. "Can we measure inflation expectations using Twitter?" *Journal of Econometrics* 228 (2):259–277. URL https://www.sciencedirect.com/science/article/pii/S0304407622000227.

Antolin-Diaz, Juan and Juan F. Rubio-Ramirez. 2018. "Narrative Sign Restrictions for SVARs." *American Economic Review* 108 (10):2802–2829. URL https://ideas.repec.org/a/aea/aecrev/v108y2018i10p2802-29.html.

Baker, Scott, Nicholas Bloom, and Steven Davis. 2016. "Measuring Economic Policy Uncertainty." *The Quarterly Journal of Economics* 131 (4):1593–1636. URL https://EconPapers.repec.org/RePEc:oup:qjecon:v:131:y:2016:i:4:p:1593-1636.

Bańbura, Marta, Elena Bobeica, and Catalina Martínez Hernández. 2023. "What drives core inflation? The role of supply shocks." Working Paper Series 2875, European Central Bank. URL https://EconPapers.repec.org/RePEc:ecb:ecbwps:20232875.

Bańbura, Marta, Michele Lenza, and Joan Paredes. 2024. "Forecasting inflation in the US and in the euro area." In *Handbook of Research Methods and Applications in Macroeconomic Forecasting*, edited by Michael P. Clements and Ana Beatriz GalvÃÂ£o, Chapters, chap. 9. Edward Elgar Publishing, 218–245. URL https://ideas.repec.org/h/elg/eechap/22222_9.html.

Baumeister, Christiane and James D. Hamilton. 2019. "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks." *American Economic Review* 109 (5):1873–1910. URL https://www.aeaweb.org/articles?id=10.1257/aer.20151569.

Baumeister, Christiane and Gert Peersman. 2013. "Time-Varying Effects of Oil Supply Shocks on the US Economy." *American Economic Journal: Macroeconomics* 5 (4):1–28. URL https://ideas.repec.org/a/aea/aejmac/v5y2013i4p1-28.html.

Bergholt, Canova-F. Furlanetto F. Maffei-Faccioli N., D. and P. Ulvedal. 2024. "What Drives the Recent Surge in Inflation? The Historical Decomposition Roller Coaster." *CEPR Discussion Papers 19005* .

Bernanke, Ben S. 2008. "Outstanding issues in the analysis of inflation: a speech at the Federal Reserve Bank of Boston's 52nd Annual Economic Conference, Chatham, Massachusetts, June 9, 2008." Speech 412, Board of Governors of the Federal Reserve System (U.S.).

Blanchard, Olivier J. and Ben S. Bernanke. 2023. "What Caused the US Pandemic-Era Inflation?" NBER Working Papers 31417, National Bureau of Economic Research, Inc.

Boudoukh, Jacob, Ronen Feldman, Shimon Kogan, and Matthew Richardson. 2013. "Which News Moves Stock Prices? A Textual Analysis." NBER Working Papers 18725, National Bureau of Economic Research, Inc. URL https://EconPapers.repec.org/RePEc:nbr:nberwo:18725.

Boughton, James and William Branson. 1991. "Commodity Prices as a Leading Indicator of Inflation." In *Leading Economic Indicators*, edited by Kajal Lahiri and Geoffrey Moore. Cambridge, UK: Cambridge University Press, 305–338.

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello. 2019. "Oil price elasticities and oil price fluctuations." *Journal of Monetary Economics* 103 (C):1–20. URL https://ideas.repec.org/a/eee/moneco/v103y2019icp1-20.html.

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

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel. 2020. "Hedging Climate Change News." *Review of Financial Studies* 33 (3):1184–1216.

Faust, Jon and Jonathan Wright. 2013. "Forecasting Inflation." chap. Chapter 1. Elsevier, 2–56. URL https://EconPapers.repec.org/RePEc:eee:ecofch:2-2.

Garratt, Anthony and Ivan Petrella. 2022. "Commodity prices and inflation risk." *Journal of Applied Econometrics* .

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78 (1):35–71. URL https://ideas.repec.org/a/ecm/emetrp/v78y2010i1p35-71.html.

Giannone, Domenico and Giorgio Primiceri. 2024. "The Drivers of Post-Pandemic Inflation." *NBER Working Papers 32859* .

Gospodinov, Nikolay and Serena Ng. 2013a. "Commodity Prices, Convenience Yields, and Inflation." *The Review of Economics and Statistics* 95 (1):206–219.

———. 2013b. "Commodity Prices, Convenience Yields, and Inflation." *The Review of Economics and Statistics* 95 (1):206–219. URL https://ideas.repec.org/a/tpr/restat/v95y2013i1p206-219.html.

Güntner, Jochen, Magnus Reif, and Maik Wolters. 2024. "Sudden stop: Supply and demand shocks in the German natural gas market." *Journal of Applied Econometrics* 39.

Hassan, Tarek, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2023a. "The Global Impact of Brexit Uncertainty." *The Journal of Finance, forthcoming* .

Hassan, Tarek, Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun. 2023b. "Sources and Transmission of Country Risk." *The Review of Economic Studies, forthcoming* .

Hassan, Tarek A, Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. "Firm-Level Political Risk: Measurement and Effects." *The Quarterly Journal of Economics* 134 (4):2135–2202. URL https://ideas.repec.org/a/oup/qjecon/v134y2019i4p2135-2202..html.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken. 2012. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." *The Review of Financial Studies* 25 (5):1588–1629.

Hoberg, Gerard and Gordon Phillips. 2010. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *Review of Financial Studies* 23 (10):3773–3811. URL https://EconPapers.repec.org/RePEc:oup:rfinst:v:23:y:2010:i:10:p:3773-3811.

Hong, Yongmiao, Fuwei Jiang, Lingchao Meng, and Bowen Xue. 2024. "Forecasting Inflation Using Economic Narratives." *Journal of Business Economic Statistics* 43:1–27.

———. 2025. "Forecasting Inflation Using Economic Narratives." *Journal of Business & Economic Statistics* 43 (1):216–231.

Hooker, Mark A. 2002. "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime." *Journal of Money, Credit and Banking* 34 (2):540–61. URL https://EconPapers.repec.org/RePEc:mcb:jmoncb:v:34:y:2002:i:2:p:540-61.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (1):161–182.

———. 2023. "Local Projections for Applied Economics." *Annual Review of Economics* 15 (1):607–631.

Joseph, Andreas, Eleni Kalamara, George Kapetanios, Galina Potjagailo, and Chiranjit Chakraborty. 2021. "Forecasting UK inflation bottom up." Bank of England working papers 915, Bank of England.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review* 111 (4):1092–1125. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190964.

Kilian, Lutz. 2008. "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *The Review of Economics and Statistics* 90 (2):216–240. URL https://ideas.repec.org/a/tpr/restat/v90y2008i2p216-240.html.

———. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3):1053–1069. URL https://ideas.repec.org/a/aea/aecrev/v99y2009i3p1053-69.html.

Kilian, Lutz and Daniel P. Murphy. 2012. "Why Agnostic Sign Restrictions Are Not Enough: Understanding The Dynamics Of Oil Market Var Models." *Journal of the European Economic Association* 10 (5):1166–1188. URL https://ideas.repec.org/a/bla/jeurec/v10y2012i5p1166-1188.html.

Kilian, Lutz, Alessandro Rebucci, and Nikola Spatafora. 2009. "Oil shocks and external balances." *Journal of International Economics* 77 (2):181–194. URL https://ideas.repec.org/a/eee/inecon/v77y2009i2p181-194.html.

Kilian, Lutz and Robert J. Vigfusson. 2017. "The Role of Oil Price Shocks in Causing U.S. Recessions." *Journal of Money, Credit and Banking* 49 (8):1747–1776. URL https://ideas.repec.org/a/wly/jmoncb/v49y2017i8p1747-1776.html.

Lenza, Michele and Giorgio E. Primiceri. 2022. "How to estimate a VAR after March 2020." *Journal of Applied Econometrics* :688–699.

Lippi, Francesco and Andrea Nobili. 2012. "Oil And The Macroeconomy: A Quantitative Structural Analysis." *Journal of the European Economic Association* 10 (5):1059–1083. URL https://ideas.repec.org/a/bla/jeurec/v10y2012i5p1059-1083.html.

Loughran, Tim, Bill McDonald, and Ioannis Pragidis. 2019. "Assimilation of oil news into prices." *International Review of Financial Analysis* 63 (C):105–118. URL https://EconPapers.repec.org/RePEc:eee:finana:v:63:y:2019:i:c:p:105-118.

Lumbanraja, Alvin, Sarah Mouabbi, Eugenia Passari, and Adrien Rousset Planat. 2025. "The Origins of Commodity Price Fluctuations." Tech. rep.

Nakamura, Emi and Jon Steinsson. 2018. "High-frequency identification of monetary non-neutrality: The information effect." *The Quarterly Journal of Economics* 3 (133):1283â1330.

Ng, Serena. 2021. "Modeling Macroeconomic Variations after Covid-19." Working Paper 29060, National Bureau of Economic Research.

Romer, Christina D. and David H. Romer. 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review* 100 (3):763–801. URL https://ideas.repec.org/a/aea/aecrev/v100y2010i3p763-801.html.

Shapiro, Adam Hale. 2020. "A Simple Framework to Monitor Inflation." Working paper series, Federal Reserve Bank of San Francisco.

Stock, James H. and Mark W. Watson. 2003. "Forecasting Output and Inflation: The Role of Asset Prices." *Journal of Economic Literature* 41 (3):788–829.

———. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *Economic Journal* 128 (610):917–948.

Tetlock, Paul C. 2007. "Giving Content to Investor Sentiment: The Role of Media in the Stock Market." *Journal of Finance* 62 (3):1139–1168. URL https://ideas.repec.org/a/bla/jfinan/v62y2007i3p1139-1168.html.

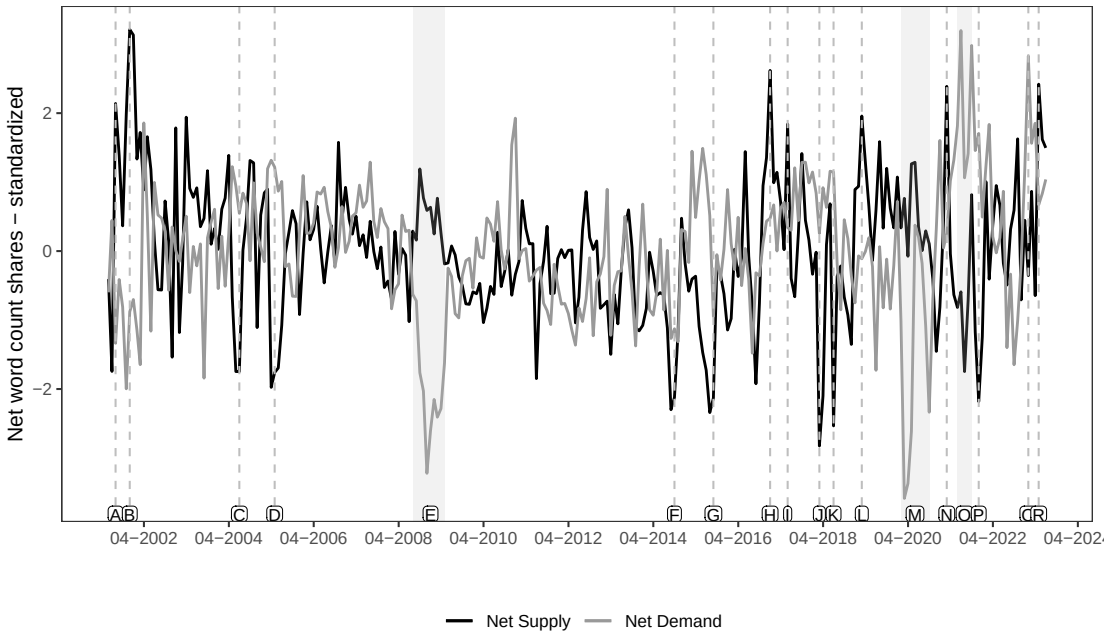
Winne, Jasmien and Gert Peersman. 2016. "Macroeconomic Effects of Disruptions in Global Food Commodity Markets: Evidence for the United States." *Brookings Papers on Economic Activity* 2016:183–286.

Wu, Tao and Michele Cavallo. 2012. "Measuring Oil-Price Shocks Using Market-Based Information." IMF Working Papers 12/19, International Monetary Fund. URL https://ideas.repec.org/p/imf/imfwpa/12-19.html.

Zhou, Xiaoqing. 2019. "Refining the Workhorse Oil Market Model." Working Papers 1910, Federal Reserve Bank of Dallas. URL https://ideas.repec.org/p/fip/feddwp/1910.html.

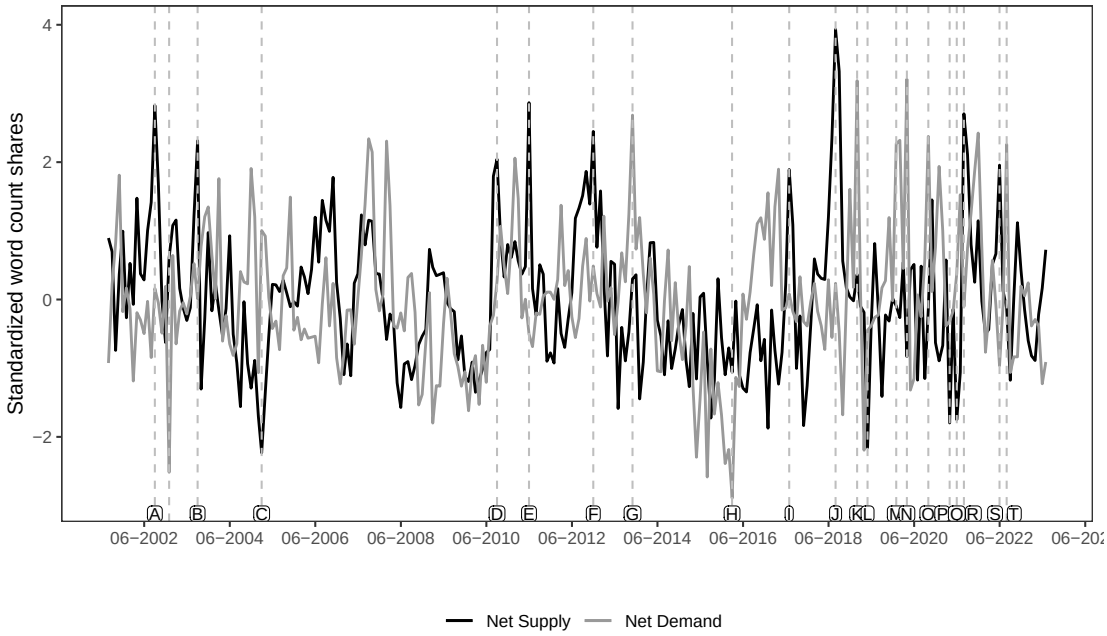
APPENDIX A. ADDITIONAL FIGURES

FIGURE A1. Crude Oil Index: Standardized Net supply and Standardized Net Demand Share of Total Number of Words per Month



[A] OPEC cuts output to lift oil prices following US 2001 recession. [B] 118th OPEC Meeting, production cut following US economic slowdown & 9/11. [C] OPEC to raise output quotas (131st E.M.). [D] OPEC pumps at 25-year highs to cater for high demand expected from China. [E] 150th & 151st OPEC E.M.: production cut (GFC). [F] Faster than expected recovery of Libyan oil production & unaffected Iraq production: perceived glut. [G] Worries about oversupplied oil market (Chinaâs slowdown); Middle East producers pump crude at record levels. [H] OPEC 171st Meeting: OPEC & non-OPEC to cut output for first time in 8 years. [I] OPEC+ extended cuts in oil output to battle global glut. [J] US record oil production: oversupply worries. [K] OPEC to raise oil production amid calls from top consumers to cool prices & support world economy. [L] OPEC production cuts & US sanctions on Iranian and Venezuelan crude. [M] COVID-19. [N] Winter storm in Texas causes US crude production drop. [O] Global economy and oil demand recover faster than expected following vaccination. [P] The US announces release of millions of barrels from strategic reserves in coordination with China, India, South Korea, Japan & Britain, to cool prices. [Q] Subsiding recession risks in major economies & reopening of China's economy boost demand. [R] Surprise decision by OPEC+ to cut output (48th JMMC).
Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

FIGURE A2. Wheat Index: Standardized Net supply and Standardized Net Demand Share of Total Number of Words per Month



[A] Severe drought in Canada cuts wheat production by more than 25%. Worrisome dry weather in Australia & US Northern Plains. [B] Hot & dry weather trims European output & Canada's crop. [C] Ample world wheat supplies amid bumper crops. [D] Due to heatwave Russia bans exports of wheat; Ukraine imposes export quotas. [E] Dry weather stresses crops in Plains; wet weather slows wheat seedings in northern Plains. [F] Weather sparks worries about world supply. [G] Prospects of strong demand from China. [H] Sluggish demand for US supplies due to China's policy shift that favors other grains. [I] Dry weather in US northern Plains threatens crops. [J] Severe drought in Europe cuts harvests. [K] Resuming trade talks between China & the US boosts higher wheat demand. [L] Abundant world supplies due to crop-friendly weather in key areas. [M] China & the US reach phase one trade deal. [N] Covid-19 cause panic shoppers around the world to stock up on wheat-based items. [O] Strong demand from China following trade deal with the US. [P] Improved US wheat condition ratings. [Q] Above-average global yield expectations pressures wheat markets. [R] Adverse weather conditions & reduced expectations for Russia's harvest. [S] Easing fears of global recession & renewed import demand. [T] US winter crop ratings fall amid droughts; hot weather curbs India's production.

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

FIGURE A3. Local Projections Impulse Responses, with controls



Note: This figure from Lumbanraja et al. (2025) shows the impulse responses to an oil supply shock (mean and 90% and 95% CI) for the Real Spot Price of Oil, World Oil Production, OECD Oil Inventories, World Industrial Production, US Industrial Production, and US CPI. The teal dashed lines show significance bands for the test that all coefficients of interest are jointly zero, as in Jordà (2023). The shocks are the oil supply shocks of Baumeister and Hamilton (2019) and Känzig (2021) and a "surprise" series, calculated from our Net Supply Crude Oil Index as the residual of an AR(1), controlling for lag Net Demand. The shock is normalized to a shock producing an impact increase of 10% in the real price of oil. The oil market variables and world IP include 12 lags of each other (and their own) as controls, while US IP and CPI include the lags of all variables.

APPENDIX B. DICTIONARIES

In this Appendix we provide the list of supply and demand words used for the creation of the "global" commodity index and the indices for crude oil, natural gas an wheat. We also present the standard list of increase and decrease words which is common across commodities.

Table B1 to Table B3 present the list of roots of words we use to search for supply and demand related expressions for the construction of the "global" commodity index as well as the indices that aim to capture individual commodities. Our results our robust to using both words and lemmas.

Table B4 and Table B5 present the list of roots of increase and decrease words.

TABLE B1. Supply and demand words: Global Commodity Index

Supply		
suppl*	produc*	output

Demand			
demand*	consum*	buy*	purchas*

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

TABLE B2. Supply and demand words: Crude Oil

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

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

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

TABLE B3. Supply and demand words: Wheat

Table B4: Supply and demand words
Wheat

Supply

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

Demand

demand*	consum*	buy*	purchas*
flour*	feed*	miller*	

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

TABLE B4. Increase words

accru*	climb*	improve*	rais*	soar*
accumulat*	elevat*	increas*	rall*	spik*
add*	enlarg*	inflat*	reach*	spring*
advanc*	escalat*	jackup*	rebuil*	spurt*
augment*	expand*	jump*	recoup*	strengthen*
bolster*	firm*	leap*	recover*	surg*
boom*	gain*	lift*	regain*	surpass*
boost*	grow*	perk*	resurgenc*	swell*
buil*	heigh*	pickup	reviv*	up*
bullish*	high*	pop*	ris*	
buoyant*	hik*	propell*	*rocket*	

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

TABLE B5. Decrease words

abat*	dent*	evaporat*	pullback*	slump*
bearish*	depress*	fad*	reced*	small*
below	deteriorat*	fall*	reduc*	squeez*
collaps*	diminish*	falter*	restrict*	stumbl*
compress*	dip*	halt*	retre*	sub*
contract*	disappear*	landslid*	sink*	suppress*
crash*	disappoint*	less*	short*	suspend*
crimp*	disrupt*	los*	shrink*	tank*
crush*	div*	low*	shut*	tap*
curtail*	down*	meltdown*	slack*	tight*
cut*	drawdown*	nosediv*	slash*	
dampe*	drop*	outage*	slid*	
declin*	dwindl*	plummet*	slip*	
decreas*	ebb*	plung*	slow*	

Source: Lumbanraja, Mouabbi, Passari, and Rousset Planat (2025)

APPENDIX C. ADDITIONAL TABLES

TABLE C1. Forecast performance (RMSE), Core CPI inflation

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	0.327	1.260	1.121	1.109	1.374	1.136	1.420
	h=6	0.527	1.086	1.050	1.025	1.133	1.021	1.162
	h=3	0.574	1.046	1.024	1.012	1.097	1.062	1.155
Controls, Autoregressive								
	h=12	0.577	1.196	1.065	1.111	1.190	1.162	1.286
	h=6	0.634	1.094	1.024	1.057	1.163	1.093	1.173
	h=3	0.613	1.073	1.030	1.041	1.083	1.121	1.153
Panel B: with commodities								
Autoregressive								
	h=12	0.333	1.260	1.121	1.098	1.367	1.136	1.426
	h=6	0.526	1.091	1.050	1.030	1.163	1.025	1.198
	h=3	0.572	1.045	1.024	1.012	1.106	1.065	1.160
Controls, Autoregressive								
	h=12	0.585	1.194	1.082	1.106	1.178	1.177	1.278
	h=6	0.635	1.095	1.025	1.063	1.164	1.095	1.162
	h=3	0.615	1.071	1.034	1.041	1.082	1.125	1.147
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.464	1.015	1.004	1.009	1.073	1.035	1.074
	h=6	1.397	1.014	0.999	1.016	1.047	1.005	1.058
	h=3	1.640	1.007	0.997	1.008	1.021	1.007	1.033
Controls, Autoregressive								
	h=12	1.592	1.012	1.012	0.999	1.027	1.040	1.013
	h=6	1.481	1.005	1.009	0.999	1.033	1.038	1.045
	h=3	1.661	1.010	1.006	1.005	1.025	1.030	1.011
Panel B: with commodities								
Autoregressive								
	h=12	1.458	1.022	1.009	1.012	1.073	1.039	1.077
	h=6	1.395	1.010	1.001	1.010	1.043	1.007	1.056
	h=3	1.623	1.003	1.001	1.002	1.013	1.012	1.026
Controls, Autoregressive								
	h=12	1.591	1.010	1.013	0.997	1.023	1.041	1.005
	h=6	1.471	1.014	1.012	1.008	1.039	1.041	1.049
	h=3	1.624	1.026	1.010	1.018	1.032	1.035	1.031
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of core CPI inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

TABLE C2. Forecast performance (RMSE), Core PCE inflation

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
			PRE-COVID SAMPLE					
Panel A: without commodities								
Autoregressive								
	h=12	0.312	0.963	1.002	0.914	0.985	1.018	0.933
	h=6	0.440	1.050	1.022	1.008	1.114	1.009	1.170
	h=3	0.520	1.090	0.977	1.052	1.115	0.994	1.192
Controls, Autoregressive								
	h=12	0.415	1.066	1.164	0.944	0.993	1.174	1.066
	h=6	0.486	1.033	1.004	1.044	1.061	1.046	1.094
	h=3	0.548	1.123	1.042	1.043	1.052	1.078	1.165
Panel B: with commodities								
Autoregressive								
	h=12	0.308	0.989	1.006	0.941	1.005	1.022	0.959
	h=6	0.440	1.056	1.030	1.019	1.127	1.011	1.175
	h=3	0.497	1.086	0.991	1.058	1.112	1.003	1.170
Controls, Autoregressive								
	h=12	0.434	1.119	1.212	0.917*	0.970	1.219	1.096
	h=6	0.486	1.046	1.010	1.052	1.068	1.049	1.107
	h=3	0.521	1.134	1.051	1.044	1.058	1.097	1.186
			FULL SAMPLE					
Panel A: without commodities								
Autoregressive								
	h=12	1.282	0.967	0.993	0.969	1.015	1.022	1.011
	h=6	1.097	0.987	0.990	0.990	1.039	1.004	1.061
	h=3	1.147	1.007	0.979*	1.016	1.045	0.980	1.051
Controls, Autoregressive								
	h=12	1.335	0.938	1.008	0.930	0.941	1.021	0.923
	h=6	1.114	0.954	0.998	0.956	0.989	1.018	1.000
	h=3	1.182	1.002	0.999	0.994	1.018	1.010	1.013
Panel B: with commodities								
Autoregressive								
	h=12	1.272	0.965	0.997	0.966	1.008	1.031	1.006
	h=6	1.081	0.987	0.995	0.991	1.040	1.011	1.060
	h=3	1.091	1.008	0.985	1.018	1.041	0.988	1.045
Controls, Autoregressive								
	h=12	1.329	0.934	1.014	0.922	0.933	1.026	0.913
	h=6	1.090	0.965	1.001	0.969	1.003	1.021	1.009
	h=3	1.114	1.021	1.004	1.011	1.032	1.020	1.038
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of core PCE inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

APPENDIX D. TURNING POINTS

TABLE D1. Forecast performance (RMSE) across different inflation regimes, CPI inflation

	h=12						h=6					
	Headline			Core			Headline			Core		
	Fall	Incr	Stable	Fall	Incr	Stable	Fall	Incr	Stable	Fall	Incr	Stable
Panel A: without commodities												
Autoregressive	1.861	4.458	1.766	1.389	1.599	1.633	2.529	3.555	2.003	0.792	1.371	1.577
model 1	0.704	0.847	0.775	1.087	1.032	0.988	0.637	0.953	0.914	1.168	1.000	1.008
model 2	1.007	1.015	0.996	1.008	1.023	1.008	1.026	1.002	1.011	1.130	0.993	1.016
model 3	0.731	0.837	0.781	1.105	1.008	0.975	0.639	0.970	0.895	1.102	1.008	1.000
model 4	0.747	0.875	0.808	1.091	1.028	1.015	0.632	1.009	0.911	1.172	1.091	1.020
model 5	1.032	1.013	1.006	1.074	1.044	1.033	1.038	1.052	1.006	1.113	1.010	1.050
model 6	0.731	0.864	0.748	0.983	1.057	0.997	0.707	1.016	0.935	1.072	1.027	1.056
Panel B: with commodities												
Autoregressive	1.895	4.455	1.771	1.366	1.607	1.632	2.550	3.572	1.898	0.833	1.374	1.547
model 1	0.735	0.852	0.759	1.060	1.030	0.992	0.698	0.935	0.961	1.156	1.002	1.027
model 2	1.020	1.011	0.997	0.991	1.029	1.010	1.044	1.015	1.022	1.100	0.995	1.029
model 3	0.752	0.841	0.766	1.091	1.003	0.978	0.692	0.947	0.934	1.116	1.009	1.018
model 4	0.752	0.877	0.792	1.076	1.020	1.015	0.693	0.980	0.950	1.179	1.091	1.035
model 5	1.050	1.011	1.010	1.062	1.049	1.034	1.068	1.055	1.020	1.118	1.012	1.062
model 6	0.766	0.860	0.733	0.949	1.054	0.991	0.760	0.992	0.990	1.053	1.022	1.073

Note: The table shows results from forecasts of headline CPI inflation across different monthly horizons, using the composite commodity indices. Each column shows the root mean squared error (RMSE) of each forecast relative to the baseline. The autoregressive specification is given by (7), with all controls. Each specification includes some of our indexes: i) model 1: net supply and net demand; ii) model 2: net supply; iii) model 3: net demand; iv) model 4: demand increase and decrease; v) model 5: supply increase and decrease; vi) model 6: supply and demand increase and decrease. Panel B also includes log monthly change in commodity prices. Fall: 3-month moving average % change in inflation negative for at least 5 consecutive months; Incr: same measure positive for at least 5 consecutive months; Stable: neither falling nor rising.

TABLE D2. Forecast performance (RMSE) across different inflation regimes, PCE inflation

	h=12						h=6					
	Headline			Core			Headline			Core		
	Fall	Incr	Stable	Fall	Incr	Stable	Fall	Incr	Stable	Fall	Incr	Stable
Panel A: without commodities												
Autoregressive	1.302	3.143	1.394	0.890	1.799	1.193	1.671	2.680	1.507	0.450	1.829	0.901
model 1	0.745	0.864	0.833	1.097	0.950	0.884	0.741	0.850	0.995	1.319	0.917	0.947
model 2	0.991	1.026	0.995	1.006	1.025	0.989	1.053	1.003	0.998	1.048	0.984	0.972
model 3	0.783	0.848	0.837	1.124	0.929	0.884	0.737	0.843	1.011	1.348	0.906	0.963
model 4	0.801	0.886	0.870	1.031	0.933	0.929	0.759	0.855	1.049	1.384	0.917	1.019
model 5	1.015	1.027	0.987	1.036	1.042	0.991	1.037	1.054	0.996	1.064	1.009	0.986
model 6	0.729	0.882	0.788	0.887	0.940	0.909	0.791	0.844	1.079	1.364	0.935	1.022
Panel B: with commodities												
Autoregressive	1.345	3.151	1.394	0.862	1.805	1.185	1.577	2.577	1.512	0.535	1.784	0.870
model 1	0.746	0.869	0.800	1.056	0.964	0.868	0.776	0.863	0.975	1.105	0.934	0.964
model 2	1.005	1.023	0.996	0.984	1.037	0.991	1.078	1.005	1.005	0.988	0.984	0.989
model 3	0.772	0.845	0.805	1.097	0.927	0.875	0.760	0.862	0.986	1.170	0.921	0.985
model 4	0.774	0.883	0.838	1.001	0.933	0.918	0.785	0.871	1.017	1.197	0.932	1.044
model 5	1.035	1.026	0.988	1.025	1.055	0.990	1.073	1.059	1.002	1.066	1.010	0.999
model 6	0.725	0.877	0.749	0.827	0.951	0.883	0.853	0.860	1.051	1.153	0.950	1.040

Note: The table shows results from forecasts of headline CPI inflation across different monthly horizons, using the composite commodity indices. Each column shows the root mean squared error (RMSE) of each forecast relative to the baseline. The autoregressive specification is given by (7), with all controls. Each specification includes some of our indexes: i) model 0: net supply and net demand; ii) model 1: net supply; iii) model 2: net demand; iv) model 4: demand increase and decrease; v) model 5: supply increase and decrease; vi) model 6: supply and demand increase and decrease. Panel B also includes log monthly change in commodity prices. Fall: 3-month moving average % change in inflation negative for at least 5 consecutive months; Incr: same measure positive for at least 5 consecutive months; Stable: neither falling nor rising.

APPENDIX E. OTHER FORECASTS

TABLE E1. Forecast performance (RMSE), CPI inflation - alternative criterion

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	0.836	0.648**	1.022	0.652***	0.712***	1.153	0.689***
	h=6	1.055	0.857	1.082	0.837	0.892	1.162	1.065
	h=3	1.320	1.016	1.125	0.951	1.009	1.209	1.210
Controls, Autoregressive								
	h=12	1.125	0.694***	1.021	0.702***	0.725**	1.042	0.690***
	h=6	1.256	0.923	1.012	0.919	0.934	1.092	1.025
	h=3	1.550	1.078	1.027	1.060	1.066	1.101	1.133
Panel B: with commodities								
Autoregressive								
	h=12	0.776	0.717**	1.035	0.707***	0.814	1.178	0.788**
	h=6	1.105	0.872	1.062	0.853*	0.887	1.099	1.012
	h=3	1.370	1.021	1.081	0.945	0.994	1.123	1.153
Controls, Autoregressive								
	h=12	1.111	0.745***	1.019	0.741***	0.773*	1.057	0.759**
	h=6	1.260	0.926	1.039	0.923	0.936	1.101	1.020
	h=3	1.485	1.039	1.038	1.005	1.015	1.084	1.099
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.449	0.787***	1.012	0.801***	0.855**	1.061	0.831**
	h=6	1.599	0.881	1.038	0.881*	0.922	1.078	0.998
	h=3	1.885	1.001	1.064	0.970	1.010	1.103	1.105
Controls, Autoregressive								
	h=12	1.853	0.748***	1.010	0.756***	0.777***	1.019	0.740***
	h=6	1.882	0.872	1.006	0.871	0.892	1.044	0.936
	h=3	2.100	1.036	1.015	1.030	1.049	1.056	1.072
Panel B: with commodities								
Autoregressive								
	h=12	1.388	0.835***	1.021	0.843***	0.916	1.080	0.890
	h=6	1.613	0.900*	1.032	0.888***	0.918	1.050	0.986
	h=3	1.823	1.012	1.046	0.972	1.006	1.062	1.092
Controls, Autoregressive								
	h=12	1.848	0.765***	1.011	0.766***	0.789***	1.030	0.759***
	h=6	1.864	0.891	1.027	0.884*	0.902	1.053	0.950
	h=3	1.938	1.040	1.024	1.017	1.035	1.054	1.087
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of CPI inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Each specification includes 3 lags of the regressors. Columns (1)-(6) show the mean absolute error (MAE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

TABLE E2. Forecast performance, PCE inflation - alternative criterion

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	0.651	0.688*	1.022	0.669**	0.733**	1.124	0.673**
	h=6	0.783	0.826	1.069	0.808	0.831	1.119	0.977
	h=3	0.959	0.970	1.083	0.935	0.964	1.175	1.134
Controls, Autoregressive								
	h=12	0.864	0.762**	1.043	0.744***	0.808*	1.027	0.694***
	h=6	0.928	0.938	0.995	0.911	0.910	1.062	0.956
	h=3	1.097	1.068	1.018	1.053	1.054	1.070	1.118
Panel B: with commodities								
Autoregressive								
	h=12	0.629	0.732*	1.026	0.724**	0.813*	1.133	0.750**
	h=6	0.822	0.833	1.049	0.823*	0.838*	1.065	0.930
	h=3	0.943	1.003	1.082	0.953	0.985	1.104	1.104
Controls, Autoregressive								
	h=12	0.849	0.791**	1.046	0.765***	0.822*	1.048	0.736***
	h=6	0.884	0.926	1.014	0.906	0.914	1.077	0.967
	h=3	0.987	1.010	1.001	0.997	1.005	1.058	1.061
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.155	0.826**	1.014	0.823***	0.876*	1.048	0.845*
	h=6	1.176	0.883	1.032	0.881*	0.913	1.056	0.984
	h=3	1.300	0.990	1.046	0.978	1.008	1.094	1.095
Controls, Autoregressive								
	h=12	1.433	0.800***	1.019	0.795***	0.832***	1.014	0.766***
	h=6	1.342	0.879	0.999	0.874	0.888	1.032	0.911
	h=3	1.400	1.062	1.013	1.058	1.073	1.046	1.102
Panel B: with commodities								
Autoregressive								
	h=12	1.152	0.845**	1.018	0.844***	0.903	1.057	0.867*
	h=6	1.195	0.883**	1.029	0.880**	0.907	1.037	0.960
	h=3	1.239	1.009	1.048	0.985	1.014	1.058	1.083
Controls, Autoregressive								
	h=12	1.434	0.799***	1.023	0.789***	0.818***	1.027	0.758***
	h=6	1.300	0.878*	1.012	0.867*	0.885	1.043	0.917
	h=3	1.285	1.020	1.002	1.016	1.031	1.036	1.056
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of PCE inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Each specification includes 3 lags of the regressors. Columns (1)-(6) show the mean absolute error (MAE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

TABLE E3. Real-time Forecast performance (RMSE), CPI inflation

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.042	0.685*	1.018	0.712**	0.837***	1.095	0.771**
	h=6	1.414	0.840	1.061	0.826*	0.893	1.114	0.974
	h=3	1.742	1.030	1.132	0.950	1.023	1.208	1.180
Controls, Autoregressive								
	h=12	1.281	0.733***	1.028	0.722***	0.791**	1.050	0.765**
	h=6	1.630	0.999	1.024	1.001	1.008	1.094	1.055
	h=3	1.994	1.024	1.024	1.006	1.011	1.080	1.084
Panel B: with commodities								
Autoregressive								
	h=12	0.986	0.733	1.020	0.752**	0.905	1.122	0.850*
	h=6	1.415	0.888*	1.066	0.860**	0.916	1.083	0.984
	h=3	1.693	1.054	1.131	0.956	1.009	1.176	1.176
Controls, Autoregressive								
	h=12	1.254	0.763***	1.015	0.749***	0.813**	1.051	0.803**
	h=6	1.618	0.981	1.033	0.984	0.993	1.102	1.053
	h=3	1.881	1.021	1.036	0.991	1.000	1.089	1.092
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	2.176	0.875*	0.994	0.887*	0.914	1.004	0.887
	h=6	2.234	0.891*	1.008	0.893**	0.929	1.028	0.954
	h=3	2.562	1.007	1.040	0.984	1.028	1.064	1.080
Controls, Autoregressive								
	h=12	2.553	0.807*	1.002	0.811*	0.842*	1.004	0.815**
	h=6	2.552	0.893	1.003	0.897	0.920	1.023	0.932
	h=3	2.813	0.994	1.002	0.991	1.019	1.022	1.026
Panel B: with commodities								
Autoregressive								
	h=12	2.147	0.897*	0.998	0.906*	0.933	1.017	0.910
	h=6	2.212	0.917	1.017	0.909*	0.937	1.026	0.965
	h=3	2.373	1.030	1.052	0.989	1.021	1.066	1.095
Controls, Autoregressive								
	h=12	2.561	0.806*	1.001	0.808*	0.836*	1.007	0.808*
	h=6	2.525	0.899	1.012	0.898	0.919	1.032	0.938
	h=3	2.577	1.016	1.014	1.001	1.025	1.037	1.057
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of headline CPI inflation across different monthly horizons, using real time data for inflation and the unemployment gap. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1,^{**} p < 0.05,^{***} p < 0.01$.

TABLE E4. Forecast performance (RMSE), Headline CPI inflation, Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
		NDNS	ND	EFFR	YC	USD	GAP	IP	VIX	SPX	SPF	CLEV	MICH
Panel A: Pre-COVID													
w/o commodities													
Autoregressive	1.034	0.675*	0.696**	0.945	1.010	0.920	1.207	1.026	1.244	1.000	1.060	0.999	1.043
Net D+S (NDNS)	0.698			1.103	1.061	0.992	1.061	1.083	1.093	1.044	1.144	1.101	1.194
Net Demand (ND)	0.719			1.148	1.043	0.961	1.045	1.075	1.078	1.035	1.142	1.102	1.205
with commodities													
Autoregressive	0.979	0.722*	0.734**	0.944	0.992	0.972**	1.186	1.084	1.238	1.056	1.030	0.984	1.054
Net D+S (NDNS)	0.707			1.090	1.067	0.980*	1.062	1.069	1.106	1.030	1.121	1.078	1.176
Net Demand (ND)	0.719			1.132	1.049	0.961***	1.050	1.075	1.100	1.036	1.128	1.089	1.201
Panel B: Full Sample													
w/o commodities													
Autoregressive	2.177	0.872*	0.884*	1.058	1.018	1.006	1.023	0.986	0.996	0.995	1.132	1.001	1.090
Net D+S (NDNS)	1.898			1.023	1.029	1.029	1.010	1.031	0.959	1.014	1.076	1.052	1.060
Net Demand (ND)	1.924			1.039	1.031	1.026	1.005	1.023	0.947	1.008	1.066	1.057	1.044
with commodities													
Autoregressive	2.149	0.895*	0.904*	1.076	1.016	1.019	1.022	0.999	0.994	1.008	1.127	1.000	1.085
Net D+S (NDNS)	1.922			1.017	1.026	1.016	1.012	1.018	0.968	1.001	1.060	1.035	1.054
Net Demand (ND)	1.942			1.037	1.028	1.017	1.008	1.014	0.959	0.999	1.056	1.040	1.041

Notes: The table shows results from forecasts of headline CPI inflation with different predictors and samples, using the composite commodity indices. Panel A shows results for the sample until December 2019, and Panel B for the full sample. The two subpanels show separate results for specifications where the baseline includes or not commodity price growth as a predictor. Column (1) shows the root mean squared error (RMSE) of the baseline forecast, for different models. For the "autoregressive" model, it shows the RMSE for the simple autoregressive model with no additional predictors, as defined in the text (7). The "NDNS" and "ND" models include Net Demand+Net Supply or Net Demand only to the autoregressive model. Columns 2-13 show RMSEs relative to the baseline models from specifications which include additional predictors to the baseline models, one at a time. Columns 2-3 also show the RMSEs of the NDNS and ND models relative to the autoregressive one, for convenience. The additional predictors are: EFFR=effective federal funds rate; YC=10 minus 2 year government bond yield; USD=trade-weighted US dollar log monthly change; GAP: unemployment gap; IP=monthly log change in industrial production; VIX: CBOE Volatility Index; SPF: 12 month CPI inflation expectations from the Survey of Professional Forecasters; SPX: log monthly change in the SP500 index; CLEV: Cleveland Fed options implied 12-month inflation expectations; MICH: Household 12-month inflation expectations from the Michigan Survey. Sample: March 2001-June 2023. $^{*}p < 0.1$,** $p < 0.05$,*** $p < 0.01$.

TABLE E5. Forecast performance (RMSE), Headline PCE inflation, Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
		NDNS	ND	EFFR	YC	USD	GAP	IP	VIX	SPX	SPF	CLEV	MICH
Panel A: Pre-COVID													
w/o commodities													
Autoregressive	0.834	0.699	0.719*	0.948	0.996	0.926	1.203	1.043	1.272	1.032	1.204	0.991	1.067
Net D+S (NDNS)	0.583			1.089	1.044	0.990	1.111	1.104	1.137	1.088	2.517	1.121	1.220
Net Demand (ND)	0.600			1.142	1.017	0.949	1.100	1.106	1.115	1.082	2.189	1.124	1.247
with commodities													
Autoregressive	0.796	0.745	0.756*	0.940	0.979	0.970	1.179	1.093	1.275	1.082	1.275	0.968	1.086
Net D+S (NDNS)	0.593			1.065	1.050	0.973	1.100	1.084	1.147	1.069	2.435	1.102	1.197
Net Demand (ND)	0.602			1.114	1.022	0.946**	1.089	1.103	1.140	1.078	2.184	1.112	1.239
Panel B: Full Sample													
w/o commodities													
Baseline	1.741	0.899**	0.906**	1.017	1.005	1.023	1.018	0.993	0.996	1.014	1.094	0.988	1.067
Net D+S (NDNS)	1.565			0.990	1.014	1.031	1.006	1.018	0.953	1.019	1.335	1.056	1.034
Net Demand (ND)	1.578			1.009	1.015	1.030	1.004	1.016	0.943	1.015	1.260	1.060	1.027
with commodities													
Baseline	1.739	0.910*	0.915**	1.028	1.002	1.024	1.013	0.994	1.000	1.015	1.098	0.984*	1.057
Net D+S (NDNS)	1.582			0.973	1.007	1.020	1.003	1.008	0.968	1.008	1.308	1.033	1.027
Net Demand (ND)	1.592			0.996	1.009	1.022	1.001	1.007	0.961	1.006	1.248	1.038	1.024

Notes: The table shows results from forecasts of headline PCE inflation with different predictors and samples, using the composite commodity indices. Panel A shows results for the sample until December 2019, and Panel B for the full sample. The two subpanels show separate results for specifications where the baseline includes or not commodity price growth as a predictor. Column (1) shows the root mean squared error (RMSE) of the baseline forecast, for different models. For the "autoregressive" model, it shows the RMSE for the simple autoregressive model with no additional predictors, as defined in the text (7). The "NDNS" and "ND" models include Net Demand+Net Supply or Net Demand only to the autoregressive model. Columns 2-13 show RMSEs relative to the baseline models from specifications which include additional predictors to the baseline models, one at a time. Columns 2-3 also show the RMSEs of the NDNS and ND models relative to the autoregressive one, for convenience. The additional predictors are: EFFR=effective federal funds rate; YC=10 minus 2 year government bond yield; USD=trade-weighted US dollar log monthly change; GAP: unemployment gap; IP=monthly log change in industrial production; VIX: CBOE Volatility Index; SPF: 12 month CPI inflation expectations from the Survey of Professional Forecasters; SPX: log monthly change in the SP500 index; CLEV: Cleveland Fed options implied 12-month inflation expectations; MICH: Household 12-month inflation expectations from the Michigan Survey. Sample: March 2001-June 2023. $^{*}p < 0.1$,** $p < 0.05$,*** $p < 0.01$.

TABLE E6. Forecast performance (RMSE), CPI inflation - alternative lag structure

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.045	0.663*	1.049	0.682*	0.845**	1.161	0.808**
	h=6	1.420	0.874	1.089	0.845	0.928	1.158	0.980
	h=3	1.743	1.047	1.165	0.955	1.036	1.269	1.197
Controls, Autoregressive								
	h=12	1.368	0.763**	1.092	0.735***	0.876	1.182	0.891
	h=6	1.626	1.004	1.070	0.996	1.008	1.165	1.094
	h=3	2.019	1.015	1.041	0.994	0.989	1.123	1.084
Panel B: with commodities								
Autoregressive								
	h=12	0.993	0.705*	1.036	0.717**	0.914	1.214	0.898
	h=6	1.439	0.933	1.100	0.886	0.964	1.126	1.017
	h=3	1.724	1.070	1.169	0.986	1.035	1.206	1.172
Controls, Autoregressive								
	h=12	1.335	0.804*	1.056	0.764***	0.907	1.157	0.942
	h=6	1.601	1.009	1.084	0.993	1.016	1.147	1.113
	h=3	1.854	1.028	1.057	1.002	1.001	1.115	1.108
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	2.207	0.884*	1.003	0.892*	0.929	1.014	0.910
	h=6	2.237	0.903*	1.021	0.895*	0.942	1.048	0.980
	h=3	2.561	1.008	1.058	0.976	1.022	1.095	1.091
Controls, Autoregressive								
	h=12	2.609	0.813*	1.017	0.810**	0.851*	1.031	0.837*
	h=6	2.562	0.903	1.019	0.900	0.930	1.058	0.970
	h=3	2.824	1.007	1.013	0.997	1.019	1.046	1.046
Panel B: with commodities								
Autoregressive								
	h=12	2.170	0.917*	1.019	0.922*	0.959	1.048	0.946
	h=6	2.212	0.945	1.050	0.924**	0.965	1.063	1.012
	h=3	2.376	1.031	1.071	0.990	1.019	1.087	1.104
Controls, Autoregressive								
	h=12	2.604	0.806*	1.017	0.800*	0.838	1.033	0.823
	h=6	2.512	0.925	1.043	0.915	0.945	1.067	0.993
	h=3	2.547	1.037	1.025	1.019	1.036	1.053	1.092
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of CPI inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Each specification includes 3 lags of the regressors. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-July 2023. $^{*}p < 0.1$,$^{**}p < 0.05$,$^{***}p < 0.01$.

TABLE E7. Forecast performance (RMSE), PCE inflation - alternative lag structure

	horizon	baseline	(1)	(2)	(3)	(4)	(5)	(6)
PRE-COVID SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	0.846	0.701	1.027	0.715*	0.874*	1.094	0.830
	h=6	1.043	0.831	1.072	0.818	0.872	1.109	0.917
	h=3	1.256	1.018	1.125	0.952	1.012	1.205	1.139
Controls, Autoregressive								
	h=12	1.070	0.868	1.150	0.787**	0.950	1.216	0.933
	h=6	1.190	0.990	1.061	1.005	0.999	1.113	1.071
	h=3	1.432	1.019	1.029	0.998	0.997	1.087	1.073
Panel B: with commodities								
Autoregressive								
	h=12	0.808	0.749	1.024	0.754*	0.938	1.143	0.913
	h=6	1.036	0.894	1.080	0.879	0.930	1.081	0.967
	h=3	1.182	1.028	1.136	0.975	1.010	1.152	1.112
Controls, Autoregressive								
	h=12	1.047	0.932	1.122	0.829	0.991	1.205	1.000
	h=6	1.131	0.996	1.057	1.017	1.013	1.093	1.100
	h=3	1.247	1.009	1.030	1.006	1.003	1.072	1.079
FULL SAMPLE								
Panel A: without commodities								
Autoregressive								
	h=12	1.759	0.916**	1.006	0.918**	0.948	1.002	0.933
	h=6	1.622	0.914	1.021	0.910*	0.959	1.036	0.996
	h=3	1.735	1.009	1.046	0.985	1.035	1.078	1.094
Controls, Autoregressive								
	h=12	2.025	0.865	1.029	0.850*	0.895	1.040	0.871
	h=6	1.808	0.926	1.018	0.932	0.964	1.046	0.996
	h=3	1.884	1.023	1.006	1.016	1.047	1.033	1.069
Panel B: with commodities								
Autoregressive								
	h=12	1.748	0.928*	1.020	0.926**	0.955	1.029	0.942
	h=6	1.603	0.940*	1.037	0.928**	0.973	1.043	1.017
	h=3	1.605	0.999	1.044	0.973	1.010	1.050	1.076
Controls, Autoregressive								
	h=12	2.032	0.849	1.032	0.828*	0.872	1.046	0.845
	h=6	1.761	0.928	1.026	0.931	0.963	1.046	1.004
	h=3	1.689	1.015	0.997	1.009	1.034	1.017	1.073
Composite Indices								
Net Supply			X	X				
Net Demand			X		X			
Supply Increase/Decrease							X	X
Demand Increase/Decrease						X		X

Note: The table shows results from forecasts of PCE inflation across different monthly horizons, using the composite commodity indices. The upper half shows results for the sample until December 2019, and the lower half for the full sample. Each specification includes 3 lags of the regressors. Columns (1)-(6) show the root mean squared error (RMSE) of each forecast relative to the baseline. Model 1 includes as predictors net supply and net demand; model 2 includes net supply; model 3 includes net demand; model 4 includes demand increase and demand decrease; model 5 includes supply increase and supply decrease; and model 6 includes supply increase, supply decrease, demand increase and demand decrease. The specification is given by (7). Each specification includes some of our indexes, as indicated. Where noted, controls are: FFR, 10-year minus 2-year US treasury, trade-weighted US dollar log monthly change, and the unemployment gap. Panel B also includes log monthly change in commodity prices. Sample: March 2001-June 2023. $^{*}p < 0.1$,** $p < 0.05$,*** $p < 0.01$.

TABLE E8. Forecast performance - alternative weighting

	(1)	(2)	(3)	(4)	(5)	(6)
	baseline	news-weighting	production-weighting	baseline	news-weighting	production-weighting
	Mean Squared Error criterion					
Panel A: Pre-COVID	CPI			PCE		
without commodities	1.034	0.696**	0.832**	0.834	0.719*	0.860**
with commodities	0.979	0.734**	0.861*	0.796	0.756*	0.886*
Panel B: Full Sample						
without commodities	2.177	0.884*	0.943	1.741	0.906**	0.968
with commodities	2.149	0.904*	0.968	1.739	0.915**	0.984
	Mean Absolute Error criterion					
Panel A: Pre-COVID	CPI			PCE		
without commodities	0.836	0.652***	0.758**	0.651	0.669**	0.792**
with commodities	0.776	0.707**	0.793**	0.629	0.724**	0.810**
Panel B: Full Sample						
without commodities	1.449	0.801***	0.888**	1.155	0.823***	0.915
with commodities	1.388	0.843***	0.933	1.152	0.844***	0.936

Note: The table shows results from forecasts of CPI and PCE inflation across 12-month horizons. Columns (1) and (4) show the forecast performance of the basic autoregressive model including log monthly change in commodity prices. Columns (2) and (5) show the forecast performance (relative to the baseline) of models that include our net demand composite index with (implicit) news-weighting. Columns (3) and (6) instead show the forecast performance (relative to the baseline) of models that include the net demand composite index constructed by production-based weights. The top part of the table shows results using the root mean squared error (MSE) criterion, and the bottom half using the mean absolute error (MAE) criterion. Sample: March 2001-June 2023. $^{*}p < 0.1$,** $p < 0.05$,*** $p < 0.01$.

TABLE E9. Forecast performance, Headline inflation - factors

	(1) Baseline	(2) NetDS	(3) NetD	(4) BC	(5) GR	(6) NDis	(7) ND+ NetDS	(8) ND+ NetD
Panel A: Mean Squared Error criterion								
CPI pre-COVID								
w/o commodities	1.034	0.675*	0.696**	1.204	1.103	0.992	0.673*	0.684*
with commodities	0.979	0.722*	0.734**	1.177	1.089	0.984	0.721	0.726
CPI full sample								
w/o commodities	2.177	0.872*	0.884*	1.003	1.020	0.850	0.789	0.783
with commodities	2.149	0.895*	0.904*	0.992	1.023	0.851	0.813	0.802
PCE - pre-COVID								
w/o commodities	0.834	0.699	0.719*	1.233	1.104	0.987	0.697	0.700*
with commodities	0.796	0.745	0.756*	1.224	1.087	0.975	0.738	0.735
PCE - full sample								
w/o commodities	1.741	0.899**	0.906**	1.002	1.001	0.817	0.787*	0.776*
with commodities	1.739	0.910**	0.915**	0.992	0.990	0.824	0.806	0.792
Panel B: Mean Absolute Error criterion								
CPI pre-COVID								
w/o commodities	0.836	0.648**	0.652***	1.186	1.094	1.022	0.648**	0.660**
with commodities	0.776	0.717**	0.707***	1.155	1.089	1.000	0.716*	0.725*
CPI full sample								
w/o commodities	1.449	0.787***	0.801***	1.117	1.052	0.879	0.734***	0.739***
with commodities	1.388	0.835***	0.843***	1.107	1.070	0.874	0.775**	0.774**
PCE - pre-COVID								
w/o commodities	0.651	0.688*	0.669**	1.231	1.108	1.028	0.691*	0.665**
with commodities	0.629	0.732*	0.724**	1.202	1.075	0.981	0.734*	0.715**
PCE - full sample								
w/o commodities	1.155	0.826**	0.823***	1.119	1.044	0.852	0.741***	0.728***
with commodities	1.152	0.845**	0.844***	1.100	1.023	0.821	0.757**	0.738**

Notes: The table shows results from forecasts of headline CPI and PCE inflation with different predictors and samples, using the unconditional commodity factor, and, as denoted, the composite commodity indices. Panel A shows results using the root mean squared error (MSE) criterion, and Panel B with the mean absolute deviation (MAE) criterion. The subpanels show separate results for specifications where the baseline includes or not commodity price growth as a predictor. Column (1) shows the prediction error of the baseline autoregressive model with no additional predictors, or with commodity price growth, as defined in the text (7). Columns 2-8 include further predictors, as denoted: BC=business cycle factor; GR: geopolitical risk factor; NDis: natural disaster factor. The "NetDS" and "NetD" models include Net Demand+Net Supply or Net Demand only composite indices to the autoregressive model. Columns 7 and 8 include the composite indices, as denoted, to the geopolitical risk factor. Sample: March 2001-June 2023. $^{*}p < 0.1$,** $p < 0.05$,*** $p < 0.01$.