Estimating the Macroeconomic Effects of Oil Supply News

Lorenzo Mori
Bank of Italy

Gert Peersman
Ghent University

SUPPLEMENTARY APPENDIX

A. OPEC Response to News Channel and ex ante Predictability

The predictability documented in Table 1 may appear surprising from the perspective of standard asset-pricing arguments, but aligns with studies documenting predictability of US monetary policy surprises (Miranda-Agrippino and Ricco, 2023; Bauer and Swanson, 2023a). As emphasized by Bauer and Swanson (2023a), predictability of announcement surprises can arise from imperfect private-sector knowledge of policymakers' reaction functions. More generally, ex post predictability in financial markets is a standard implication of models with imperfect information and learning (Timmermann, 1993). Applied to oil markets, this perspective suggests that investors may not fully observe how OPEC adjusts supply in response to macroeconomic and financial conditions. Market participants update their beliefs only after the announcement, generating ex post predictability in observed oil price surprises that reflects the discrepancy between realized policy actions and prior expectations.

To assess whether such an "OPEC response to news" channel may be the source of the predictability, we provide additional evidence on the relationship between financial uncertainty and oil price surprises around OPEC announcements. We begin by focusing on the VXO, which emerges as a strong predictor of the instrumental variable (IV) in Table 1 and is also included in our VAR specification. Specifically, we estimate the following regression:

$$IV_t = constant + \sum_{i=1}^{n} \beta_i \triangle VXO_{t-i} + \varepsilon_t \tag{1}$$

The results are reported in Table A1 for $n = 6$. The first four lags of changes in the VXO enter with negative and statistically significant coefficients, while lags five and six are insignificant. The sum of the coefficients is highly significant ($p<0.01$). The consistency in sign and significance across multiple lags indicates that the predictability is not driven by isolated observations. In Mori and Peersman

(2024), we further show that this negative relationship holds symmetrically for increases and decreases in the VXO, and also applies to subsamples. Overall, periods of elevated financial uncertainty are systematically associated with more negative oil price surprises around OPEC announcements.

Table A1: Predictability of the Instrumental Variable

Variable	*#obs*	**Sign**	**Overall significance (p-value)**
VXO	116	−	0.00***
Fresh news	116	−	0.83
Oil VXO	30	−	0.05**
Fresh news	30	−	0.48
Implied oil volatility (1m)	38	−	0.07*
Fresh news	38	−	0.71
Implied oil volatility (3m)	38	−	0.06*
Fresh news	38	−	0.14
Realized oil volatility	80	−	0.03**
Fresh news	80	−	0.14
S&P 500	116	+	0.22
Fresh news	116	−	0.70
MSCI	116	+	0.19
Fresh news	116	+	0.65
Oil prices	116	+	0.02**
Fresh news	116	+	0.04**
Commodity prices	116	+	0.38
Fresh news	116	+	0.03**

Note: The sample includes only months with OPEC announcements. "Fresh news" is defined as the change in the corresponding daily indicator within the announcement month, measured from the last day of the previous month up to the day preceding the announcement. For VXO, S&P 500, MSCI, oil prices and commodity prices six lags are included; for the remaining indicators, one lag is used due to shorter data availability. All variables are expressed in first differences, except for the S&P 500, MSCI, oil prices and commodity prices, which are expressed in log differences. The "overall significance" column reports the (robust) p-value from the joint test of the lag coefficients in the baseline specification, and the (robust) p-value of the "fresh news" regressor in the augmented specification. Sign indicates the sign of the sum of (lag) coefficients. Significance levels are denoted by $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$.

Table A1 also assesses the robustness of this result using measures of oil-market uncertainty: the CBOE crude oil volatility index (available from 2007M5), the 60-day historical volatility of Brent prices (Bloomberg, available from 1995M1), and one- and three-month implied volatility of one-month oil futures (Bloomberg, available from 2005M1). Owing to shorter sample availability, we include only one lag for these measures. Despite this limitation, all indicators predict the IV with

negative coefficients. This suggests that OPEC's supply decisions may respond to uncertainty about oil prices, which is strongly correlated with the VXO, rather than macroeconomic uncertainty. However, given the broader and longer coverage of the VXO, we use this measure in the VAR estimations.

We next examine whether the predictability also applies to very recent information. Following Bauer and Swanson (2023a), we construct "fresh news" measures based on changes in daily indicators from the beginning of the month up to the day preceding the OPEC meeting. We find no evidence that such fresh news predicts the IV. Instead, predictability is driven by information arriving earlier, consistent with the view that OPEC reacts to news perceived as sufficiently persistent or economically relevant. This pattern closely mirrors the evidence documented for monetary policy announcements by Bauer and Swanson (2023a). Extending the analysis to daily stock price indicators such as the S&P 500 and the MSCI index yields similar conclusions: neither their lagged components nor their within-month "fresh news" components predict the IV, suggesting that the contemporaneous exogeneity of the instrument remains intact. The only exceptions are oil prices and the Bloomberg Commodity Index, for which within-month changes appear to have some predictive power. A closer inspection, however, reveals that this result is mainly driven by the October 2008 observation. When this month is excluded from the regressions, the p-values associated with both "fresh news" variables exceed 0.10. We therefore do not further adjust the instrument to account for this predictability. When we nonetheless pre-clean the instrument for these within-month changes, the resulting impulse responses remain very similar to those obtained from our benchmark VAR augmented with financial variables and, if anything, further accentuate the differences relative to Känzig (2021).

Finally, it is important to emphasize that ex post predictability does not imply ex ante predictability. We follow the approach of Bauer and Swanson (2023b) to explore this for the VXO. Specifically, using expanding estimation windows starting from an initial sample of five years, we find that forecasting models incorporating lagged changes in the VXO do not outperform a simple zero-forecast benchmark. Forecasting gains are economically small and statistically insignificant according to Diebold–Mariano tests. For example, the out-of-sample RMSE (MAE) changes only from 2.61 (1.90) under the benchmark to 2.60 (1.97) under the VXO-based model, corresponding to a gain (loss) of just 0.32 percent (3.9 percent), with Diebold–Mariano p-values of 0.94 (0.54). Similar conclusions hold for oil price innovations and oil supply news shocks identified in Känzig (2021): while lagged VXO changes significantly predict these series in sample, we find no statistically significant evidence of out-of-sample predictability. Taken together, the coexistence of strong in-sample predictability and the absence of out-of-sample forecasting power supports the interpretation that our findings reflect an "OPEC response to news" channel.

B. Results Covering the Post-COVID Era

Figure A1 illustrates why estimating monthly oil-market SVAR-IV models identified using OPEC announcements becomes problematic when the sample includes the COVID-19 period. During this episode, world and US industrial production experienced unprecedented swings relative to their historical distribution growth rates that were overwhelmingly driven by the pandemic rather than by oil-market fundamentals. The contractions in world industrial production observed in January and March 2020, and the sharp rebound during the reopening phase in June and July represent the largest monthly declines and increases in the entire sample. A similar pattern emerges for the United States.

Figure A1: Industrial production growth and instrumental variable during COVID

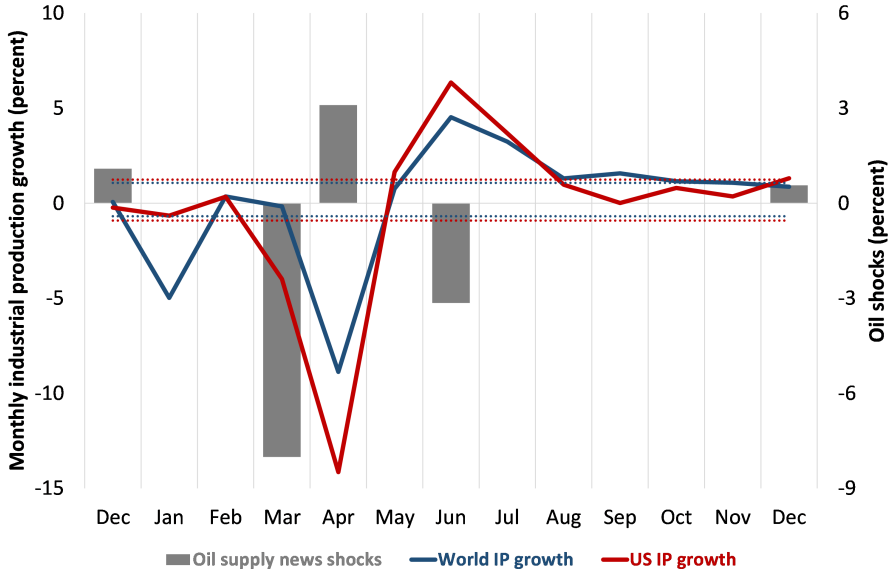

Note: World and US monthly industrial production growth from 2019M12 to 2020M12 (percent). The dotted lines indicate the 5th–95th percentiles of each series over the full sample period. The grey bars report oil futures price changes around OPEC announcements for the corresponding months (percent).

At the same time, the period was characterized by exceptionally large high-frequency oil supply news shocks. Following the breakdown of negotiations between OPEC and Russia over production cuts—which led to Russia's exit from the agreement and the temporary collapse of the OPEC+ alliance—oil futures prices fell by 8.05 percent on 6 March 2020, one of the largest daily price movements in decades. Moreover, the instrument took values of 3.09 and −3.15 percent in April and June 2020, respectively, closely mirroring the pandemic-driven swings in industrial production. Absent appropriate adjustments, this spurious comovement between macroeconomic outcomes and the instrument can severely distort the estimation and inference, as confirmed by our empirical results.

Figure A2: Impact of oil supply news shocks including the post-COVID era

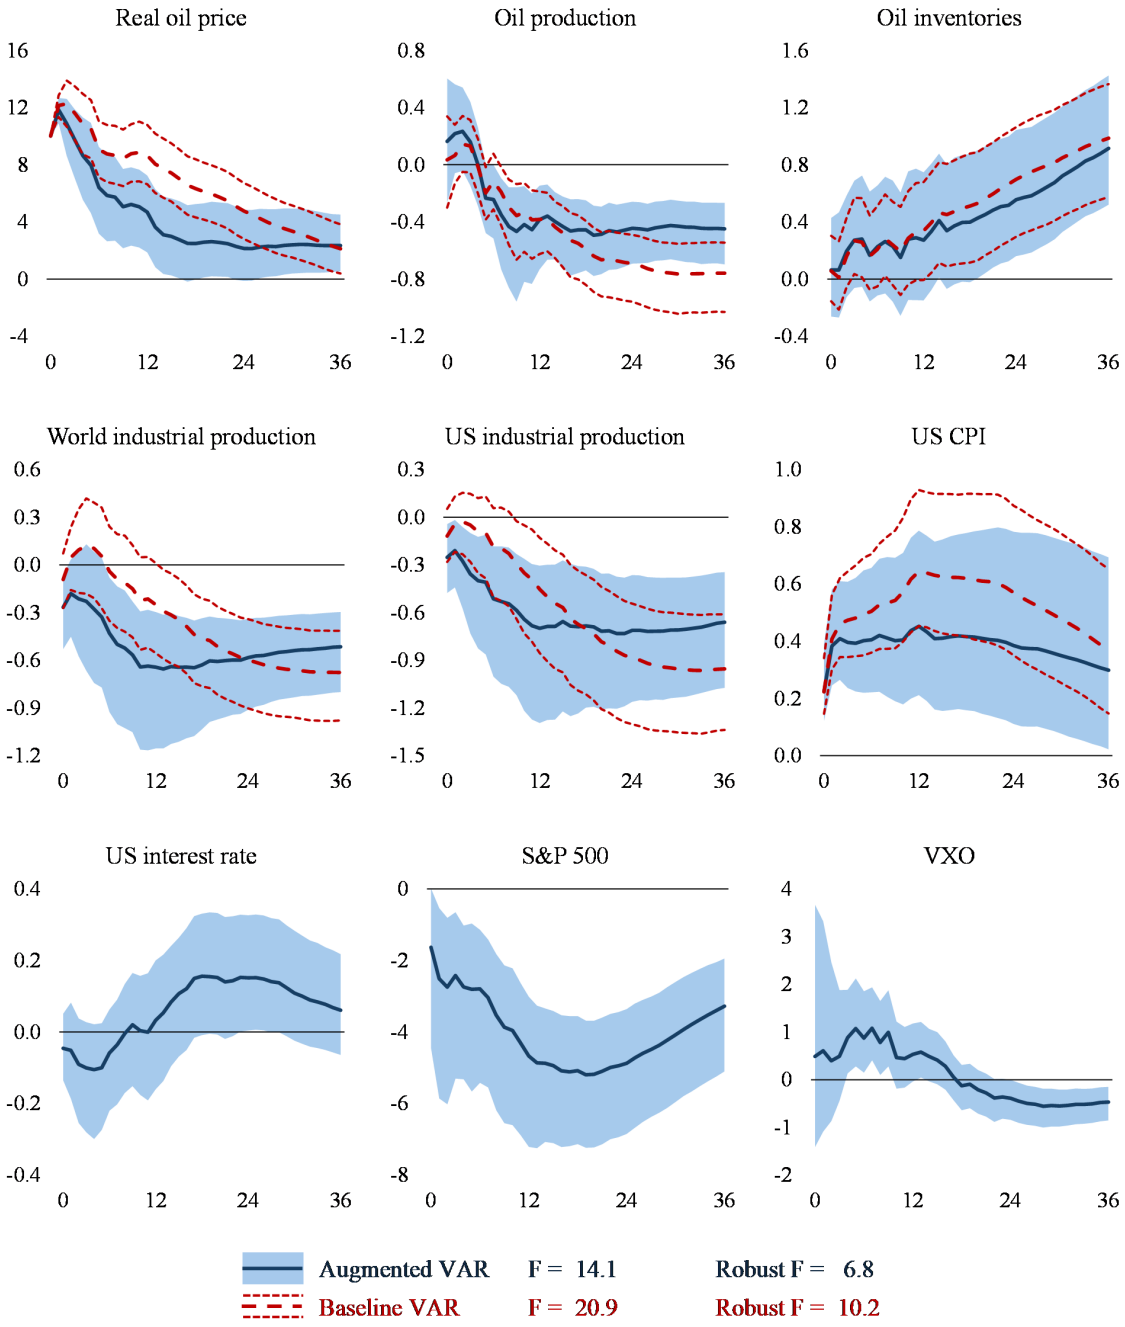

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Baseline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

To account for these extreme observations and to assess the robustness of our findings in the post-COVID period, we include dummy variables for the months 2020M1–2020M6 and re-estimate both

the six-variable baseline and the nine-variable augmented SVAR-IV specifications over the extended sample period 1974M1–2024M8. We also formally test the significance of the differences between the two specifications, corresponding to Figures 1 and 2 in the paper. As shown in Figures A2 and A3, our main findings remain valid when the sample is extended to include the post-COVID era.

Figure A3: Tests for differences between augmented and baseline VAR impulse responses including post-COVID era

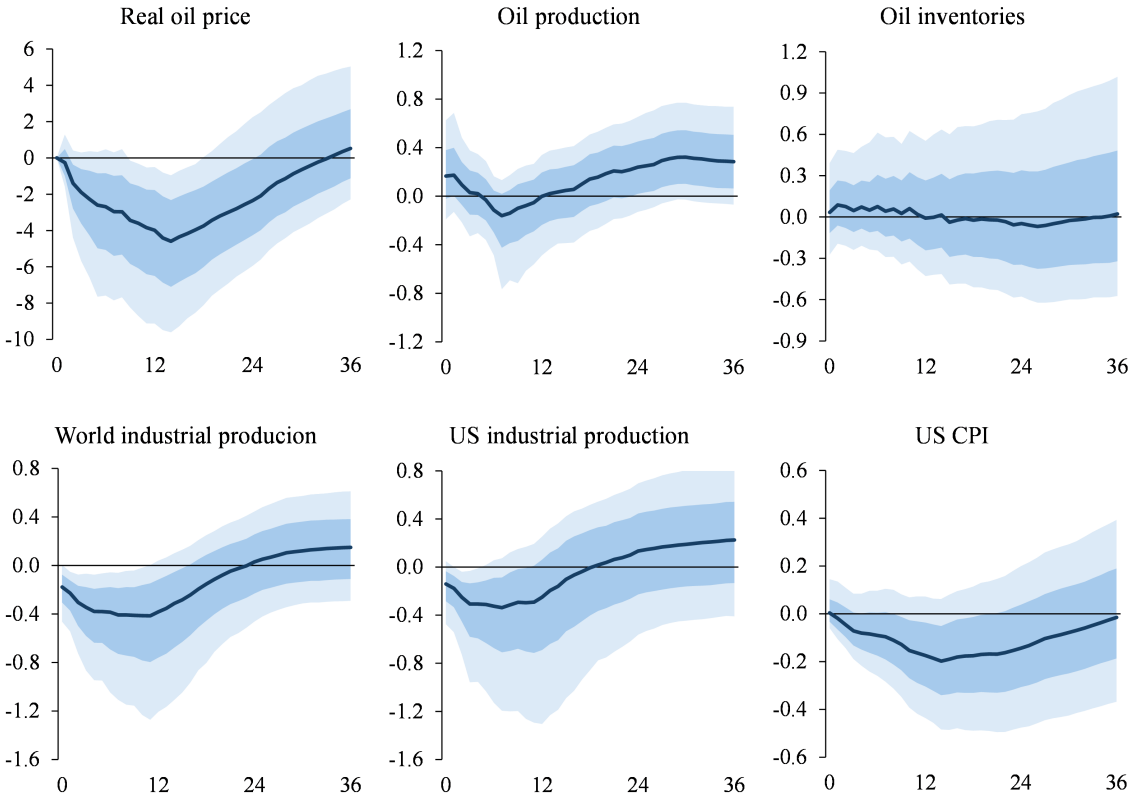

Note: Estimated differences between impulse responses of the VAR model augmented with financial variables and the baseline VAR for oil supply news shocks that raise oil prices by 10% on impact. Monthly horizon. 68% and 90% confidence intervals constructed using a moving block bootstrap.

C. Alternative Inference Methods

Figure A4 and A5 display the impulse responses of the baseline and augmented VAR models using alternative inference methods: the Bayesian approach from Miranda-Agrippino and Ricco (2023), and the weak-instrument robust inference method of Montiel Olea et al. (2021), respectively.

Figure A4: Impact of oil supply news shocks - Bayesian SVAR

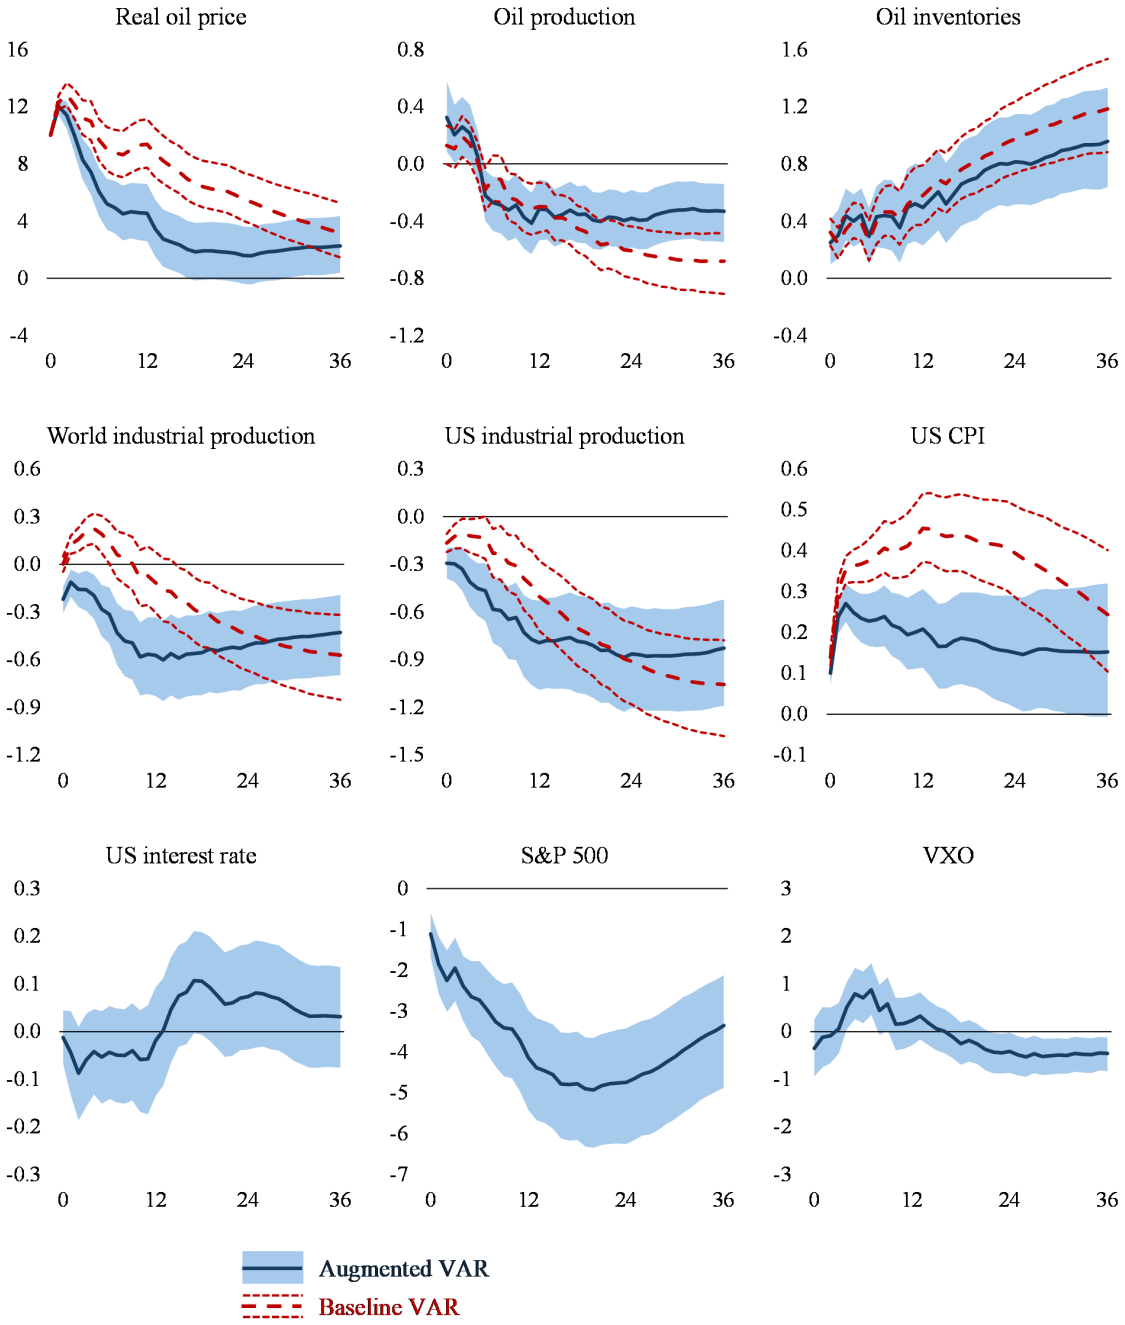

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Baseline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% credible intervals constructed using the Bayesian approach of Miranda-Agrippino and Ricco (2023), using a flat prior.

Figure A5: Impact of oil supply news shocks - Weak instrument robust inference

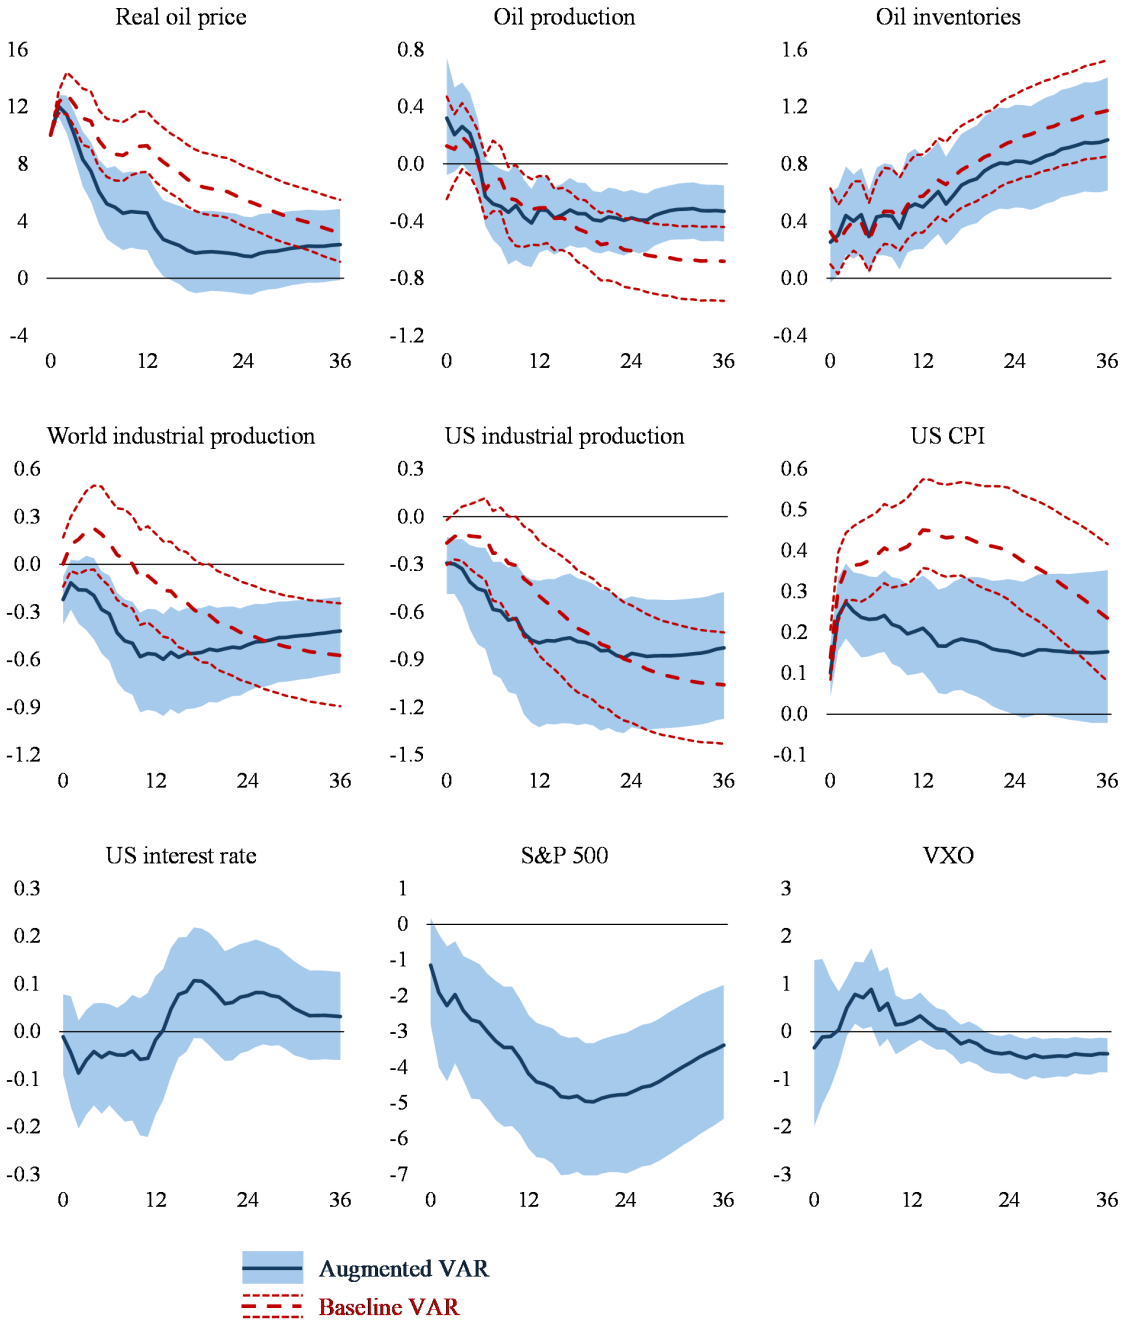

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Basline VAR (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% confidence intervals constructed using the weak instrument robust inference approach of Montiel Olea et al. (2021).

D. Historical Contribution and Relevance of Oil Supply News Shocks

Figure A6 shows for the original (top panel) and the augmented VAR model (middle panel), the cumulative historical contribution of oil supply news shocks to the real price of oil, together with the deviation of the actual real price from its baseline evolution implied in the VAR. The bottom panel presents the contribution to the S&P 500 in the augmented VAR. Table A2 reports the forecast error variance of the VAR variables attributed to oil supply news shocks at horizons 0, 12, and 36 months.

Figure A6: Historical contribution of oil supply news shocks

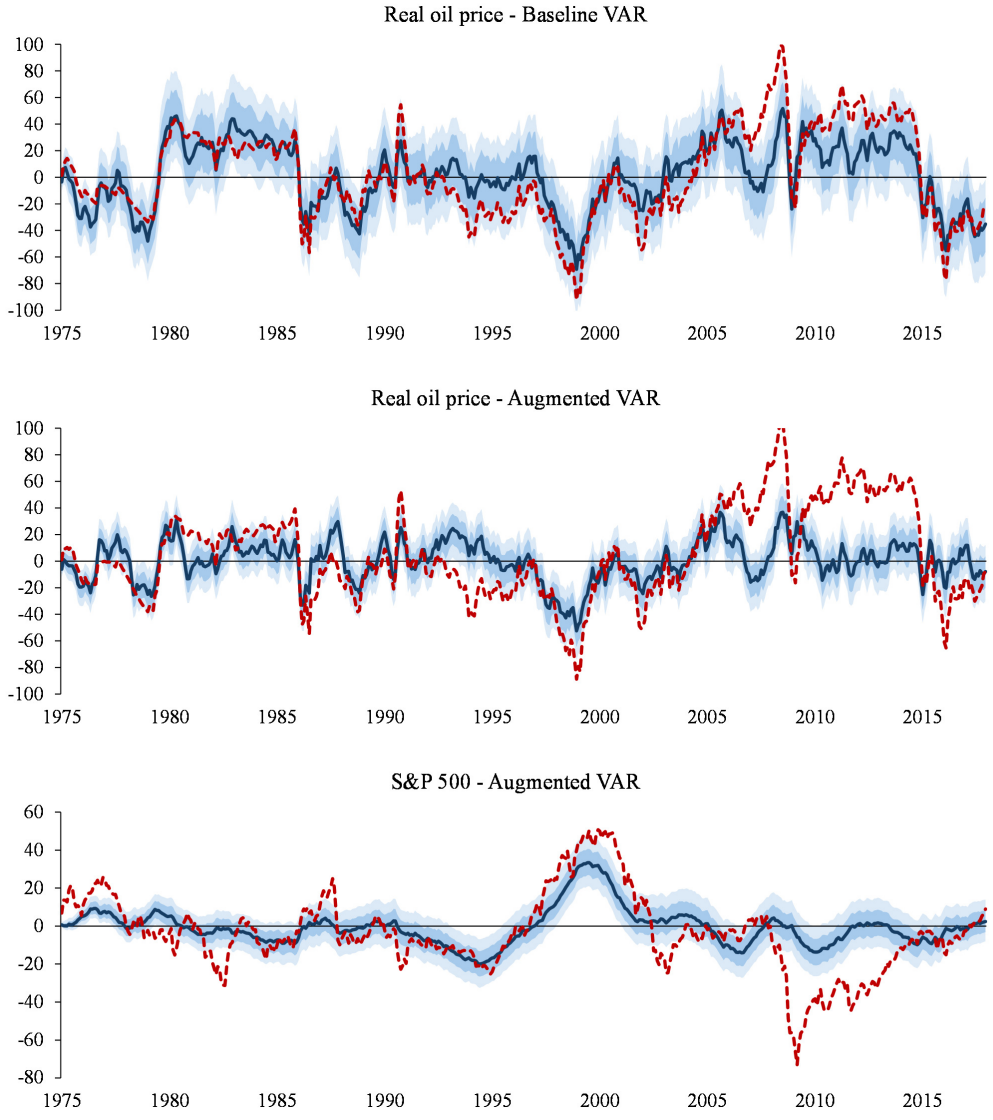

Note: Percentage points cumulative historical contribution, together with the actual deviation of the variables from their baseline evolution implied in the VAR. The top panel is based on the baseline VAR, while the middle and bottom panel are based on the VAR augmented with financial variables. 68% and 90% confidence intervals.

Table A2: Forecast error variance decompositions

		Baseline VAR		VAR with financial variables	
Real oil price	0	0.83	[0.39 0.91]	0.70	[0.12 0.81]
	12	0.75	[0.24 0.81]	0.41	[0.05 0.55]
	36	0.73	[0.20 0.75]	0.27	[0.05 0.43]
Oil production	0	0.00	[0.00 0.10]	0.02	[0.00 0.13]
	12	0.03	[0.02 0.17]	0.05	[0.02 0.17]
	36	0.24	[0.05 0.40]	0.12	[0.03 0.23]
Oil inventories	0	0.05	[0.00 0.21]	0.03	[0.00 0.18]
	12	0.11	[0.01 0.28]	0.08	[0.01 0.26]
	36	0.37	[0.04 0.52]	0.24	[0.02 0.42]
World industrial production	0	0.00	[0.00 0.19]	0.08	[0.00 0.29]
	12	0.01	[0.01 0.23]	0.10	[0.01 0.33]
	36	0.11	[0.04 0.45]	0.19	[0.02 0.38]
US industrial production	0	0.03	[0.00 0.22]	0.10	[0.00 0.32]
	12	0.04	[0.01 0.27]	0.18	[0.02 0.40]
	36	0.32	[0.07 0.52]	0.31	[0.04 0.44]
US CPI	0	0.20	[0.01 0.50]	0.10	[0.00 0.33]
	12	0.48	[0.11 0.68]	0.18	[0.01 0.38]
	36	0.38	[0.06 0.50]	0.09	[0.01 0.25]
US interest rate	0			0.00	[0.00 0.11]
	12			0.01	[0.00 0.11]
	36			0.02	[0.01 0.14]
S&P 500	0			0.04	[0.00 0.36]
	12			0.20	[0.02 0.46]
	36			0.41	[0.04 0.53]
VXO	0			0.00	[0.00 0.27]
	12			0.02	[0.02 0.23]
	36			0.04	[0.03 0.20]

Note: Forecast error variance of the variables explained by oil supply news shocks at horizons 0, 12 and 36 months, with 90% confidence intervals shown in the brackets.

E. Stability of Oil Supply News Shocks across Sample Periods

Figure A7 shows the instability in the baseline VAR-model across sample periods. The stability of the augmented SVAR-model including the post-COVID era can be observed in Figure A2.

Figure A7: Instabilities in the baseline VAR-model

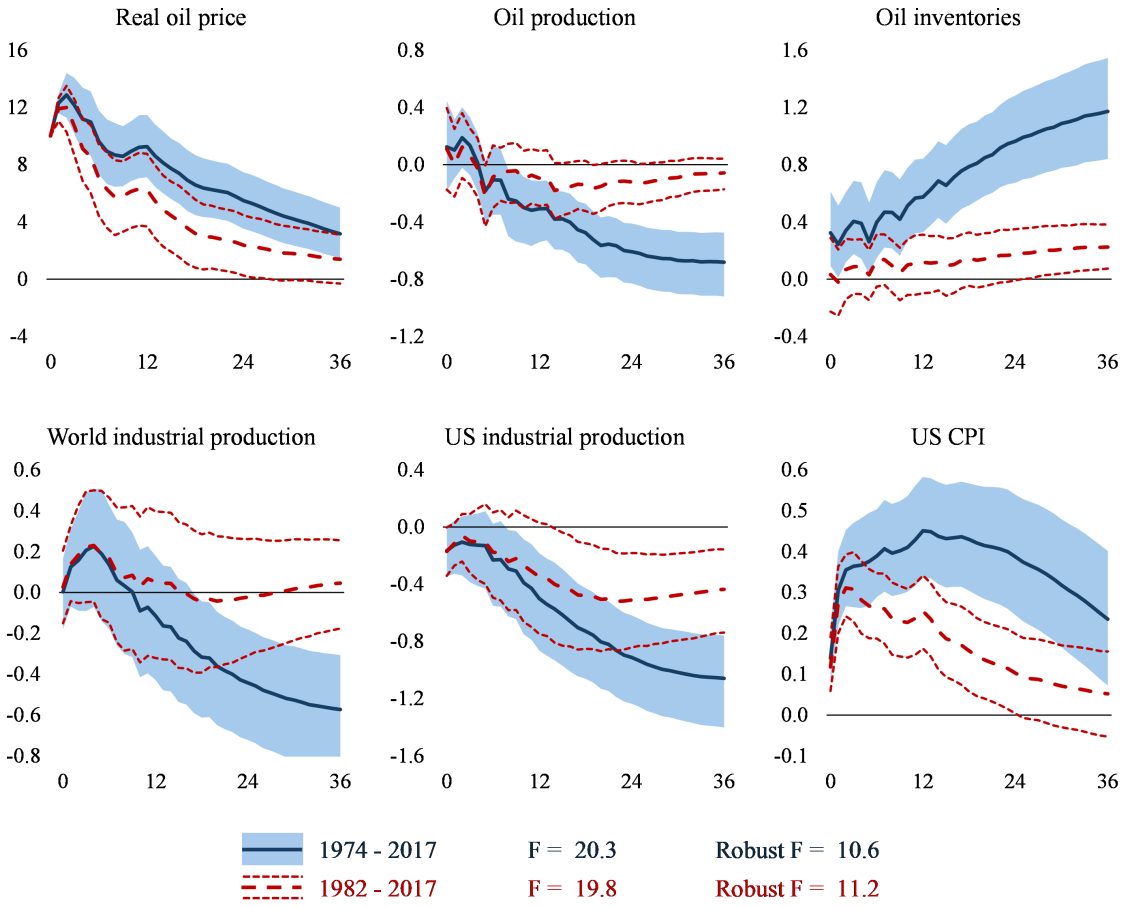

Note: Impulse responses to oil supply news shock that raises oil prices by 10% on impact, estimated over sample periods 1974-2017 and 1982-2017, respectively. 68% confidence intervals constructed using a moving block bootstrap.

F. The Presence of Information Effects

Degasperi (2025) argues that OPEC announcements may induce imperfectly informed markets to revise their beliefs about demand conditions, implying that surprises in oil futures prices around these announcements can reflect both supply- and demand-driven components. To disentangle these effects, he classifies oil futures price surprises based on their contemporaneous co-movement with equity

prices. Shocks associated with a negative co-movement of equity prices on the day of the announcement are interpreted as oil supply news and used as an instrument for supply shocks, whereas shocks accompanied by a positive co-movement are attributed to demand-related information effects. Consistent with our results, Degasperi (2025) finds that negative oil supply news shocks generate a much larger contraction in economic activity than those reported in Känzig (2021).

On one hand, Degasperi's refined instrumental variable does not seem to resolve the omitted-variables problem or the contamination of the instrument.[1] Specifically, the refined instrument for oil supply news shocks—after orthogonalization with respect to the lagged variables of the baseline VAR—is still predictable by uncertainty measures: it is Granger-caused by the VXO ($p < 0.02$) and by the financial uncertainty index of Ludvigson et al. (2021) ($p < 0.01$). Moreover, the resulting oil supply news shocks remain predictable by the joint common factors of McCracken and Ng (2016), including factors F6 and F7, the Global Financial Cycle of Miranda-Agrippino and Rey (2020), as well as by the S&P 500, the MSCI World index, the VXO, and the excess bond premium of Gilchrist and Zakrajšek (2012). In addition, the instability of the baseline VAR across sample periods persists when the refined instrument is used. By contrast, once we augment the information set with financial indicators—US interest rates, the S&P 500, and the VXO—the Granger-causality tests no longer reject the null that any of the variables in Table 1 predict either the identified structural shocks or the orthogonalized instrumental variable, and the sample-period instability disappears. Put differently, refining the instrument alone does not address the omitted-variables problem documented in this paper.

On the other hand, conditioning on the information set that includes financial variables substantially attenuates the role of potential information effects, rendering them negligible in sample periods starting in the early 1980s. Specifically, although the decline in world and US industrial production is larger when the SVAR is identified using the refined instrument (Figure A8), Figure A9 shows that the differences between the impulse responses obtained with the original instrument of Känzig (2021) and those based on Degasperi's refined instrument are statistically insignificant at the 90% confidence level once the VAR is augmented with financial variables.[2] Moreover, as shown in Figure A10 and A11, the impulse responses obtained with both instruments become virtually indistinguishable when the sample period starts in the early 1980s. Notably, the same applies when the SVAR is estimated over the sample period 1982M7–2023M5 with COVID dummies, which corresponds to the baseline sample used in Degasperi (2025). Taken together, these results indicate that the larger contraction in economic activity relative to Känzig (2021) primarily reflects omitted-variables bias rather than in-

[1] The refined instrumental variable is available at: https://github.com/riccardo-degasperi/OPEC-surprises.

[2] These tests are conducted by nesting both estimations within the bootstrap procedure. Specifically, for each bootstrap replication of the VAR model, we identify the shocks using the original and the refined instrument, respectively, and compute the difference between the corresponding impulse response functions. The differences and the confidence intervals for these differences are reported in Figure A9 and A11.

formation effects embedded in OPEC announcements and that incorporating financial variables in the SVAR model is sufficient to disentangle supply from demand news.

Figure A8: Impact of oil supply news shocks: original (Känzig) versus refined (Degasperi) instrument, 1974–2017 sample

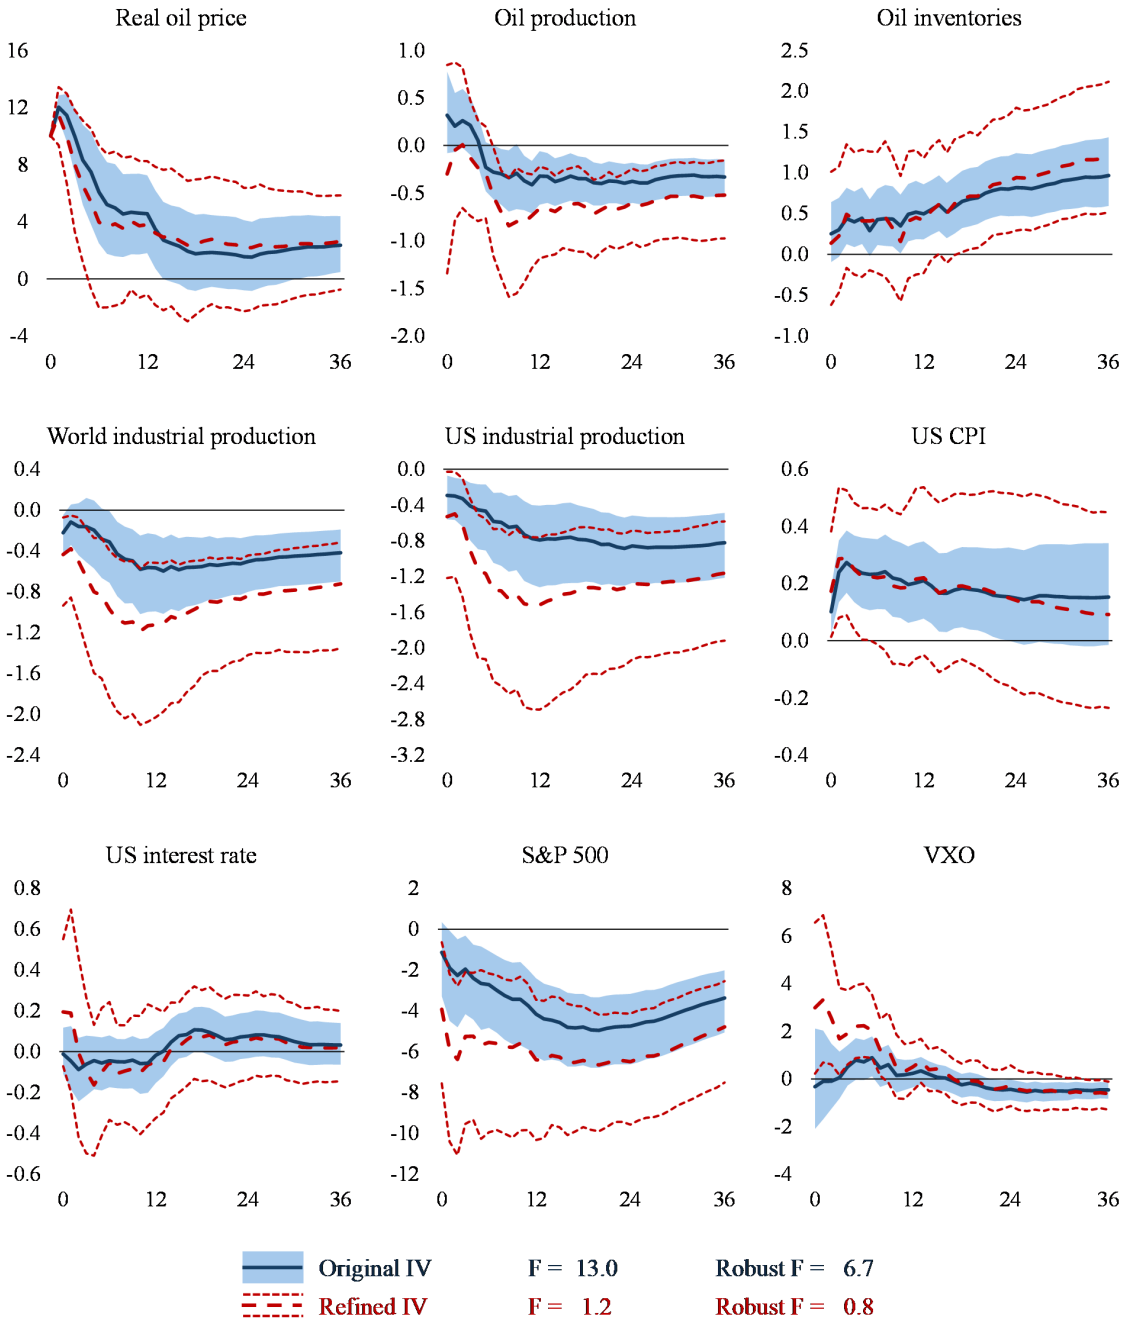

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. VAR identified with the refined instrument (red dotted responses) versus VAR model identified with the original instrument (blue full responses). Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Figure A9: Differences in impulse responses from the augmented SVAR model: original (Känzig) versus refined (Degasperi) instrument, 1974–2017 sample

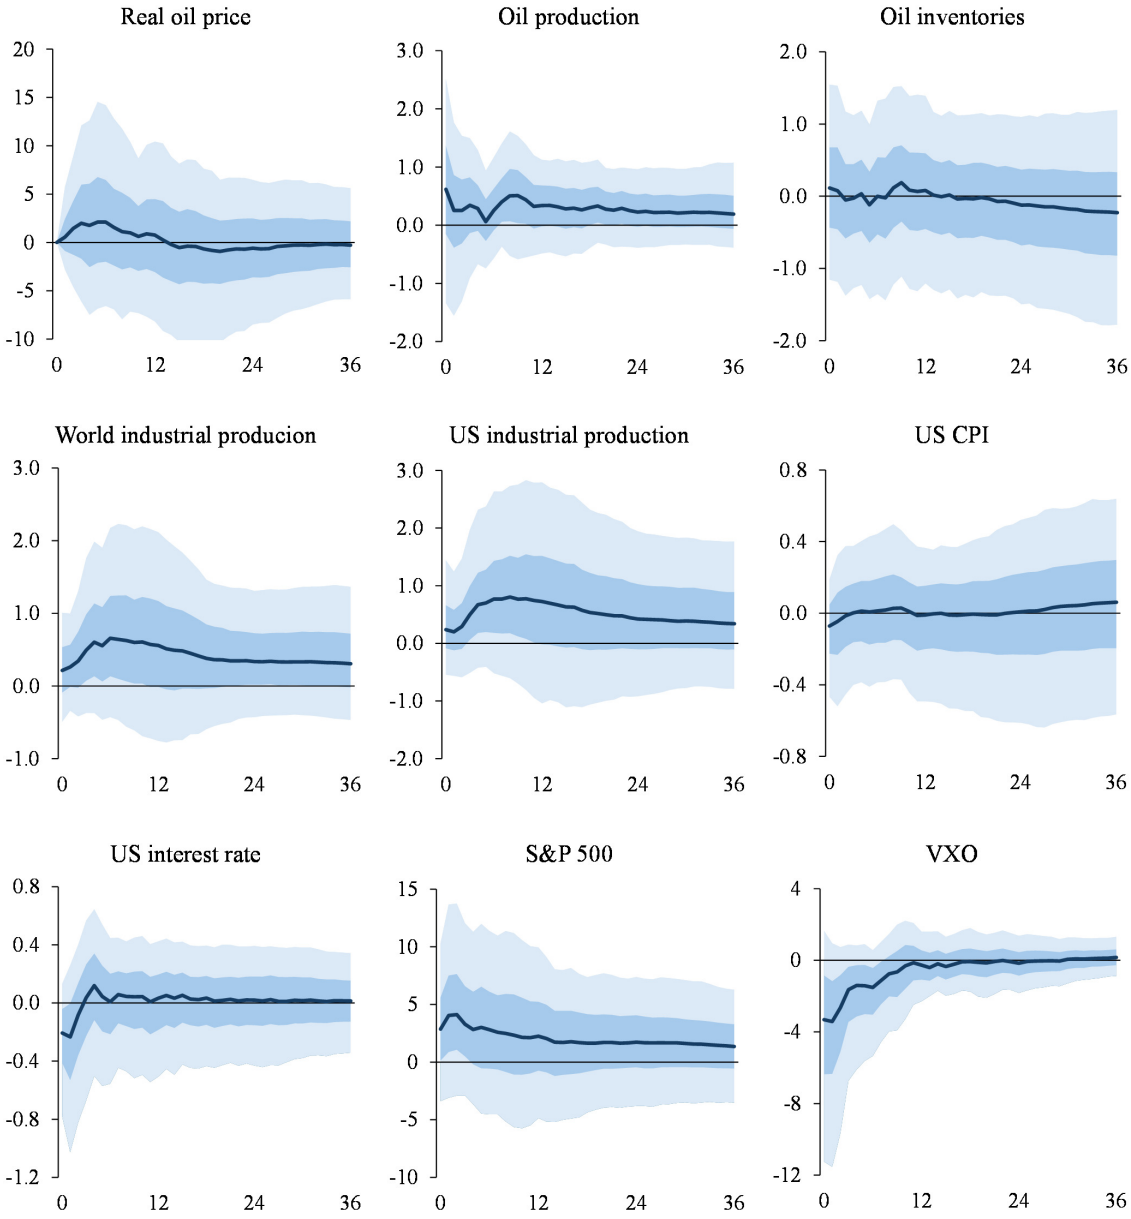

Note: Estimated differences between impulse responses of the VAR model augmented identified with the original instrument of Känzig (2021) versus the VAR model identified with the refined instrument of Degasperi (2021) for oil supply news shocks that raise oil prices by 10% on impact. Monthly horizon. 68% and 90% confidence intervals constructed using a moving block bootstrap.

Figure A10: Impact of oil supply news shocks: original (Känzig) versus refined (Degasperi) instrument, 1982–2017 sample

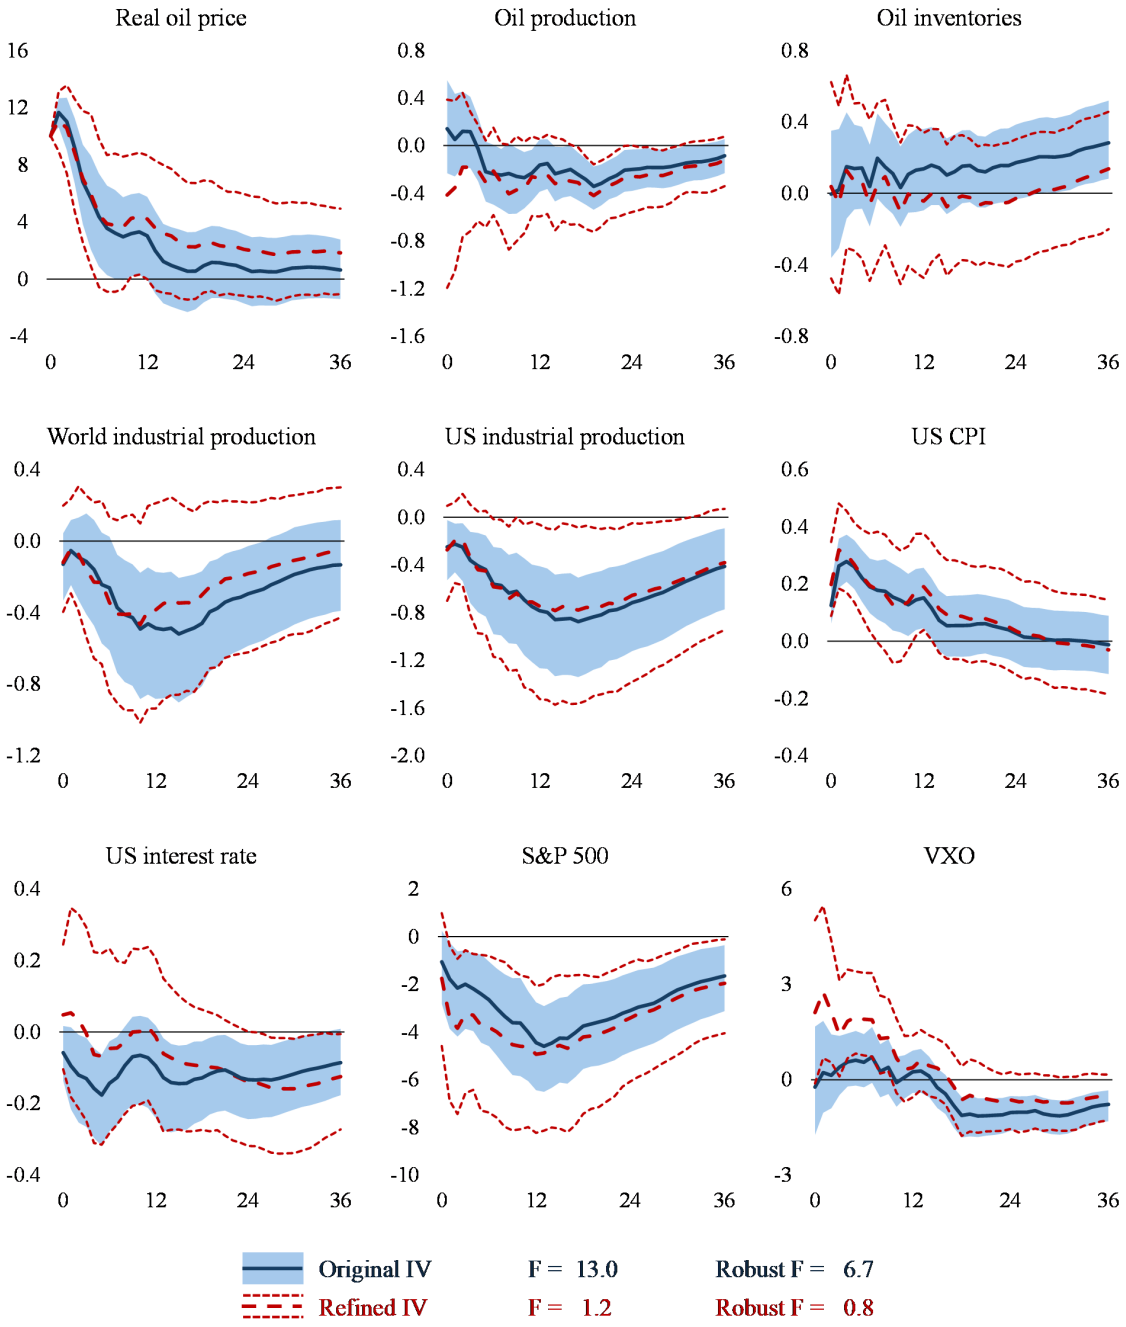

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. VAR identified with the refined instrument (red dotted responses) versus VAR model identified with the original instrument (blue full responses). Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Figure A11: Differences in impulse responses from the augmented SVAR model: original (Känzig) versus refined (Degasperi) instrument, 1982–2017 sample

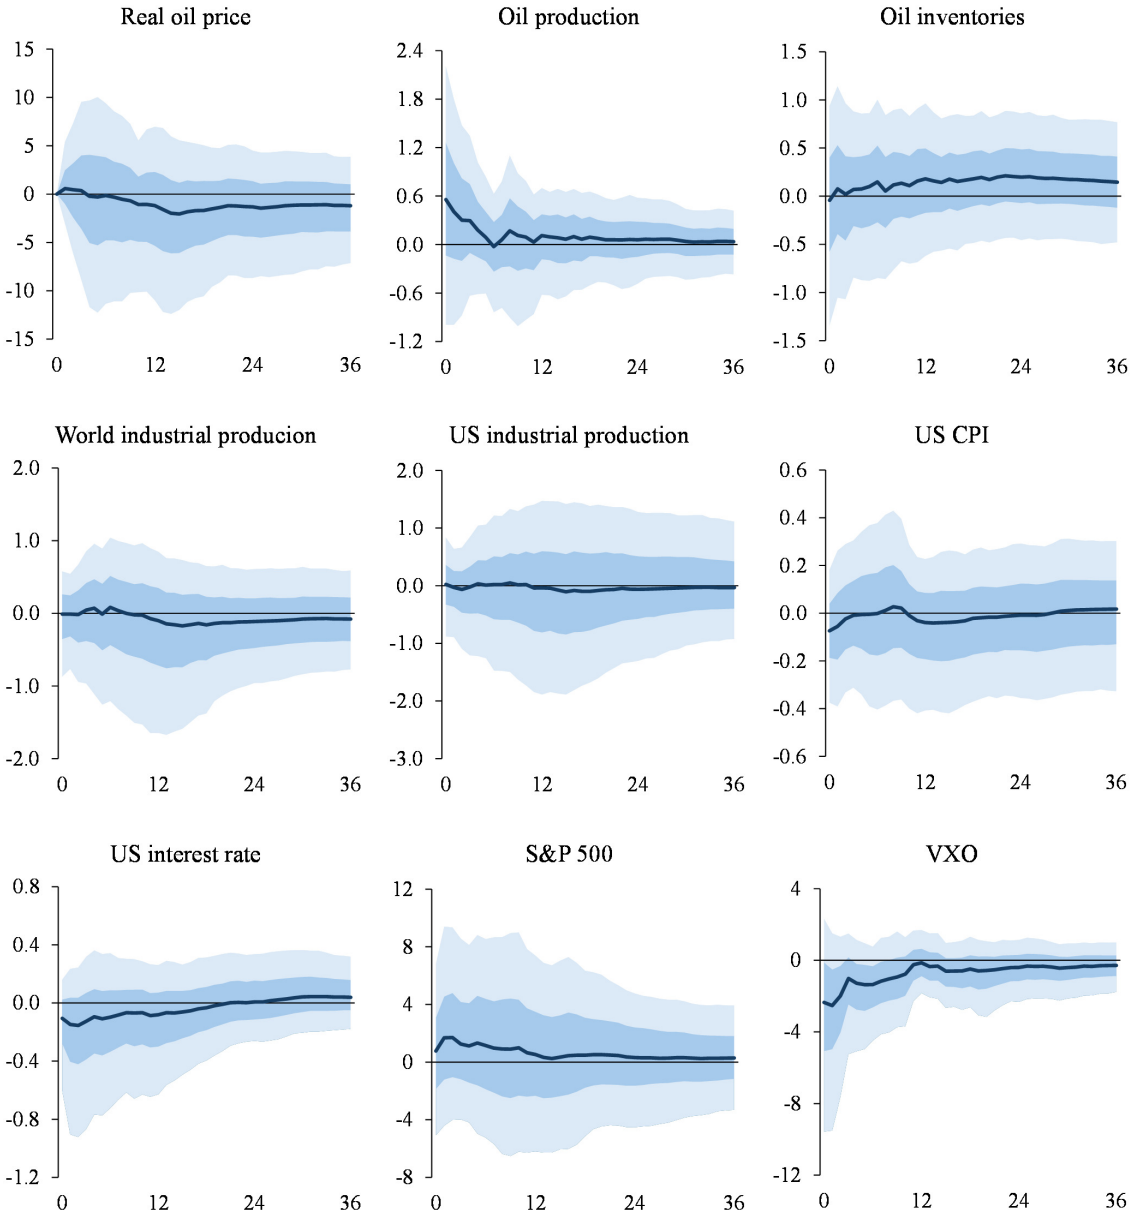

Note: Estimated differences between impulse responses of the VAR model augmented identified with the original instrument of Känzig (2021) versus the VAR model identified with the refined instrument of Degasperi (2021) for oil supply news shocks that raise oil prices by 10% on impact. Monthly horizon. 68% and 90% confidence intervals constructed using a moving block bootstrap.

G. Additional Analysis

A natural question for practitioners is whether it suffices to orthogonalize the external instrument with respect to its predictable components and then estimate the baseline VAR—excluding financial variables—using this "cleaned" instrument. To assess this possibility, Figure A12 compares the impulse responses from the augmented VAR with those obtained from the baseline six-variable VAR, where the high-frequency OPEC announcement surprises are first regressed on financial variables—specifically the S&P 500, the VXO, and the US one-year interest rate—and the resulting residuals are used as the external instrument. Figure A13 reports the differences between the impulse responses from the two specifications, along with confidence intervals constructed by nesting both estimations within the bootstrap procedure. Several impulse responses remain statistically different, indicating that instrument contamination alone cannot fully account for the distortions in the baseline results. This conclusion is reinforced by the analysis in Section 3, which shows that financial variables also play a crucial role contemporaneously, that is, at the moment of impact.

Figure A12: Impact of oil supply news shocks: augmented VAR versus baseline VAR with pre-cleaned instrument

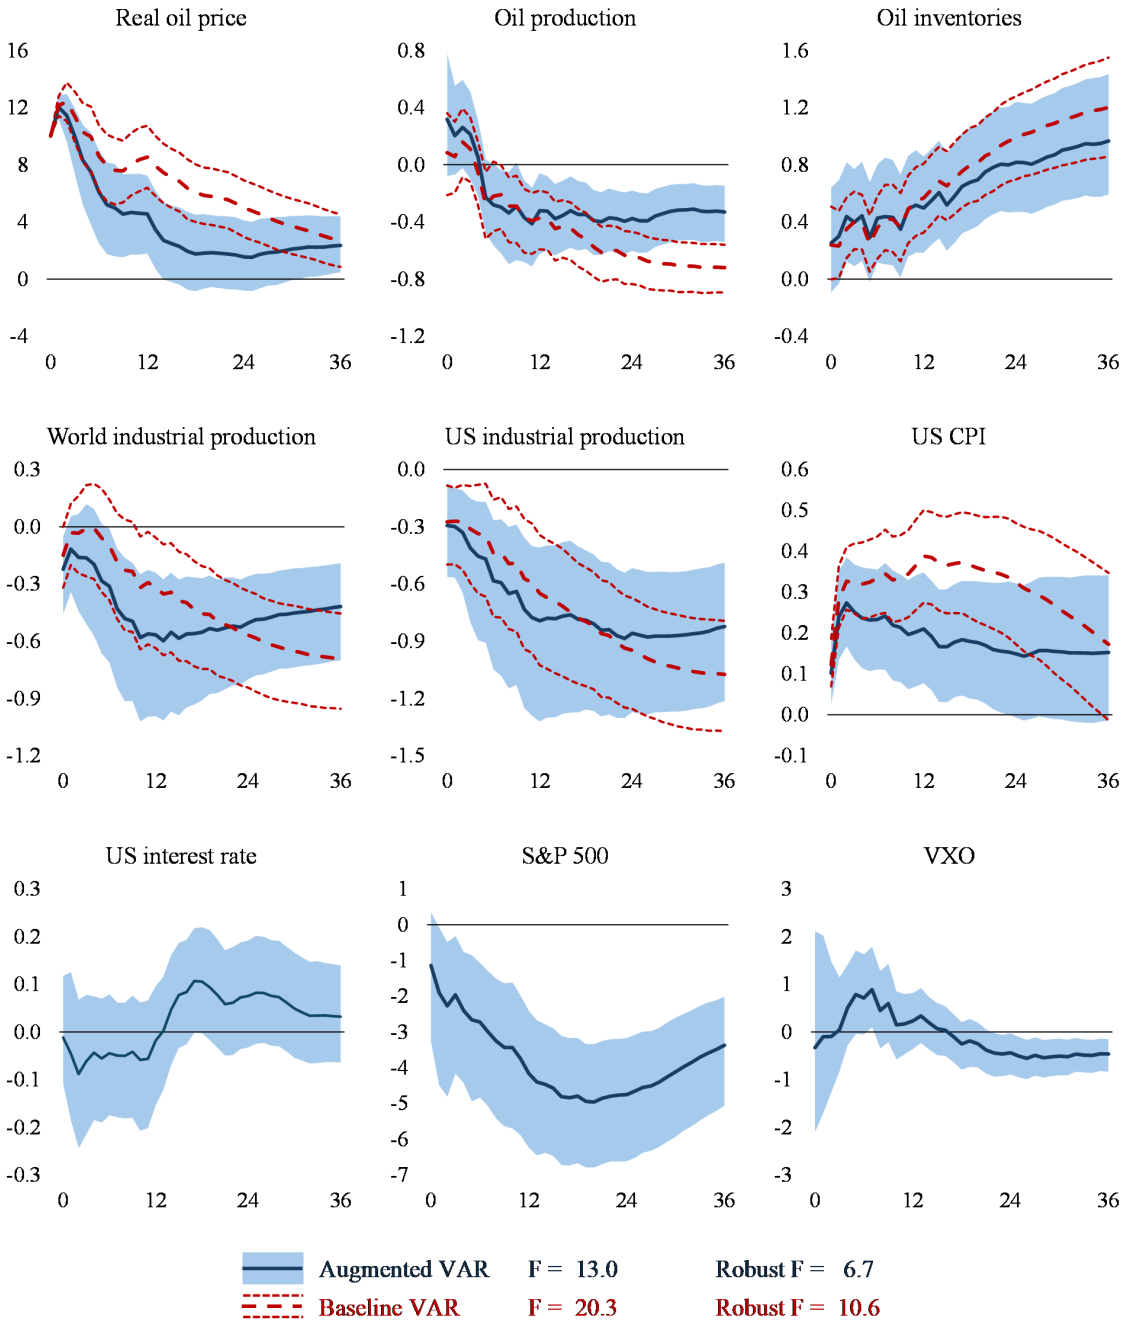

Note: Impulse responses to oil supply news shock that raise oil prices by 10% on impact. Baseline VAR identified with instrument orthogonalized to financial variables (red dotted responses) versus the VAR model augmented with financial variables (blue full responses). Monthly horizon. 68% confidence intervals constructed using a moving block bootstrap.

Figure A13: Differences in impulse responses from the augmented SVAR model and the baseline VAR with pre-cleaned instrument

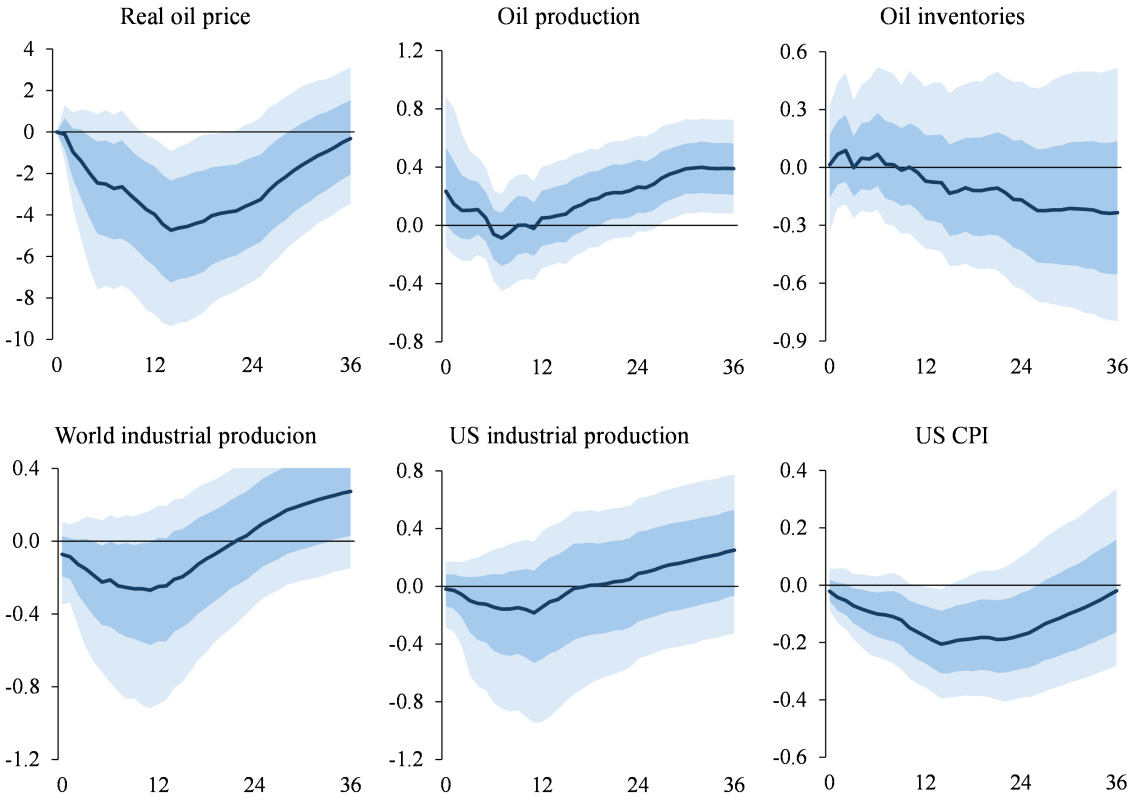

Note: Estimated differences between impulse responses of the baseline VAR identified with instrument orthogonalized to financial variables versus the augmented VAR model for oil supply news shocks that raise oil prices by 10% on impact. Monthly horizon. 68% and 90% confidence intervals constructed using a moving block bootstrap.

References

Bauer, M. and E. Swanson (2023a) "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, Vol. 113, pp. 664–700.

——— (2023b) "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, Vol. 37, pp. 87–155.

Degasperi, Riccardo (2025) "Identification of Expectational Shocks in the Oil Market using OPEC Announcements," *Available at: https://www.riccardo-degasperi.com/research*.

Gilchrist, Simon and Egon Zakrajšek (2012) "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, Vol. 102(4), pp. 1692–1720.

Känzig, Diego (2021) "The macroeconomic effects of oil supply news: Evidence from OPEC announcements," *American Economic Review*, Vol. 111(4), pp. 1092–1125.

Ludvigson, S., S. Ma, and S. Ng (2021) "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?" *American Economic Journal: Macroeconomics*, Vol. 13(4), pp. 369–410.

McCracken, Michael W. and Serena Ng (2016) "FRED-MD: A Monthly Database for Macroeconomic Research," *Journal of Business and Economic Statistics*, Vol. 34(4), pp. 574–589.

Miranda-Agrippino, S. and H. Rey (2020) "US Monetary Policy and the Global Financial Cycle," *Review of Economic Studies*, Vol. 87(6), pp. 2754–2776.

Miranda-Agrippino, Silvia and Giovanni Ricco (2023) "Identification with External Instruments in Structural VARs," *Journal of Monetary Economics*, Vol. 135, pp. 1–19.

Montiel Olea, J, J. Stock, and M. Watson (2021) "Inference in SVARs Identified with an External Instruments," *Journal of Econometrics*, Vol. 225 (1), pp. 74–87.

Mori, Lorenzo and Gert Peersman (2024) "Estimating the Macroeconomic Effects of Oil Supply News," Mimeo.

Timmermann, A. (1993) "How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices," *Quarterly Journal of Economics*, Vol. 108, pp. 1135–1145.