Supplemental Appendix: Cleansing Oil

Rui Sousa*

Department of Economics,
Northwestern University

June 7, 2026

A Empirical Validation of Oil Shock Identification

Figure A.1 compares the impulse responses functions generated from estimating equation (1) for the real WTI oil price and OECD+6 industrial production index using both the oil news shock of Känzig (2021) and the structural oil supply shock of Baumeister and Hamilton (2019), both normalised to a 10% increase in oil prices on impact. The two series generate virtually identical price dynamics: oil prices rise sharply on impact and decay to zero within eight months, so agents face the same price signal regardless of which shock is used. The two series diverge, however, in their implied activity effects. The news shock delivers the expected contractionary response: industrial production falls on impact and remains depressed for the full two-year horizon. The structural supply shock, by contrast, yields a positive and statistically significant increase in industrial production over the 2002–2025 sample — the opposite of what theory and the broader literature predict.

*Contact: *rui.agm.sousa@u.northwestern.edu.*

Figure A.1: Oil News Shock vs. Structural Supply Shock: Price and Activity Responses

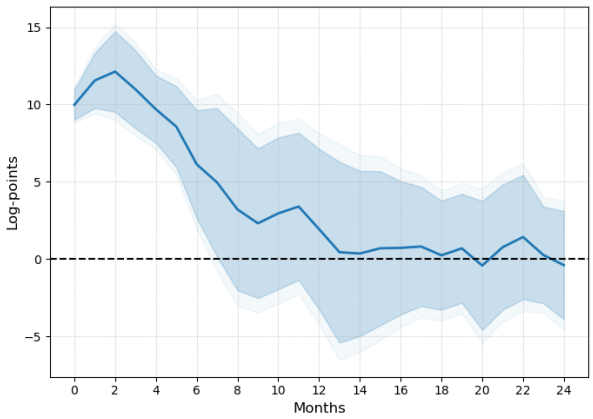

(a) Real WTI — Oil News Shock

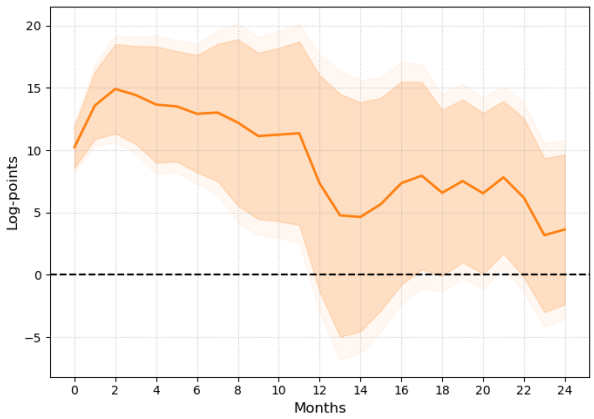

(b) Real WTI — Structural Supply Shock

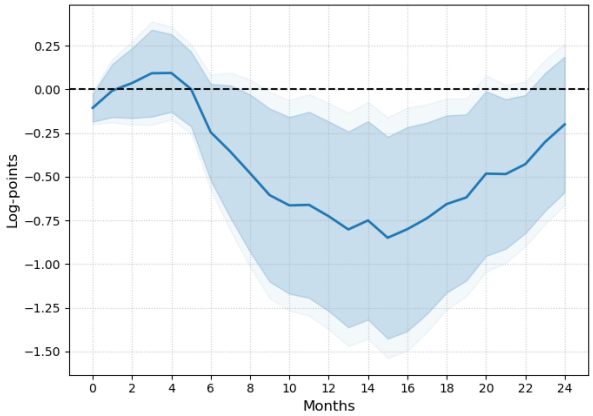

(c) Industrial Production — Oil News Shock

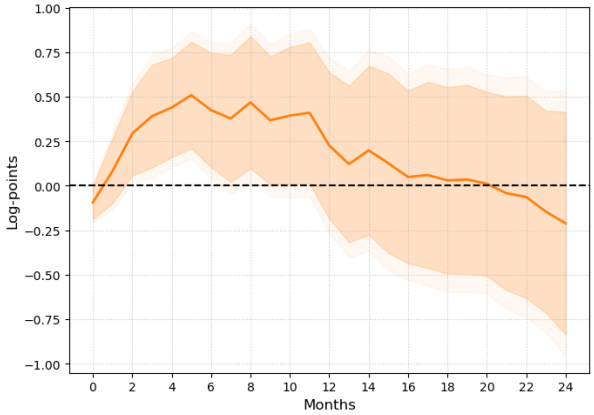

(d) Industrial Production — Structural Supply Shock

Notes: Impulse responses to each shock series, normalised to a 10% increase in oil prices on impact. Left column: oil supply news shock of Känzig (2021). Right column: structural oil supply shock of Baumeister and Hamilton (2019). Top row: real WTI oil price (log-points). Bottom row: OECD+6 industrial production (log-points). Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded. Sample: January 2002 – December 2025.

There are many reasons why this could be happening. One is the contamination of the shocks, as argued by Mori and Peersman (2024). Nevertheless, including more financial controls does not solve the problem. When using a shorter sample starting in 2010, consistent with the panel data, the problem becomes less severe, though the response is more delayed than what has been documented in the literature. The results are plotted in figure A.2. The oil news shock, by contrast, generates coherent dynamics across the samples used in the main analysis: starting in 2002, when most time-series begin, or in 2010, when panel coverage broadens significantly, the shock induces the same response in industrial production and oil prices, albeit with different persistence. I

therefore use the oil news shock throughout to ensure consistency.

Figure A.2: Response of Industrial Production and Oil Prices to the Oil News Shock — 2010-sample

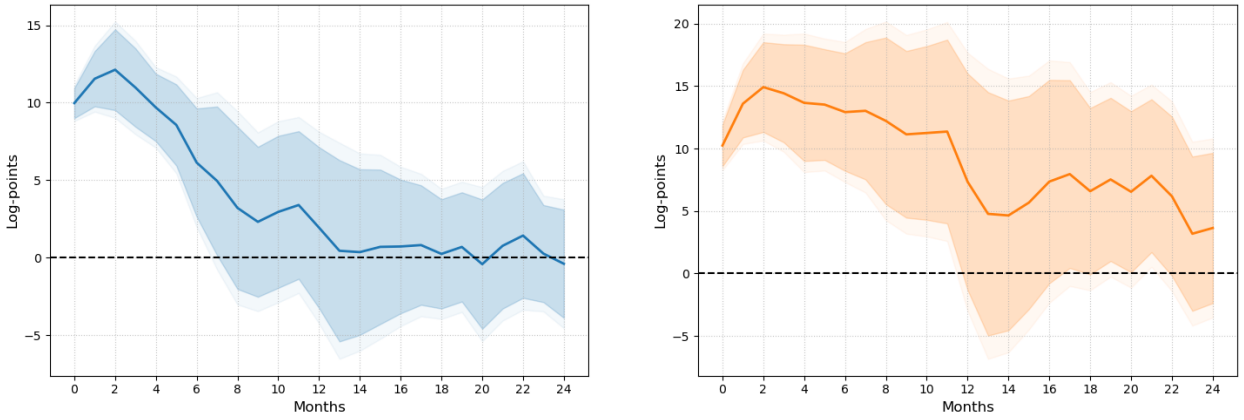

(a) WTI Price

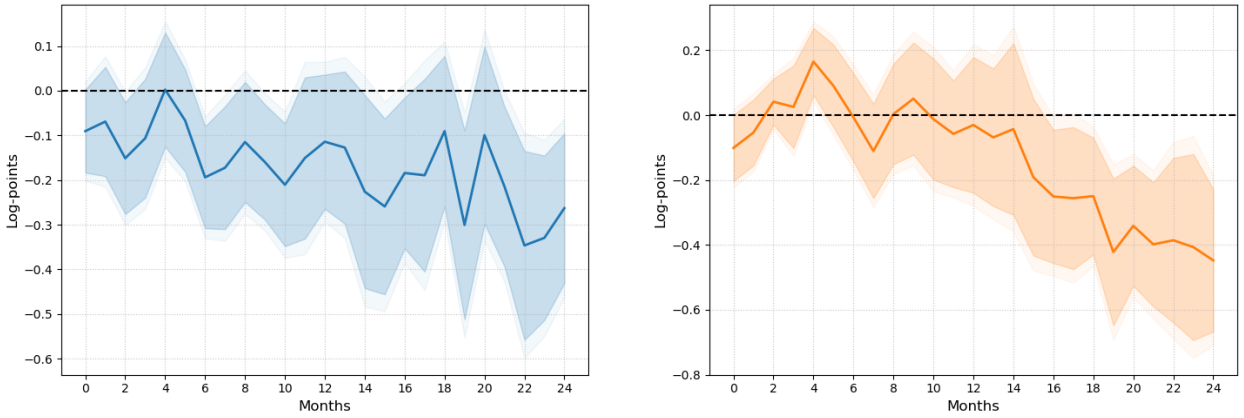

(b) Industrial Production

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. Left panel: OECD+6 Industrial Production (log-points). Right panel: Real WTI oil price (log-points). Each row restricts the sample to the labelled start year. Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded from all samples.

A.1 Construction of the CO_2-Equivalent Concentration Index

The time-series analysis uses a CO_2-equivalent (CO_2e) concentration index that aggregates the atmospheric burden of four major greenhouse gases—CO_2, CH_4, N_2O, and SF_6—into a single, radiatively consistent metric. Monthly global-mean concentrations of CO_2 (ppm), CH_4 (ppb), N_2O (ppb), and SF_6 (ppt) are obtained from NOAA.

Radiative forcing. Each gas contributes to total radiative forcing (RF) relative to its pre-industrial (1750) reference concentration. I use the simplified stratospheric-temperature-adjusted RF (SARF) expressions of Etminan et al. (2016) as adopted in Intergovernmental Panel on Climate Change (IPCC), ed (2023). Let C, M, N, S denote current concentrations of CO_2 (ppm), CH_4 (ppb), N_2O (ppb), and SF_6 (ppt), with pre-industrial reference values $C_0 = 278.3$, $M_0 = 729.2$, $N_0 = 270.1$, $S_0 = 0$. The individual SARF terms (W m^{-2}) are:

$$\Delta F_{\mathrm{CO_2}}(C, N) = \left[\alpha(C) + c_1\sqrt{N}\right] \ln\!\left(\tfrac{C}{C_0}\right) \cdot e_1, \qquad \alpha(C) = d_1 + a_1(C - C_0)^2 + b_1(C - C_0) \tag{1}$$

$$\Delta F_{\mathrm{CH_4}}(M, N) = \left[a_2\sqrt{M} + b_2\sqrt{N} + d_2\right]\left(\sqrt{M} - \sqrt{M_0}\right) \cdot e_2 \tag{2}$$

$$\Delta F_{\mathrm{N_2O}}(C, M, N) = \left[a_3\sqrt{C} + b_3\sqrt{N} + c_3\sqrt{M} + d_3\right]\left(\sqrt{N} - \sqrt{N_0}\right) \cdot e_3 \tag{3}$$

$$\Delta F_{\mathrm{SF_6}}(S) = \omega\,(S - S_0) \tag{4}$$

The scalars e_1, e_2, e_3 are the stratospheric temperature adjustment factors that convert instantaneous RF to SARF. Total forcing is $\Delta F_{\mathrm{total}} = \sum_g \Delta F_g$. Parameter values are listed in table A.1.

Table A.1: SARF parameter values (Etminan et al. 2016; Intergovernmental Panel on Climate Change (IPCC), ed 2023). †Units of W m^{-2} ppt^{-1}.

Gas	Parameter	Value
CO_2	a_1	-2.48×10^{-7}
	b_1	7.59×10^{-4}
	c_1	-2.15×10^{-3}
	d_1	5.2488
	e_1 (adj. factor)	1.05
CH_4	a_2	-8.96×10^{-5}
	b_2	-1.25×10^{-4}
	d_2	0.045194
	e_2 (adj. factor)	0.86
N_2O	a_3	-3.42×10^{-4}
	b_3	2.546×10^{-4}
	c_3	-2.44×10^{-4}
	d_3	0.12173
	e_3 (adj. factor)	1.07
SF_6	$\omega^\dagger$	5.67×10^{-4}

CO_2-equivalent concentration. Total forcing ΔF_{total} is inverted to a CO_2e concentration C_{eq} defined as the level of CO_2 that, holding N_2O at its pre-industrial value N_0, would produce the same forcing as all four gases combined:

$$\Delta F_{CO_2}(C_{\text{eq}},\, N_0) = \Delta F_{\text{total}}. \tag{5}$$

Since equation (1) is not analytically invertible in C, C_{eq} is solved numerically each period using Brent's method over the interval $[C_0,\, 5000]$ ppm.

B Panel Data Description

The empirical analysis draws on six main data sources. Table B.1 summarises the key characteristics of each dataset. For time-series variables, the table reports the sample span and number of non-missing monthly observations. For panel variables it additionally reports the number of countries contributing at least one non-missing observation over the sample period. All samples are restricted to January 2002 onwards to coincide with the availability of the oil supply news shock series.

Electricity generation and CO_2 emissions come from Ember, a monthly panel covering over eighty countries from the early 2000s onwards. Wind and solar installed capacity, also from Ember, is available for a narrower sample of twenty-six countries starting in 2016. Industrial production at the country level is sourced from the IMF, covering the OECD and selected emerging economies. Oil product consumption data (gasoline, diesel, jet fuel, LPG, and residual fuel) are from the Joint Organisations Data Initiative (JODI), as is natural gas consumption by the power sector. Coal production at the country level comes from the International Energy Agency (IEA). Coverage across these datasets is unbalanced: the number of countries reporting data varies across variables and over time, as documented in Figures B.1 to B.5.

Table B.1: Data Sources and Summary Statistics

Dataset	Source	Countries	Start	End	Obs	Min	Mean	Median	Max	Std
CO_2e Concentration	NOAA	—	2002	2025	288	416.8	451.6	449.6	491.6	21.0
Real WTI Oil Price	FRED	—	2002	2025	287	0.2	0.8	0.7	1.7	0.3
Oil News Shock	KZ	—	2002	2024	276	-3.1	0.0	-0.0	1.9	0.6
Oil Supply Shock	BH	—	2002	2024	276	-1.9	0.1	0.1	5.3	0.6
Ind Prod	BH	—	2002	2025	288	60.5	86.0	86.4	110.7	13.6
Global Fin Cycle	MR	—	2002	2024	276	-2.9	0.2	0.1	2.6	1.0
Financial Uncertainty	JLN	—	2002	2025	288	0.7	1.0	0.9	1.4	0.1
Electricity Gen	Ember	87	2002	2025	10301	0.0	29.8	5.1	1015.2	93.3
Wind & Solar Capacity	Ember	25	2016	2025	2874	0.2	67.4	18.1	2082.1	181.5
(IMF) Ind Prod	IMF	65	2002	2025	16548	32.3	107.1	103.8	450.7	25.4
Oil Products	JODI	103	2002	2025	23768	-202.0	7790.8	1242.2	306022.0	31202.4
Natural Gas	JODI	77	2009	2025	12059	-1.0	2961.8	628.0	111046.0	8886.5
Coal Production	IEA	33	2016	2025	1188	0.0	65121.0	7768.0	1283191.0	183687.0

Notes: Sample restricted to years $\geq$ 2002. COVID-19 (2020) is excluded from regressions but included in the statistics above. Obs. counts reflect non-missing observations for the key variable in each row. For panel variables, *Countries* reports the number of countries with at least one non-missing observation. KZ: Känzig (2021); BH: Baumeister and Hamilton (2019); MR: Miranda-Agrippino and Rey (2020); JLN:Jurado et al. (2015).

Figure B.1: Country Coverage — Electricity Generation & Capacity

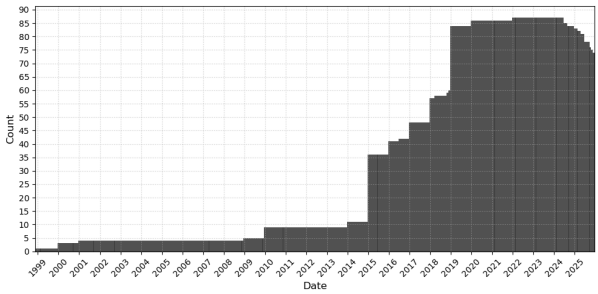

(a) Total Generation (TWh)

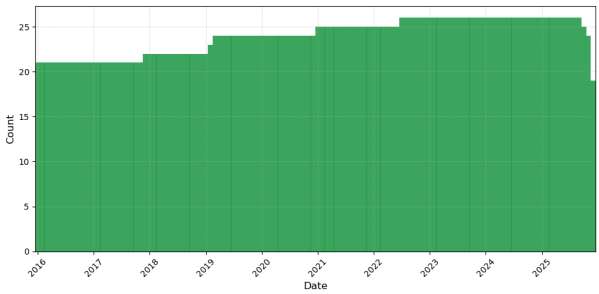

(b) Wind & Solar Capacity (GW)

Notes: Each bar shows the number of countries with non-missing data in that month. Source: Ember.

Figure B.2: Country Coverage — JODI Oil Products

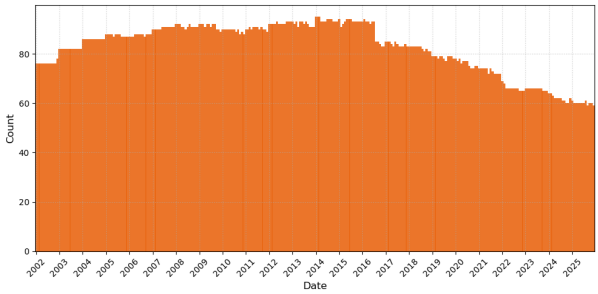

(a) Gasoline

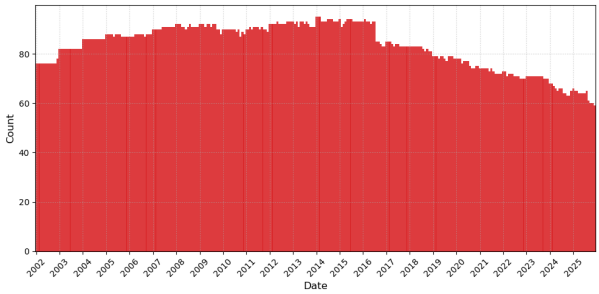

(b) Diesel

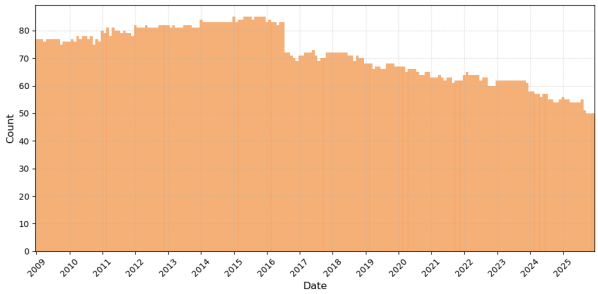

(c) Jet Fuel / Kerosene

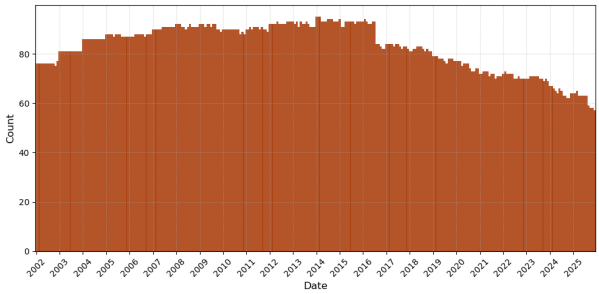

(d) Residual Fuel

Notes: Each bar shows the number of countries with non-missing data in that month. Source: JODI.

Figure B.3: Country Coverage — Natural Gas & Coal

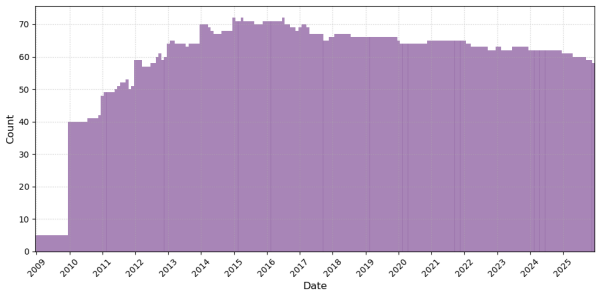

(a) Natural Gas — Total (JODI)

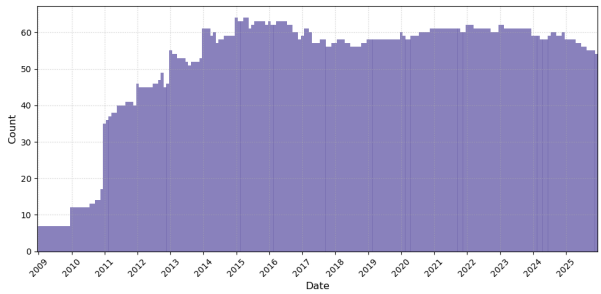

(b) Natural Gas — Power Sector (JODI)

Notes: Each bar shows the number of countries with non-missing data in that month. Source: JODI.

Figure B.4: Country Coverage — Coal Production

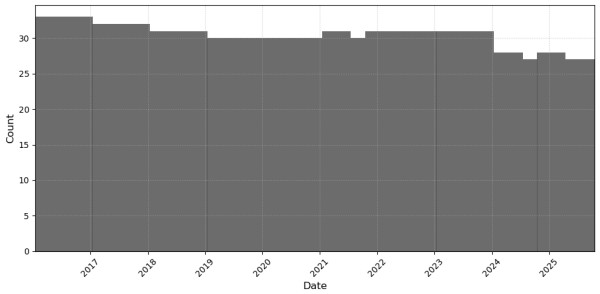

(a) Total Coal

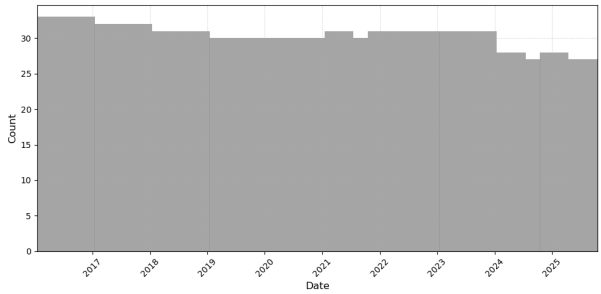

(b) Lignite

Notes: Each bar shows the number of countries with non-missing data in that month. Source: IEA.

Figure B.5: Country Coverage — Vehicle Registrations

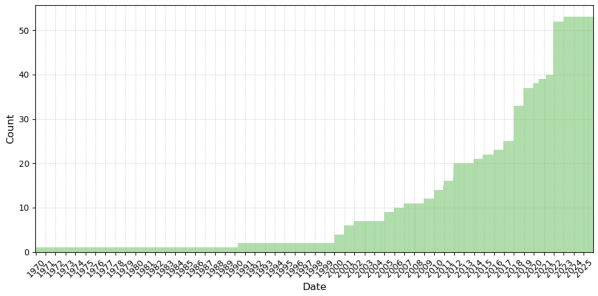

(a) Non-ICE Vehicles

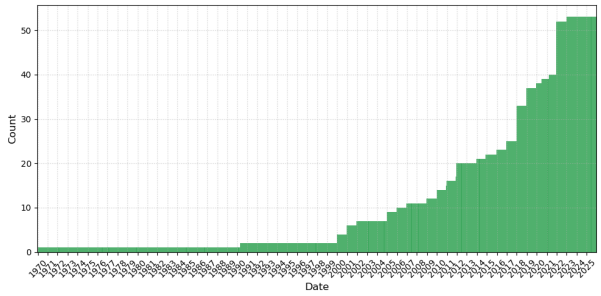

(b) Total Registrations

Notes: Each bar shows the number of countries with non-missing data in that month.

B.1 Ember–EIA Comparison

To validate the Ember dataset, I compare it against the annual international statistics published by the US Energy Information Administration (EIA). For each country and variable present in both datasets, I aggregate Ember's monthly data to an annual total and compare it against the corresponding EIA annual figure. Figure B.6 plots this comparison for electricity generation by fuel type, and figure B.7 extends the comparison to installed wind and solar capacity.

Figure B.6: Ember vs. EIA Annual Electricity Generation

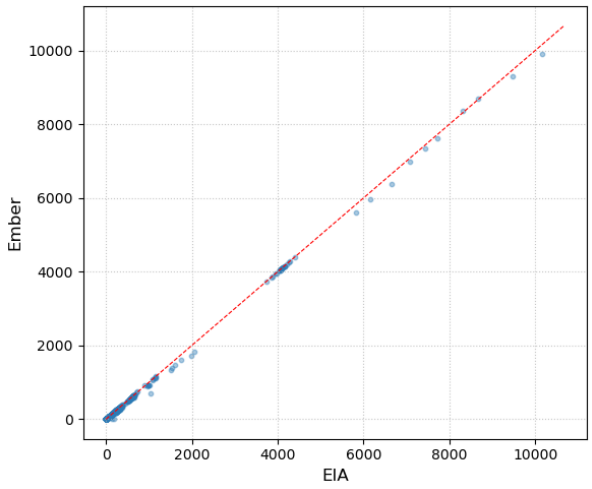

(a) Total Generation

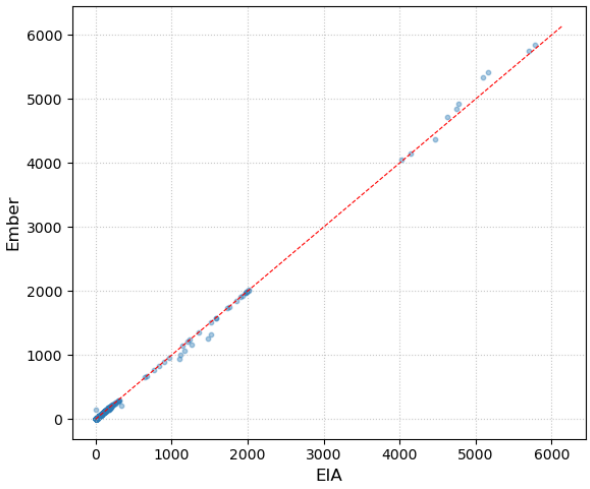

(b) Coal

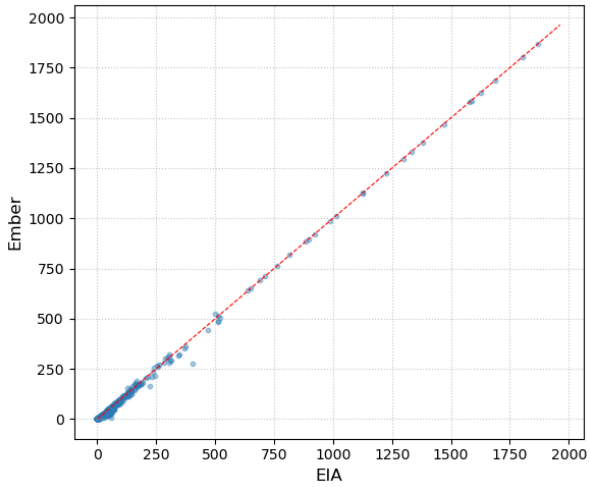

(c) Gas

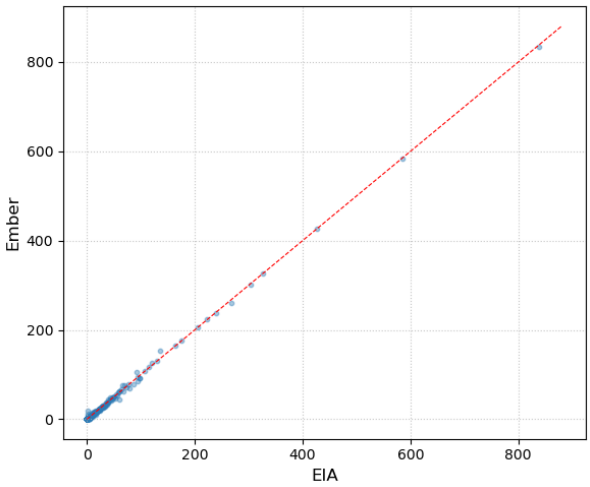

(d) Solar

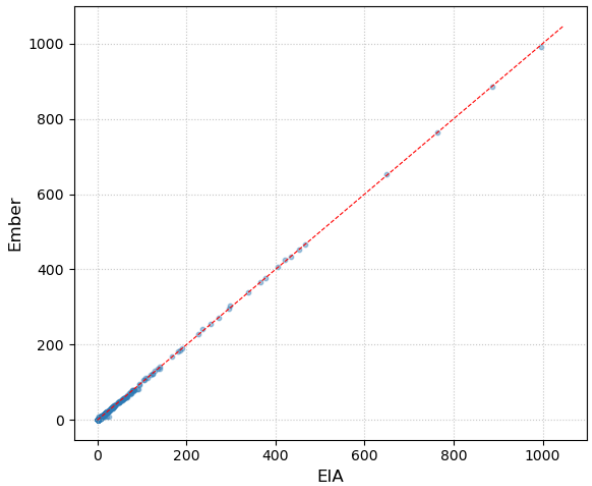

(e) Wind

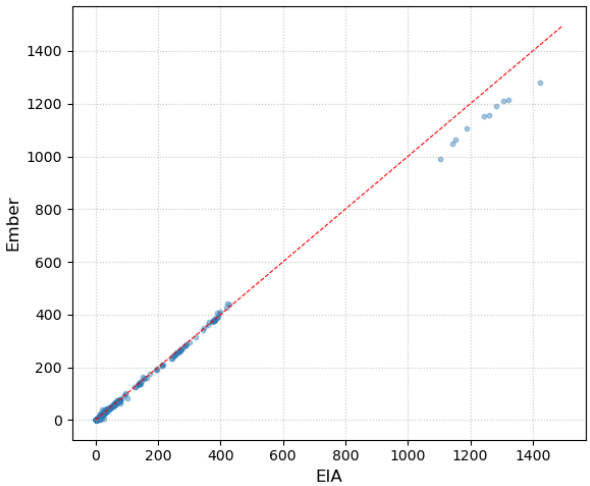

(f) Hydro

Notes: Each point represents one country-year observation. Ember monthly generation is aggregated to annual totals and compared against EIA annual reporting. Both axes in TWh.

Figure B.7: Ember vs. EIA Annual Installed Capacity

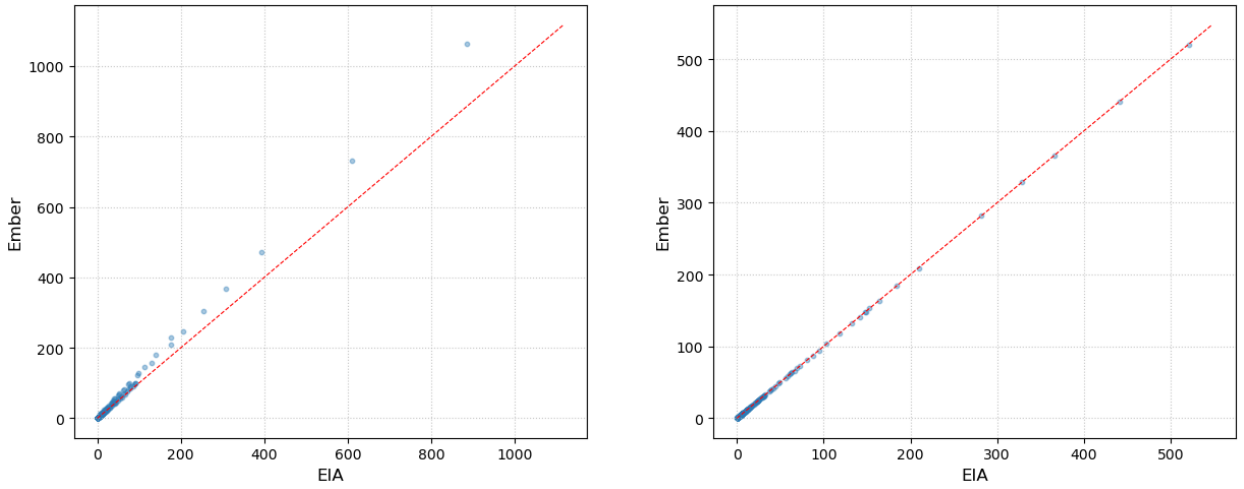

(a) Solar Installed Capacity (GW)

(b) Wind Installed Capacity (GW)

Notes: Each point represents one country-year observation. Ember annual capacity is compared against EIA annual reporting. Both axes in GW.

C Panel Analysis Robustness Checks

The baseline panel specification includes 12 lags of both the outcome variable and the oil shock series ($P = K = 12$). While this lag structure absorbs serial dependence in the outcome and controls for delayed shock propagation, it comes at a cost: the first 12 months of each country's time series are consumed by the lag block, reducing the effective sample size. The global controls — world financial cycle, financial uncertainty, and the exchange rate — are available for the full sample and do not drive this reduction; the loss is attributable entirely to the lags of the outcome and the shock. For datasets with short country-level coverage — such as coal production and installed capacity — this loss is non-trivial.

This section reports results from an alternative specification that sets $P = 0$ and $K = 12$ and instead includes a country-specific linear time trend. Dropping outcome lags recovers the full sample for every country, while the 12 lags of the oil shock series are retained alongside the global controls — world financial cycle, financial uncertainty, and the exchange rate — to preserve delayed shock propagation. The country-specific trend absorbs differential long-run growth paths in each outcome, which is important given that countries differ substantially in their stage of energy transition. Country–month fixed effects are retained throughout. The trade-off is twofold. First, without outcome lags, the specification is less able to absorb high-frequency serial correlation in the dependent variable; standard errors remain Driscoll–Kraay to account for cross-sectional dependence. Second, a country-specific linear trend captures growth paths conditional

on the country's observed sample, but does not fully account for differential trends driven by when countries enter the panel. Countries that join later may be at a systematically different stage of their energy transition, so the linear trend partially conflates genuine trend heterogeneity with entry-time heterogeneity. Reassuringly, most results closely mirror the baseline. One notable exception is wind installed capacity, which increases in this specification — a result that sits in tension with wind generation, which decreases — suggesting this particular finding is not robust and should be interpreted with caution. Overall, the main findings are not an artefact of the particular specification chosen.

Coal

Figure C.1: Response of Coal to an Oil Supply News Shock — Robustness 3

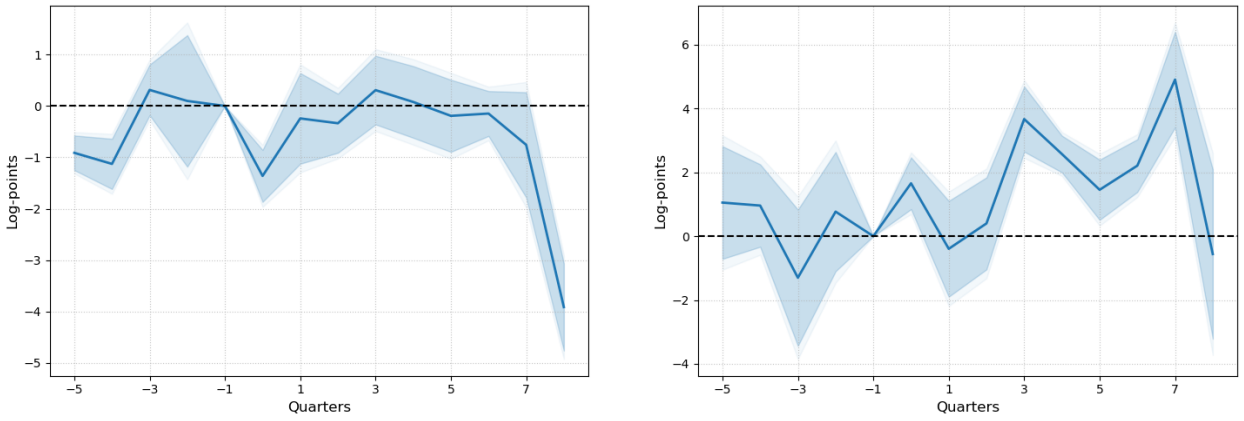

(a) Steam and Coking Coal Production

(b) Lignite Production

Notes: Conventions as in figure 6, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

Transportation

Figure C.2: Response of Vehicle Registrations to an Oil Supply News Shock — Robustness 3

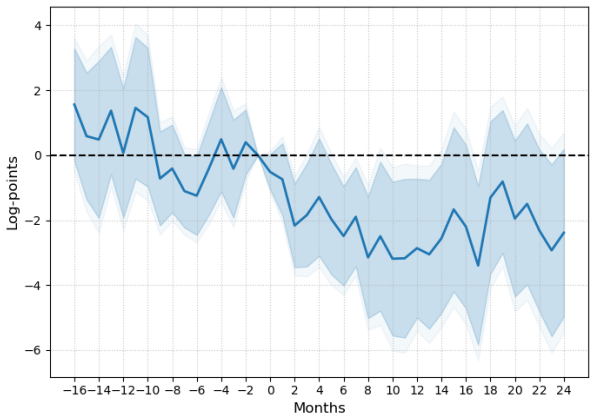

(a) ICE Registrations

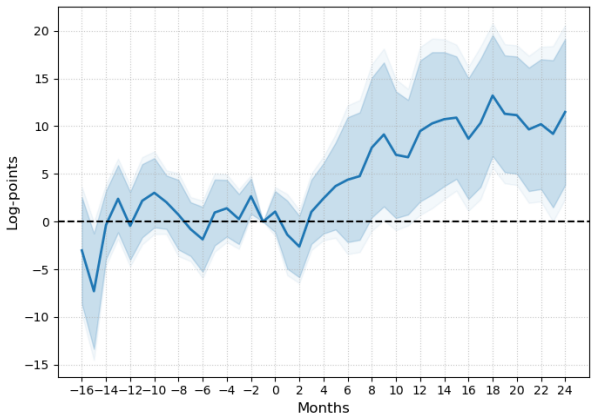

(b) Non-ICE Registrations

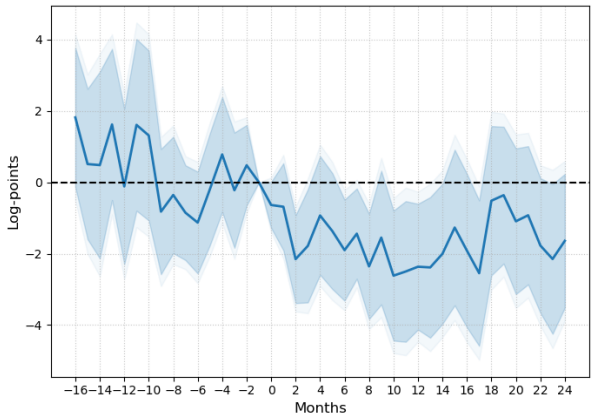

(c) Total Car Registrations

Notes: Conventions as in figure 7, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

The Power Sector

Figure C.3: Response of Total Electricity Generation to an Oil Supply News Shock — Robustness 3

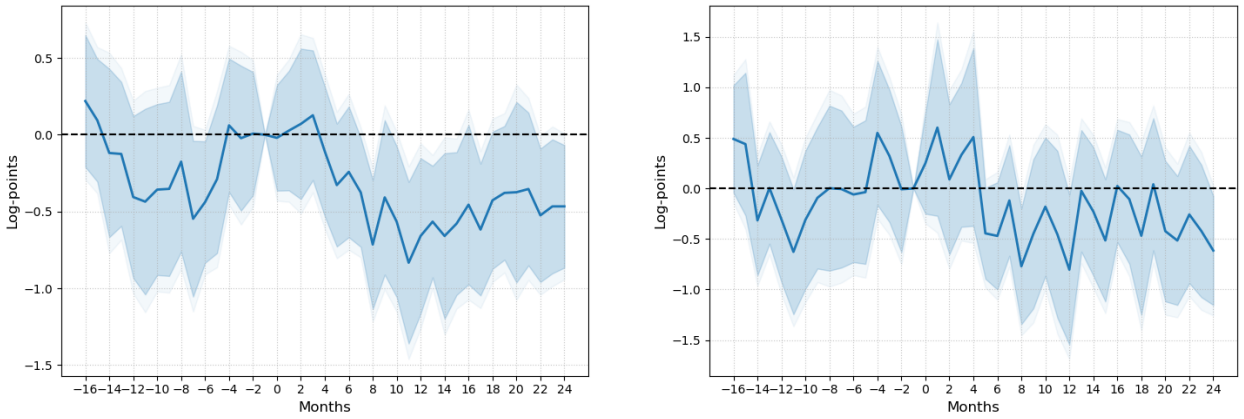

(a) Total Electricity Generation (TWh)

(b) Electricity Generation / Industrial Production

Notes: Conventions as in figure 8, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

Figure C.4: Response of Renewable Electricity Generation by Source to an Oil Supply News Shock — Robustness 3 (TWh)

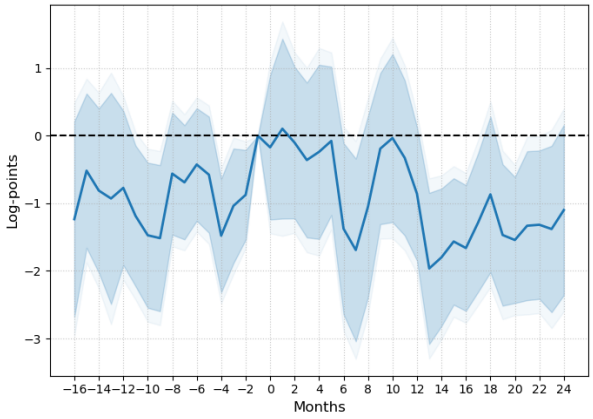

(a) Renewables x-Nuclear

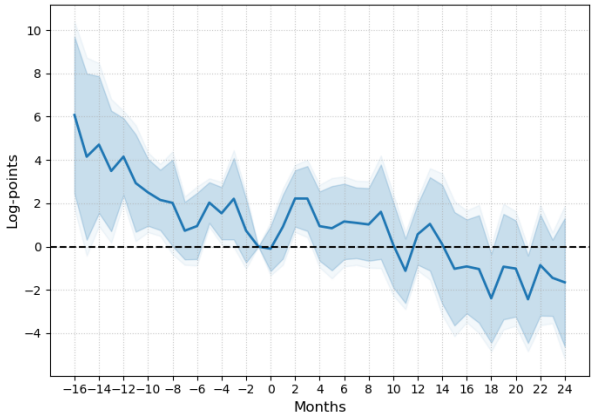

(b) Solar

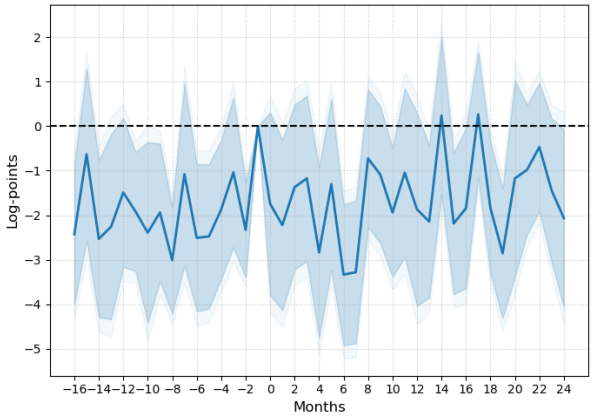

(c) Wind

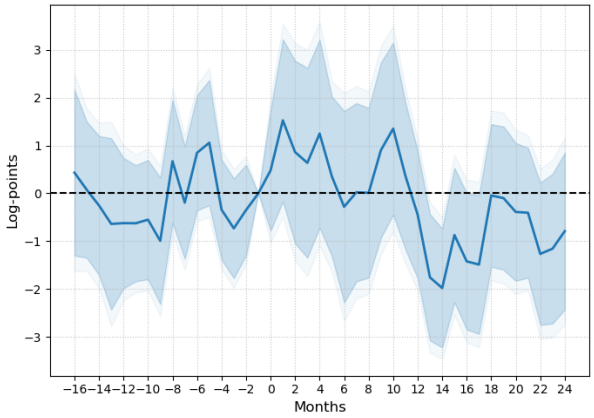

(d) Hydro

Notes: Conventions as in figure 9, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

Figure C.5: Response of Renewable Installed Capacity to an Oil Supply News Shock — Robustness 3

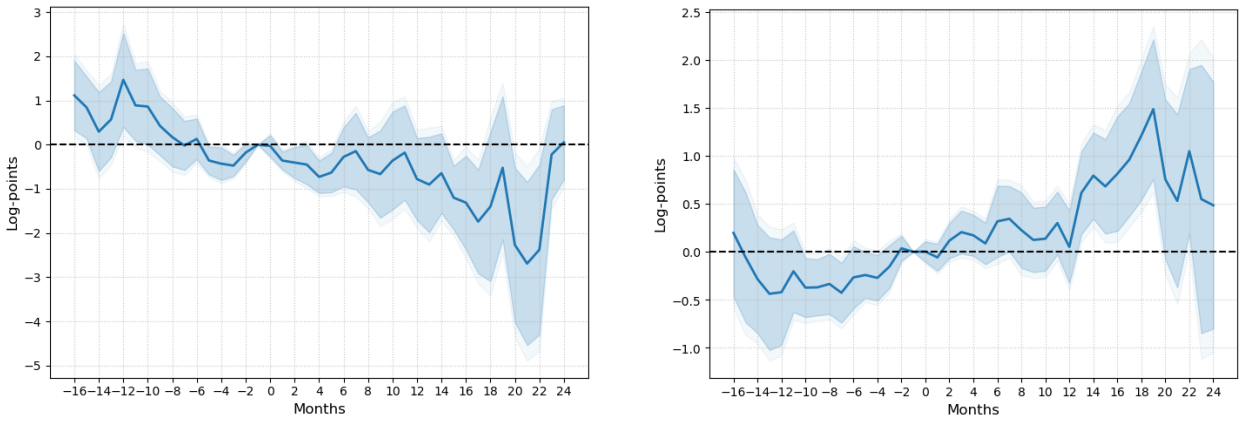

(a) Solar Installed Capacity (GW)

(b) Wind Installed Capacity (GW)

Notes: Conventions as in figure 10, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

Figure C.6: Response of Bioenergy Generation to an Oil Supply News Shock — Robustness 3 (TWh)

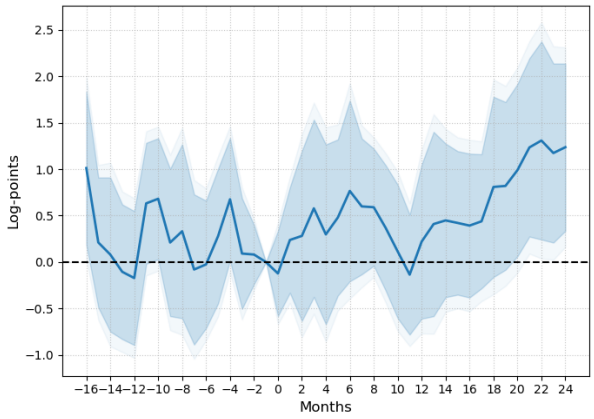

(a) Bioenergy Generation (TWh)

Notes: Conventions as in subsection 4.2, except that country-specific linear time trends replace the 12 lags of the outcome variable ($P = 0$) while retaining 12 shock lags ($K = 12$), recovering the full sample for outcome lags.

References

Baumeister, Christiane and James D. Hamilton, "Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks," *American Economic Review*, May 2019, *109* (5), 1873–1910.

Etminan, M., G. Myhre, E. J. Highwood, and K. P. Shine, "Radiative forcing of carbon dioxide, methane, and nitrous oxide: A significant revision of the methane radiative forcing," *Geophysical Research Letters*, 2016, *43* (24), 12,614–12,623. _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2016GL071930.

Intergovernmental Panel on Climate Change (IPCC), ed., "The Earth's Energy Budget, Climate Feedbacks and Climate Sensitivity," in Intergovernmental Panel on Climate Change (IPCC), ed., *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press, 2023, pp. 923–1054.

Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, "Measuring Uncertainty," *American Economic Review*, March 2015, *105* (3), 1177–1216.

Känzig, Diego R, "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements," *American Economic Review*, 2021, *111* (4).

Miranda-Agrippino, Silvia and Hélène Rey, "U.S. Monetary Policy and the Global Financial Cycle," *The Review of Economic Studies*, November 2020, *87* (6), 2754–2776.

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