Cleansing Oil[*]

Rui Sousa[†]

Department of Economics,
Northwestern University

June 7, 2026

Abstract

I study the dynamics of greenhouse gas emissions following unexpected increases in oil prices from supply disruptions. I find that CO_2-equivalent atmospheric concentrations decrease by -0.04 ppm relative to trend following a 10% increase in crude oil prices. This effect is heterogeneous across gases, reflecting a reshuffling of fossil fuel usage. Natural gas and petroleum products fall sharply across the board, while low-quality coal and bioenergy generation rise. Driven by tighter financial conditions and elevated uncertainty, solar capacity contracts. Hybrid and EV adoption increase markedly, partially substituting for conventional internal combustion engine vehicles.

1 Introduction

Oil supply disruptions are widely believed to accelerate the green transition. Whenever geopolitical tensions threaten global crude oil supplies—whether through wars in the Middle East, attacks on energy infrastructure, or concerns surrounding the closure of the Strait of Hormuz—commentators routinely argue that higher fossil fuel prices will encourage the adoption of cleaner technologies and reduce greenhouse gas emissions. This paper shows that this view is only partly correct.

I find that a 10 percent increase in oil prices caused by an exogenous oil supply shock reduces global greenhouse-gas concentrations and increases emissions efficiency. However, the adjustment does not occur through the channel most commonly emphasized in public discussions.

---

[*] I am grateful to my advisors, Martin Eichenbaum, Diego Känzig and Giorgio Primicieri. This work was conducted under the financial support of Fundação para a Ciência e a Tecnologia. All errors are my own.

[†] Contact: *rui.agm.sousa@u.northwestern.edu.*

While consumers substitute away from internal-combustion-engine vehicles toward cleaner transportation technologies, renewable electricity generation—and solar power in particular—contracts. Oil supply disruptions reduce emissions, but they do not accelerate the clean-energy transition in the power sector.

The reason is that oil supply disruptions trigger a broad reallocation across fuels rather than a simple substitution away from fossil energy. Consumption of petroleum products falls sharply, especially jet fuel, gasoline, and diesel. Natural gas consumption declines as well. At the same time, lignite production and bioenergy generation increase. These adjustments generate heterogeneous greenhouse-gas dynamics. Atmospheric concentrations of carbon dioxide and methane decline, whereas nitrous oxide concentrations rise persistently. The reduction in overall warming potential therefore masks important changes in the composition of emissions and in the underlying energy mix.

The contraction in renewable deployment is particularly surprising. Conventional reasoning suggests that higher fossil-fuel prices should improve the competitiveness of clean alternatives and stimulate their adoption. Indeed, much of the discussion surrounding oil disruptions implicitly assumes that expensive oil acts as a catalyst for the green transition. Yet oil supply shocks are not carbon taxes. The same disturbance that raises the price of fossil fuels also depresses economic activity, increases uncertainty, tightens financial conditions, and disrupts complementary energy markets. These forces work against renewable investment and can dominate the standard relative-price channel.

The distinction matters for how oil disruptions should be viewed through a climate lens. An oil shock delivers a partial carbon-price signal, but it also generates distortions that a carbon tax does not. The resulting environmental consequences are therefore more nuanced than commonly assumed. Oil-intensive activities contract and transportation becomes cleaner, yet renewable deployment weakens and electricity generation shifts toward lower-quality combustion alternatives such as lignite and bioenergy.

To study these effects, I bring the empirical tools of modern macroeconomics to the study of climate change and energy transitions. Combining identified oil supply shocks with global monthly data on greenhouse-gas concentrations, fossil-fuel consumption, electricity generation, renewable capacity, and vehicle registrations, I estimate the dynamic responses of emissions and energy use to exogenous disruptions in oil markets. The analysis spans the period from 2002 to 2024 and traces how oil supply shocks propagate through the energy system—from fuel consumption and electricity generation to technology adoption and ultimately global greenhouse-gas concentrations.

1.1 Related Literature

This paper contributes to five strands of literature. First, it relates to the large literature studying the macroeconomic consequences and transmission channels of oil-market shocks. Beginning with Hamilton (1983), a vast literature has documented that oil-price increases are associated with declines in economic activity and changes in inflation, investment, and financial conditions. More recent contributions have emphasized the importance of distinguishing between different sources of oil-price fluctuations. For example, Kilian (2009) show that supply, aggregate-demand, and precautionary-demand shocks have markedly different macroeconomic consequences, while Baumeister and Hamilton (2019) develop a Bayesian framework to incorporate prior knowledge about the different disturbances. Building on techniques typically used in the analysis of monetary policy, Känzig (2021) uses high-frequency information around OPEC announcements to identify exogenous oil-supply news shocks and study their macroeconomic effects. I leverage this literature for identification, but shift the focus towards the environmental consequences of oil shocks.

Second, the paper contributes to the literature studying the response of fossil-fuel demand to energy price changes. Much of the attention in this literature has been devoted to gasoline markets—in particular, to estimating the price elasticity of gasoline demand. A growing body of work documents how this elasticity has evolved over time and varies across households and regions, with Kilian and Zhou (2024) providing a recent account using cross-state variation in oil-to-gasoline pass-through as a new instrument. Far less attention has been devoted to other petroleum derivatives and primary fuels: diesel, jet fuel, liquefied petroleum gas, natural gas, and coal each face distinct demand structures and substitution margins that a gasoline-centric lens cannot capture. The importance of this broader coverage is underscored by recent events: the threatened closure of the Strait of Hormuz in 2025 raised acute concerns about the supply of jet fuel for global aviation and bunker oil for international maritime trade. I contribute by documenting the response of the full fossil-fuel spectrum—gasoline, diesel, jet fuel, liquefied petroleum gas, natural gas, and coal—to exogenous oil-supply shocks.

Third, this paper relates to the literature studying the relationship between greenhouse-gas emissions and macroeconomic fluctuations. Heutel (2012) establishes that carbon emissions are deeply procyclical in the United States. Using a cross-section of countries and annual data, Doda (2014) confirms this procyclicality internationally and further documents that emissions are more volatile than output. Closer to my work, Khan et al. (2019) provide a first attempt at identifying the causal effects of different macroeconomic shocks — including oil-price and technology shocks — on U.S. emissions, but do not find statistically significant effects for any of them. I document significant effects of oil-supply shocks on emissions, with the added value of conducting a worldwide analysis — important for a global problem like climate change.

Fourth, the paper contributes to the macroeconomic literature on energy efficiency, directed technical change, and capital structures. In their foundational work, Hassler et al. (2021) show that the short-run substitution elasticity between energy and capital/labor is virtually zero (near-Leontief), but that sustained energy price increases direct innovation toward energy-saving technologies, rendering long-run substitutability significantly higher. Using time-series methods, Känzig and Williamson (2024) isolate these technological drivers, finding that while energy price shocks successfully stimulate energy-saving innovation, they are highly recessionary and crowd out other productive innovations. Conversely, pure energy-saving technology shocks drive a non-recessionary decoupling of output and emissions. Furthermore, Casey (2024) introduces a "putty-clay" model of directed technical change, noting that energy efficiency is permanently embedded in physical capital ex-post. Casey (2024) argues that standard Cobb-Douglas models drastically overstate the long-run emissions reductions of simple energy taxes due to rebound effects that discourage subsequent R&D, emphasizing the need for targeted R&D policy. I complement this literature by studying whether oil-price shocks improve emissions efficiency at the global level and by tracing the underlying mechanisms through detailed adjustments in fossil-fuel use, transportation demand, and electricity generation. I show that while an oil-price shock improves emissions efficiency per unit of industrial production, this is a short-lived intensive margin adjustment that lacks the permanent, non-recessionary benefits of true technological decoupling.

Fifth, this paper contributes to the literature studying the adoption of clean-energy technologies and the structural and financial frictions that govern the green transition. While conventional theory suggests that fossil-fuel price spikes accelerate renewable adoption, macro-financial and technological frictions can disrupt this channel. First, renewable projects are highly capital-intensive and face upfront expenditures, rendering them hypersensitive to the cost of capital (Bistline et al., 2023). Inflationary energy price shocks trigger monetary policy tightening; because renewables are far more capital-intensive than their fossil-fuel substitutes such as natural gas, the resulting rise in the cost of capital creates a differential financial barrier to clean energy adoption. The financial headwinds extend to the innovation frontier: Aghion et al. (2024) document that contractionary monetary policy reduces green patenting significantly more than non-green patenting, as tight credit disproportionately constrains the young firms at the forefront of clean innovation. This contrasts with the finding of Känzig et al. (2025) that green innovation is broadly countercyclical—driven by a "green is in the future" channel that insulates long-term R&D from short-term recessions. Second, grid-level technological constraints create interdependencies across fuels. Verdolini et al. (2018) document a strong complementarity between intermittent wind and solar and fast-reacting fossil-based backup capacity (primarily natural gas). Consequently, when oil shocks simultaneously disrupt natural gas supply and spike gas prices, the adoption of solar is severely hindered. Indeed, the relationship between natural gas and renewables is complex; while the shale gas boom reduced emissions by displacing coal, Acemoglu et al. (2023) and Harstad and Holtsmark (n.d.) argue that cheap natural gas can create a ”fossil-fuel trap” that delays clean innovation, whereas Lindequist and Selent (2025) find evidence of gas acting as a green ”bridge fuel”. I contribute to this literature by documenting a stark bifurcation in the substitution channel: while high oil prices smoothly accelerate consumer transportation substitution (EV adoption), they paradoxically contract physical solar power generation and installed capacity due to the twin forces of monetary-policy-induced financing costs and natural gas supply complementarities.

1.2 Organization

The remainder of this paper is organized as follows. First, section 2 studies the implications of oil price hikes on worldwide greenhouse gas concentrations. Section 3 documents the dynamics of a broad spectrum of fossil fuels. Section 4 in turn, details the mechanisms that underlie changes in energy efficiency — particularly focusing on the extensive margin — the adoption of clean energy alternative. Section 5 concludes.

2 Global Warming Potential of Oil Price Shocks

The effect of oil price shocks on greenhouse gas emissions is non-trivial. An increase in oil prices may redirect energy usage towards more pollutant sources, such as coal, but it may also lead to increased energy efficiency and cleaner sources such as solar. In this section I evaluate the overall emissions effects of oil price shocks using standard macroeconomic time-series techniques, and thus the global warming potential changes as a result.

2.1 Data

Emissions I use monthly average global atmospheric concentration data from the National Oceanic Atmospheric Administration (NOAA) to study the global warming effects of oil price shocks. NOAA provides measurements for four of the main greenhouse gases (ghg): carbon dioxide (CO_2), Nitrous Oxide (N_2O), Methane (CH_4) and Sulfur Tetrafluoride (SF_6). The data start at different dates, so I restrict my sample to the intersection of their spans. As a result, all data start in January 2002 and end in December 2024. Starting the sample in the early *XXI* century has the advantage of coinciding with the rise of renewables as viable energy sources, a significant development for the substitution possibilities of the world economy with respect to oil.

In order to measure the global warming potential of the combined gas concentrations added, I compute the Carbon Dioxide Equivalent (CO_2e) for the four ghgs. This individual index measures how much CO_2 would be necessary, all else equal, to increase the globe's temperature by the same amount as that resulting from the combined ghg emissions since pre-industrial times. I follow the formulas and parameters provided by Etminan et al. (2016) to first measure the Stratospheric-temperature-adjusted radiative forcing (SARF). Then, I fix all gases to their pre-industrial levels according, and solve for the necessary change in CO_2 to attain the same SARF. I provide more details on the parameters and algorithm used in subsection A.1.

Oil Price Shocks To capture exogenous variations in the price of crude oil, I use the oil supply news shocks series provided by Känzig (2021). To generate this series he first exploits the institutional features of the Organization of Petroleum Exporting Countries (OPEC): at pre-determined dates OPEC systematically provides an outlook on their own crude oil production. Using high-frequency oil futures data, he can then measure the changes in the price of crude oil futures contracts in a narrow window of time around these scheduled events. This high-frequency approach thus avoids reverse causality problems and isolates supply from demand-driven fluctuations in oil price futures. In a second step, using a standard proxy vector autoregression, he is able to identify exogenous shifts in the price of crude oil due to changes in the expected supply of global crude oil by applying instrumental variables to the residuals of the oil supply equation and using the news surprise series previously generated. The resulting estimates are the exogenous shock series, used widely in the macroeconomics literature. The series is normalized so that a one-unit shock corresponds to a 10% increase in the price of crude oil on impact.

Mori and Peersman (2024) document that this identification approach is subject to informational deficiencies: the identified structural innovations are found to Granger-caused by financial variables, indicating that markets can predict OPEC decisions through information not captured by the vector auto-regressive model, thereby contaminating the instrument. They show that augmenting the econometrician's information set with financial indicators resolves these issues and yields sharper, more stable estimates. I address this concern directly by including financial controls in my local projections, as described in subsection 2.2.

In principle, the structural supply shocks of Baumeister and Hamilton (2019) would be better suited for my purposes: unlike news shocks, they are identified to produce an immediate fall in oil production, representing a cleaner physical supply disruption rather than a shift in expectations over future production, matching my motivation. However, over the sample, from 2002 to 2024, these shocks yield a positive response on industrial production regardless of the controls included, contradicting both theory and the empirical literature, and rendering them unsuitable for this analysis[1]. The oil supply news shocks of Känzig (2021) do not suffer from this problem: they deliver the expected contractionary response, confirming that the identification remains valid over the sample period. Notwithstanding, they generate identical dynamics in the price of oil —

---

[1] See section A for a full account of this phenomenon.

the signal that agents respond to. I conduct most of the exercises using both series for robustness.

2.2 Methodology

To assess the effect of oil price shocks, I resort to the local projections methodology popularized by Jordà (2005). For $h = 0, \cdots, H$, I estimate the long-differences model:

$$\Delta^{h+1} y_{t+h} = \alpha_h + \alpha_{m(t)} + \beta_h \cdot s_t + \sum_{k=1}^{12} (\rho_h \Delta y_{t-k} + \Gamma_h^{'} X_{t-k} + \tau_h s_{t-k}) \qquad (1)$$

where α_h is an intercept, to account for the average growth a horizon h. $\alpha_{m(t)}$ are month fixed-effects used to account for the seasonality in both emissions and oil prices. $m(t)$ is the calendar month at date t. s_t is the exogenous oil price shock and X_{t-k} are controls variables. $\Delta^{h+1} y_{t+h} = y_{t+h} - y_{t-1}$ is the difference between the dependent variable y at horizon h relative to the month before the shock $t-1$, and its either the concentration of individual ghg in the atmosphere, or my measure of equivalent CO_2. All confidence bands are computed using heteroskedasticity and autocorrelation consistent (HAC) standard errors (Newey and West, 1987).

As discussed in the previous subsection, Mori and Peersman (2024) show that financial variables Granger-cause the oil supply news shocks, contaminating the instrument through anticipation effects. I address this by including 12 lags of the shock and of two financial controls: the global financial cycle of Miranda-Agrippino and Rey (2020), which captures broad movements in global risk appetite and cross-border financial conditions, and the financial uncertainty index of Jurado et al. (2015), which measures the unpredictability of financial variables, reflecting forward-looking information that market participants embed in asset prices ahead of OPEC decisions. I keep this parsimonious specification uniform across all empirical sections of the paper, including the panel analyses where the available time span is considerably shorter, so that results remain comparable throughout.

The sample includes the COVID pandemic period. To handle this, I follow Baumeister and Hamilton (2025) and drop the year of 2020 completely from the data. Because of the inclusion of lags, 2011 is also effectively dropped.

2.3 Results

2.3.1 The Macroeconomic Backdrop

Before turning to the emissions results, figure 1 documents the response of key macroeconomic aggregates to the shock. This serves two purposes: it validates the shock identification for the restricted sample, 2002 to 2024 — excluding 2020, and it establishes the macroeconomic context within which all results should be interpreted.

Figure 1 shows the real oil price, computed by deflating the West Texas Intermediate price by the US consumer price index, both obtained from FRED. It rises by 10% on impact — by construction — and gradually decays, becoming statistically indistinguishable from zero at around 8 months, with point estimates centring at zero one year after - much less persistent than what Känzig (2021) finds. These price dynamics are very similar to those induced by the supply shocks of Baumeister and Hamilton (2019) over the same sample, so the implied price signals are closely comparable across the two series. Absent non-linearities introduced by supply disruptions, the two thus imply equivalent oil price signals to which economic agents react to.

Figure 1 confirms that the shock is contractionary: industrial production for OECD and 6 selected countries[2], obtained from Baumeister and Hamilton (2019), falls on impact and remains depressed over the two-year horizon, consistent with a negative aggregate supply shock. The drop peaks after 14 months, implying an almost 0.85% drop in global economic activity. The maximum effect is almost the same as in Känzig (2021). Distinctively, it is much less persistent, in line with the dynamics of the wti.

Figure 1: Macroeconomic Responses to an Oil Supply News Shock

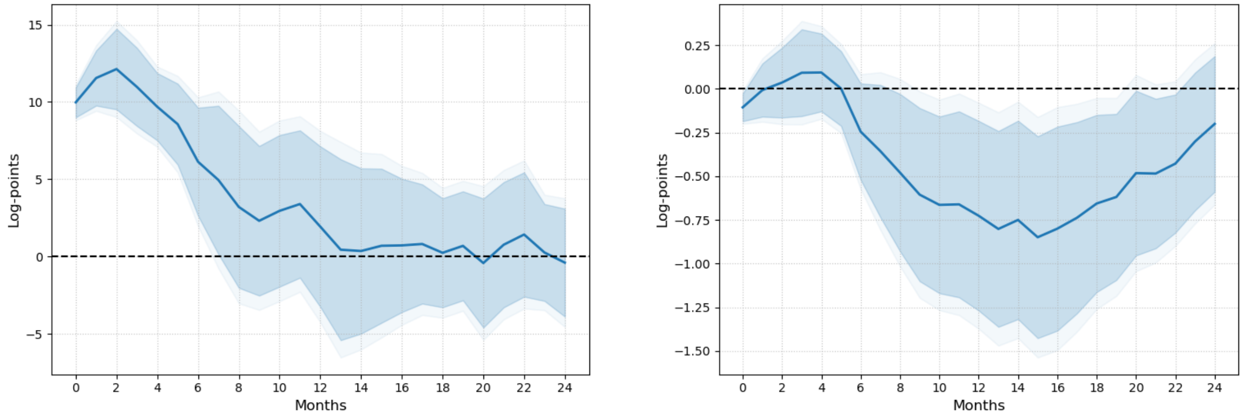

(a) Real WTI Oil Price (b) Industrial Production (OECD+6)

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. The dependent variable in each panel is the $h + 1$ change in the respective series. Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded from the sample. Sample: January 2002 – December 2024.

2.3.2 Global Warming Potential

The backdrop surrounding oil news shocks has two effects on energy consumption: first, the slowdown in economic activity decreases overall demand for energy. Second, the increase in

---

[2] Industrial production is used as a proxy for economic activity because it is available at the monthly frequency, contrarily to gross domestic product.

crude oil prices incentivizes substitution towards alternative energy sources. On the one hand, this effect can be negative, inducing more coal use - the most polluting fossil energy source. On the other, it can incentivize the uptake of renewable energy sources, such as solar. Naturally the horizons of these two substitution effects differ. The first could play out from the onset of the shock as coal plants with available capacity increase their output. The second can potentially take longer as it requires investments in new infrastructure since renewables tend to work at the limit of their useful capacity.

As a result, the overall effect of oil price shocks on overall ghg emissions is non-trivial. To assess it, figure 2 presents the impulse response functions for the aggregate CO_2e index, estimated using equation (1). Following an oil supply news shock, atmospheric CO_2e declines by 0.01 parts per million (ppm) relative to trend after 4 months. Following a short rebound, the concentration decline persists and stabilizes at 0.10 ppm. This short-term unitary semi-elasticity indicates that increases in oil prices contribute positively to global cooling and lead to immediate decreases in global warming potential relative to trend.

Naturally, the decline in emissions may be a result of the decreased economic activity and not of increased dirty energy efficiency, i.e., in the dirty energy necessary to support a certain amount of economic activity. A first way to indirectly evaluate that is by inspecting the response of emissions relative to economic activity. To do so, I compute the impulse response function of the logarithm of the ratio of the year-on-year difference in CO_2e concentration over the industrial production index for OECD+ countries. The use of the year-on-year difference deals with seasonality in concentration which may decrease despite positive emissions.

The results are plotted in figure 2 and show a decline in emissions per economic activity that is statistically significant through the 14th month — indicating that the observed decline in emissions is not just a result of the ensuing economic slowdown but of increased energy efficiency and/or a shift towards clean energy sources, in line with what (Känzig and Williamson, 2024) find for the United States. In contrast, I cannot statistically ascertain that the efficiency effect is permanent. This discrepancy may reflect sample selection: their longer sample is known to produce a more persistent decline in crude oil prices, which mechanically sustains the efficiency incentive over a longer horizon. Our subsequent reversion reflects the mean-reversion of crude oil prices themselves: as the supply disruption unwinds and prices return toward trend, the efficiency incentive dissipates.

Figure 2: Response of CO_2-Equivalent Concentration to an Oil Supply News Shock

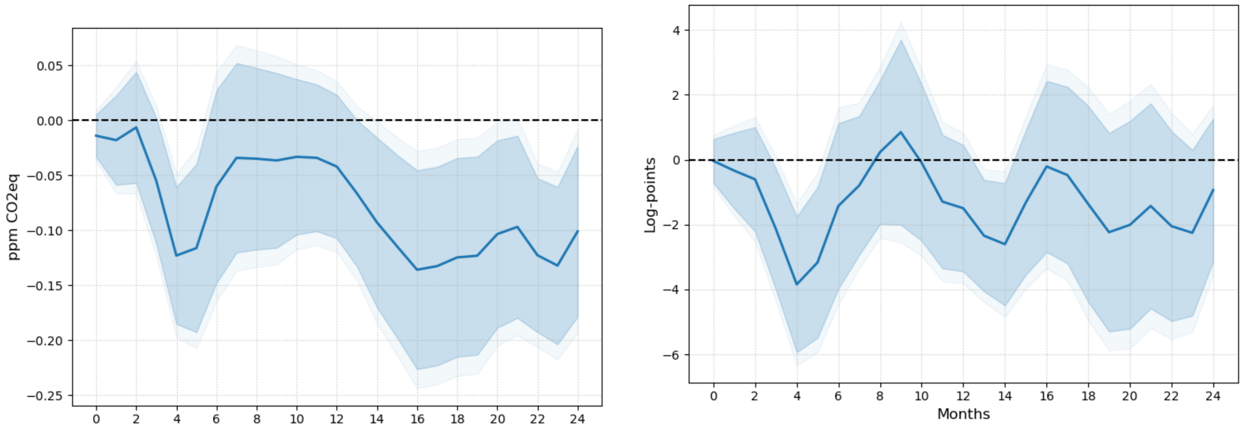

(a) CO_2e Concentration

(b) $\log \Delta^{12} CO_2e_t \big/ IP_t$

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. Panel (a): $h+1$ change in atmospheric CO_2e concentration (ppm). Panel (b): $h+1$ change in the log ratio of the 12-month difference in CO_2e concentration to the 12-month cumulative growth of industrial production. Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded from the sample. Sample: January 2002 – December 2024.

2.3.3 Individual Emissions

While overall global warming potential decreases as a result of an oil prices shock, this behavior may not be homogeneous across different gases because of the ensuing energy mix reshuffling. Moreover, understanding the individual ghg dynamics is important because they have different atmospheric lifetimes, and so different global warming potentials depending on the horizon of analysis (*Climate change 2014: synthesis report*, 2015).

To verify this, I also inspect the dynamics of each individual ghg, plotted in figure 3. The results confirm the heterogeneous dynamics. The most abundant ghg in the atmosphere, CO_2, follows a similar path to CO_2e: first it declines to 0.10 ppm and then plateaus to 0.1 ppm. Similarly, CH_4 decreases by 0.65 parts per billion (ppb) after 4 months but rebounds and its concentration mean-reverts, becoming statistical insignificant after 8 months.

Working in the opposite direction, the most powerful ghg rise: SF_6 increases 4 months after the shock, peaking at 0.004 parts per trillion (ppt), but then reverting to the mean after one year. In the same tone, N_2O increases after 4 months, peaking at 0.04 ppb before fluctuation down. In contrast, it remains elevated at around 0.05 ppb after 2 years. The persistent elevation of N_2O is consistent with the concurrent increases in lignite production and bioenergy generation documented in section 3 and subsection 4.2: both lignite combustion and biomass burning are established sources of N_2O through incomplete high-temperature combustion.

The documented heterogeneous dynamics across ghg confirm a reshuffling in the fossil fuel mix must occur, motivating the analysis of section 3.

Figure 3: Response of Individual Greenhouse Gas Concentrations to an Oil Supply News Shock

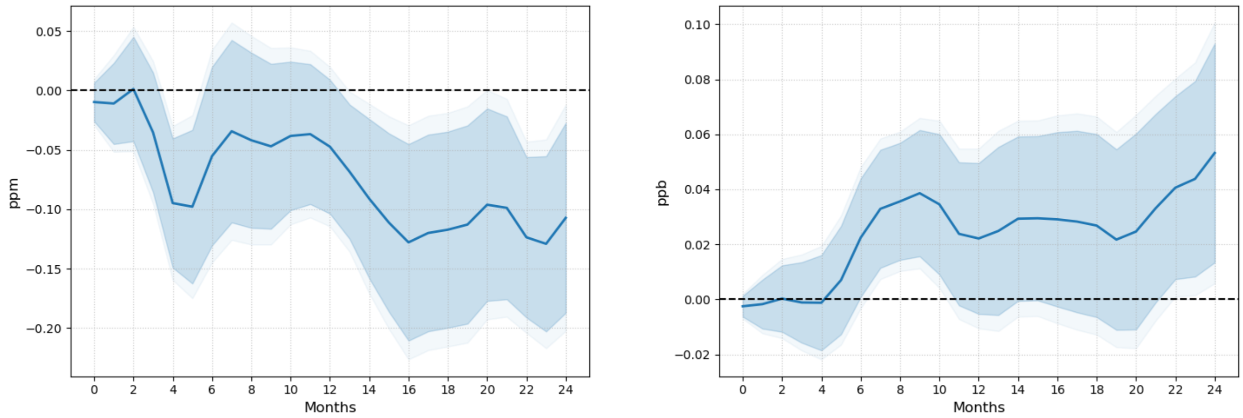

(a) CO_2 (b) N_2O

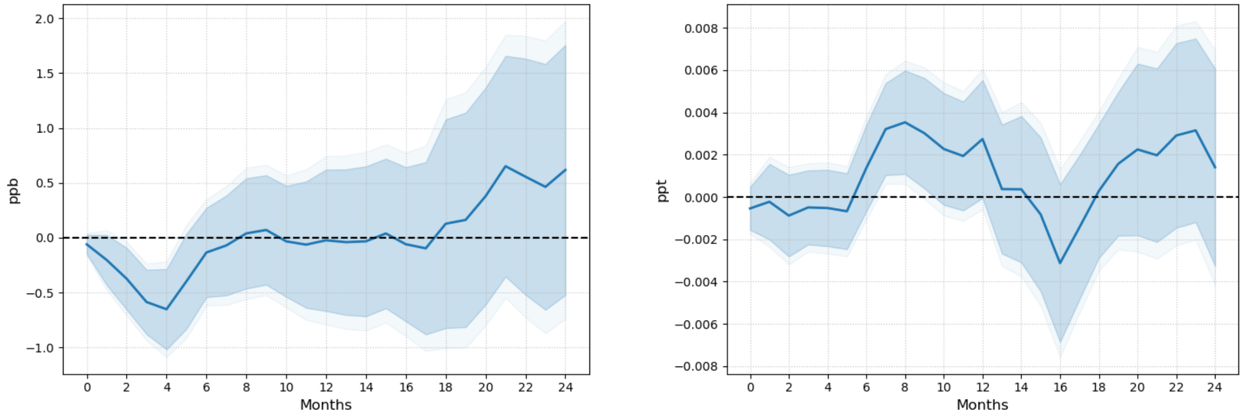

(c) CH_4 (d) SF_6

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. The dependent variable is the $h + 1$ change in atmospheric concentration. Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded from the sample. Sample: January 2002 – December 2024.

3 Hydrocarbon Dynamics

In this section I study the sources of ghg that explain the concentration dynamics presented in section 2. I focus on primary dirty energy consumption. This largely includes industrial processes and energy production, which together represented about 84% of all ghg emissions worldwide Climate Watch (2021). I thus exclude other sources of emissions like land-use, and agriculture. Specifically, I inspect the consumption and/or production dynamics of all types of fossil fuels following oil price hikes - less documented in the literature studying oil shocks.

3.1 Data

I build an unbalanced monthly panel dataset of countries from a variety of sources. All data is restricted to 2002 onwards or later, depending on its availability[3].

First, I obtain fuel consumption data from the Joint Organizations Data Initiative (JODI), a consortium that monitors and aggregates oil and gas data from national and international energy organizations. I obtain information on national refined oil fuel consumption for gasoline, diesel, jet fuel, residual oil — including bunker oil, and liquefied petroleum gas. These series thus covers primary energy derived from oil and gas (excluding electricity) used in the residential sector — specifically for heating and cooking — as well as the transportation sector, including automotive, aviation, and maritime transport. The data runs from 2002 to 2024 and covers more than 80 countries. I also obtain natural gas consumption data both for the whole economy and the electricity and heat generation sectors in particular. This data starts later in 2009 and includes more than 70 countries.

Secondly, in the absence of high-frequency consumption data, I obtain quarterly coal production data from the International Energy Agency. It starts in 2016 and covers a great share of the world's production including China, Russian, and OECD countries. Although, consumption and production of coal do not exactly match, like crude oil, their dynamics are similar - motivated by a desire for precautionary stockpiles.

3.2 Methodology

Ideally, I would build aggregate measures for each of the variables from the data collected. Nevertheless, the staggered inclusion of countries prevents such approach. A uniform sub-sample would either not be entirely representative of global activity or would have a much shorter time-span. Additionally, a panel set-up can increase the statistical power of my estimates in place of shorter series as is the case for coal, natural gas or for the power sector data. As a result, I use the individual country observations in the estimation. Following Chodorow-Reich et al. (2021); Ottonello and Winberry (2020), I adapt the methodology from subsection 2.2 to the panel set-up by running a fixed-effect version of local projections:

$$\Delta^{h+1} y_{i,t+h} = \alpha_{i,h} + \alpha_{i,m(t)} + \beta_h s_t + \sum_{k=1}^{12} (\rho_h \Delta y_{i,t-k} + \Gamma_h' X_{t-k} + \kappa_h \Delta C_{i,t-k} + \tau_h s_{t-k}) + \nu_h \iota_{i,t-1} \tag{2}$$

---

[3] Additional details about each dataset, its coverage and the raw time-series are provided in section B.

where i indexes the country. $\alpha_{i,h}$ are country-specific fixed-effects to account for differential growth trends. $\alpha_{i,m(t)}$ are country-specific month fixed-effects to capture the inherent seasonality of the dependent variable which may vary by country, especially because of different climates and geographies. X_{t-k} are the same global controls included in equation (1), and $\Delta C_{i,t-k}$ are lags of the logarithmic growth rate of country i's nominal exchange rate in USD from the International Monetary Fund. This partially for time variation in countries' economic and financial conditions and has a very broad coverage — ensuring that no observations are lost. $\iota_{i,t-1}$ is the previous month's income classification from the World Bank - coarsened into high and low income countries. All the dependent variables are defined in log-points. All regressions are weighted, so that β_h can be interpreted as a representative global estimate. The confidence bands are computed using Driscoll-Kraay standard errors (Driscoll and Kraay, 1998) with 12 month lags to account for cross-sectional and auto-correlation.

Including lags of the dependent variable can be a strenuous requirement, specifically for series with several zeros. Results are broadly consistent when I exclude lags of both the dependent variable and the shock, and instead include a time-trend. For this and additional robustness checks, see section C.

3.3 Refined Oil

I first plot the impulse response functions for refined oil products consumption in figure 4. Except for residual fuels, all series decrease on impact, albeit with different semi-elasticities and persistence. Jet fuel consumption decreases the most, approximately by 4 percentage points, and remains depressed for the two years following the oil price shock. Diesel and gasoline consumption both fall on impact reaching almost a 1 percentage point drop before rebounding and hovering around 25 basis points for the remaining horizon. Liquefied petroleum gas takes longer to drop, around 4 months, peaking at a 1 percentage point drop before mean-reverting.

In contrast, residual fuel, mainly made up of bunker oil used in shipping, actually rises on impact by 1 percentage point, but quickly reverts and the statistical uncertainty does not allow me to ascertain much about its dynamics.

Figure 4: Response of Refined Oil Consumption to an Oil Supply News Shock

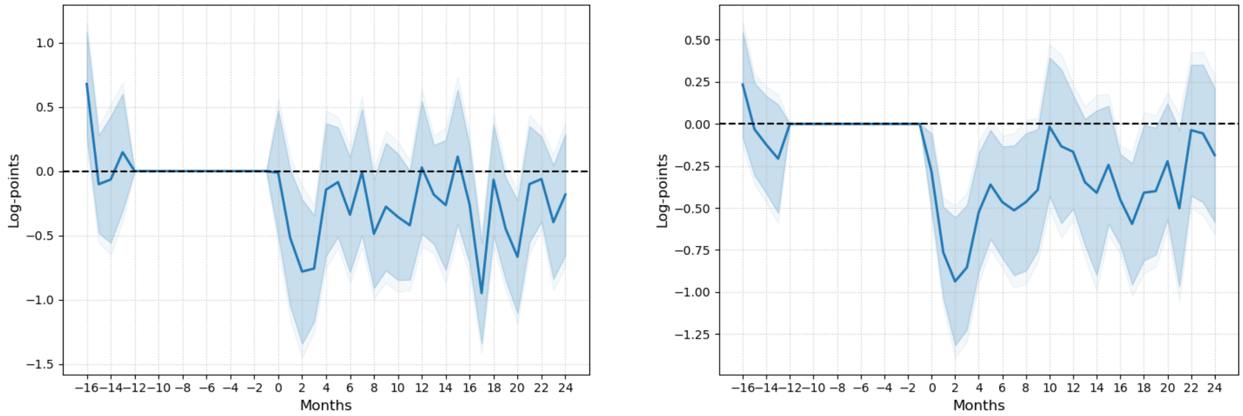

(a) Diesel

(b) Gasoline

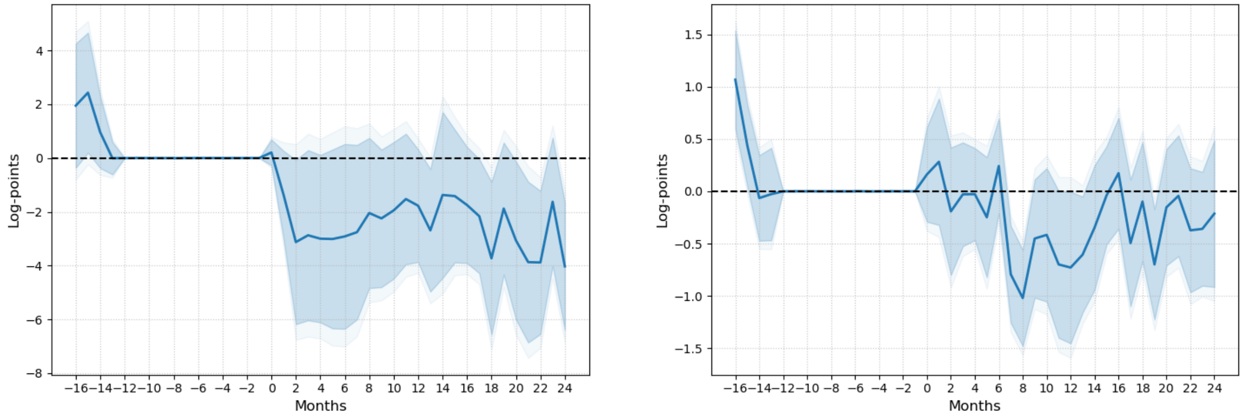

(c) Jet Fuel / Kerosene

(d) LPG

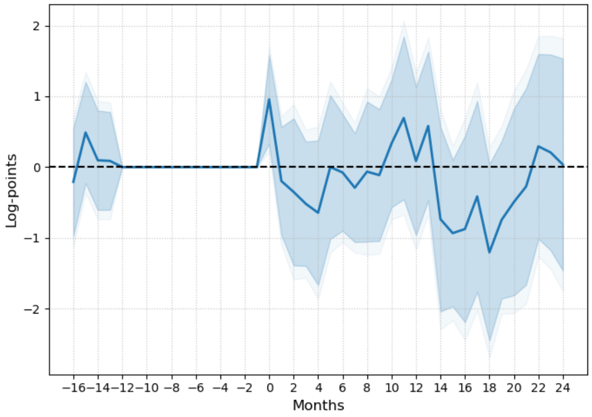

(e) Residual Fuel

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. The dependent variable is the $h + 1$ change in the log of monthly national consumption (thousand barrels). Shaded bands correspond to 90% and 95% confidence intervals based on Driscoll–Kraay standard errors (Driscoll and Kraay, 1998) with 12 lags. Regressions are weighted by the lagged 6-month moving average of the respective consumption series (thousand barrels). COVID-19 (2020) is excluded from the sample. Sample: 2002–2024.

3.4 Natural Gas

Natural gas is a partial substitute for oil in power generation and industrial heating, so an oil price increase may divert demand toward gas (Baumeister et al., 2024). At the same time, oil and gas are typically extracted from the same wells (U.S. Energy Information Administration, 2013). As a result, disruptions in the supply of the former can have consequences on the supply of the latter. Moreover, decreased economic activity can decrease the demand for natural gas. Figure 5 presents the impulse response functions for natural gas consumption following an oil supply news shock. Natural gas decreases in general, and in the power sector specifically. Total consumption continuously decays, dropping around 1 percentage point by the two year horizon. The drop is more severe in the power sector, which decreases to around 2 percentage points relative to trend.

This drop in natural gas consumption is not entirely explained by a decline in demand from slowed economic activity. In fact, the Henry Hub reference gas price, obtained from FRED, plotted in figure 5 rises by 3.64 percentage points on impact, reaching a peak of 11 pp after four months before slowly declining. This means that a decrease in supply must dominate and explain these dynamics.

Figure 5: Response of Natural Gas Consumption to an Oil Supply News Shock

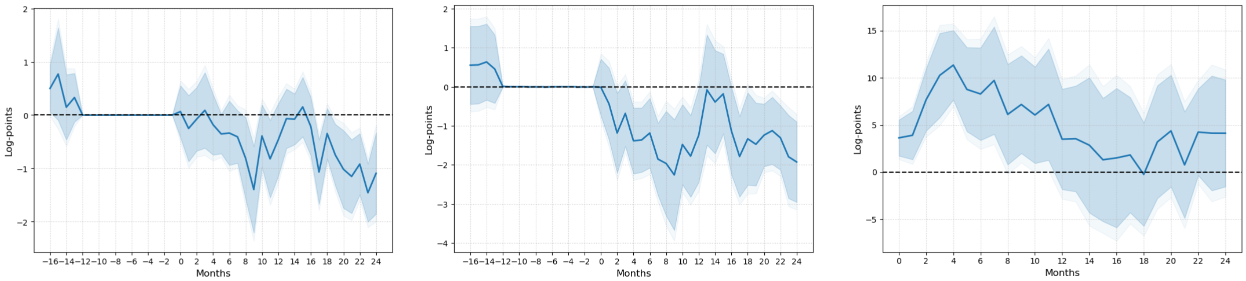

(a) Total Natgas (b) Power Sector Natgas (c) Henry Hub Price

Notes: Impulse responses to the oil supply news shock of Känzig (2021). Conventions as in figure 4, except regressions are weighted by the lagged 6-month moving average of the respective natural gas consumption series (m^3). Sample: 2009–2024.

3.5 Coal

Coal is the biggest source of energy pollution, a critical input in industrial production and a strong substitute for natural gas in power generation. As such, its dynamics can also be ambiguous. Slower industrial activity and declining energy demand decrease demand while increased energy prices induce demand substitution towards coal.

To understand how coal usage responds to oil price shocks, I adapt the methodology presented in subsection 3.2 to the quarterly frequency of the data. Figure 6 presents the impulse response functions for coal production. Steam coal, used in power generation and heating, and coking coal, used in metallurgy, decreases by percentage points on impact, reflecting the contraction in industrial production and decline in energy demand, and remain depressed after 2 years. In turn, lignite, the lowest ranked type of coal, increases by percentage points on impact and pp 3 quarters after the shock before reverting to the mean.

Figure 6: Response of Coal Productoin to an Oil Supply News Shock

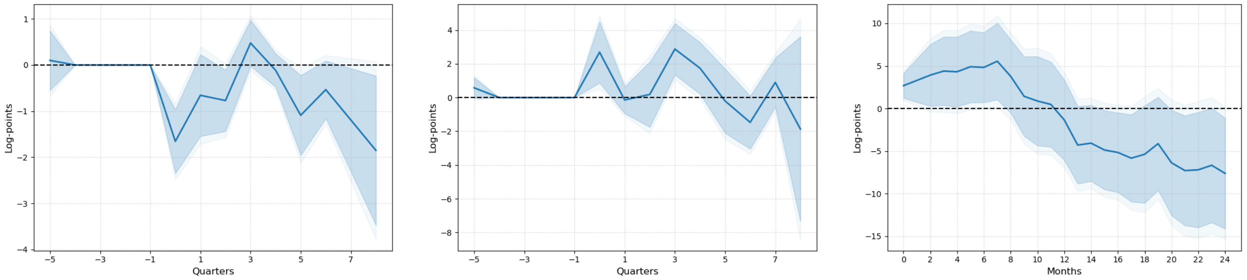

(a) Steam and Coking Coal (b) Lignite (c) Coal Price

Notes: Impulse responses to the oil supply news shock of Känzig (2021). Panels (a)–(b): dependent variable is the $h+1$ change in log quarterly coal production; horizon in quarters; Driscoll–Kraay 90% and 95% confidence intervals (Driscoll and Kraay, 1998); regressions weighted by the lagged 2-quarter moving average of the raw production level; sample: 2016–2024. Panel (c): dependent variable is the $h+1$ change in log Australian coal price (monthly); HAC 90% and 95% confidence intervals; sample: 2002–2024. COVID-19 (2020) excluded from all panels.

Together with the price dynamics presented in figure 6, our findings suggest that countries substitute away from more expensive higher quality coal to the cheaper, less carbon dense alternative, lignite.

4 Moving Away from Dirty Energy

The increase in oil & gas prices is typically viewed as an opportunity for electrification and for the adoption of alternative energy sources, in particular, renewable energy sources. In this section, I study those substitution channels. For data availability purposes, I focus on the behavior of (private) transportation, the biggest petroleum consuming sector, and electricity generation.

4.1 Transportation

I focus on the private vehicle market. On top of data availability, personal lower-emissions vehicles are much more developed than alternative transportation sectors, such as aviation or shipping, and so can considered as viable alternatives to internal combustion engine cars.

To study how this viability changes following increases in oil prices, I use vehicle registration data from (Robbie Andrew, 2026), who collects monthly new registrations from national statistical registries. This dataset starts in 1997 and covers more than 80 countries. I repeat the monthly panel methodology described in subsection 3.2.

Figure 7 plots the impulse responses of standard ICE registrations alongside all non-ICE registrations — a composite category including battery electric and hydrogen vehicles (ZEV), plug-in and non-plug-in hybrids, and alternative-fuel ICE vehicles such as LPG and flex-fuel cars. ICE registrations fall clearly and persistently, reaching the highest drop of around log-points 8 months after the shock and remaining depressed over the two-year horizon, consistent with consumers delaying or cancelling purchases of high-operating-cost vehicles. This effect is Non-ICE registrations display the opposite pattern: they rise persistenty reaching a peak of pp by the end of the considered time-horizon. While this increase in registration only compensates roughly 1/3 of the decreased registration in ICEs - total registration decreases to a peak of pp, a substitution channel is clearly at play. This reallocation is consistent with higher oil prices compressing the operating cost advantage of conventional vehicles relative to lower-emission alternatives, accelerating the green transition at the extensive margin. Moreover, because of vehicles' useful lives extending beyond the considered horizon, these dynamics have longer run effects on the transportation sector's emissions.

Figure 7: Response of Vehicle Registrations to an Oil Supply News Shock

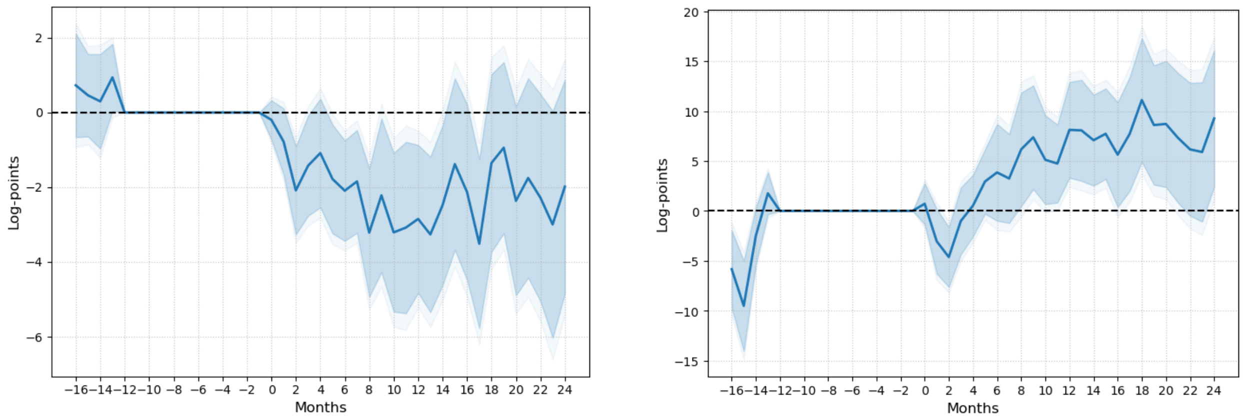

(a) ICE Registrations

(b) Non-ICE Registrations

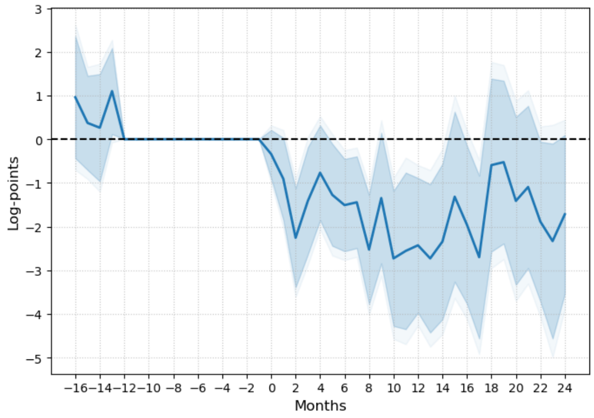

(c) Total Car Registrations

Notes: Impulse responses to the oil supply news shock of Känzig (2021). Conventions as in figure 4, except regressions are weighted by the lagged 6-month moving average of the respective vehicle registration count. Non-ICE includes zero-emission vehicles (battery electric, hydrogen), plug-in and non-plug-in hybrids, and alternative-fuel ICE vehicles (LPG, flex-fuel). Sample: 2002–2024.

4.2 The Power Sector

Two opposing forces shape the electricity sector's response to an oil supply shock. First, the contraction in economic activity documented above suppresses aggregate energy demand, pulling total electricity generation down. Second, higher oil prices compress the competitive advantage of oil-intensive processes and raise the relative attractiveness of electrification and renewable generation in particular. These substitution incentives operate at two margins: the intensive margin, where existing renewable plants run closer to installed capacity, and the extensive margin, where higher expected returns accelerate investment in new infrastructure. section 3 highlighted the former, whereby lignite coal production rises and natural gas consumption decreases in the power sector. This section focuses on the latter.

Electricity generation by source and wind and solar installed capacity come from Ember, a think tank that compiles data from national and international energy organizations. I obtain monthly TWh of electricity generated from coal, gas and other fossil fuels, nuclear, wind, solar, hydro, and other renewable sources, covering more than 80 countries. The data start in 1998, but the bulk of countries enter between 2010 and 2015. Capacity data record GW of nameplate capacity for wind and solar installations—including privately owned rooftop solar—starting in 2016 and covering 26 countries. In subsection B.1, I confirm that the Ember generation data match well the annual reporting from the US Energy Information Administration (EIA).

Figure 8 first presents the impulse response functions for total electricity generation and for the ratio of electricity generation to industrial production. Total generation contracts but by less than the proxy for economic activity, peaking at pp. The electricity-to-IP ratio, however, rises persistently: controlling for the aggregate cycle, electricity demand holds up, indicating that production increasingly relies on electrification as oil-intensive alternatives become costlier.

Figure 8: Response of Total Electricity Generation to an Oil Supply News Shock

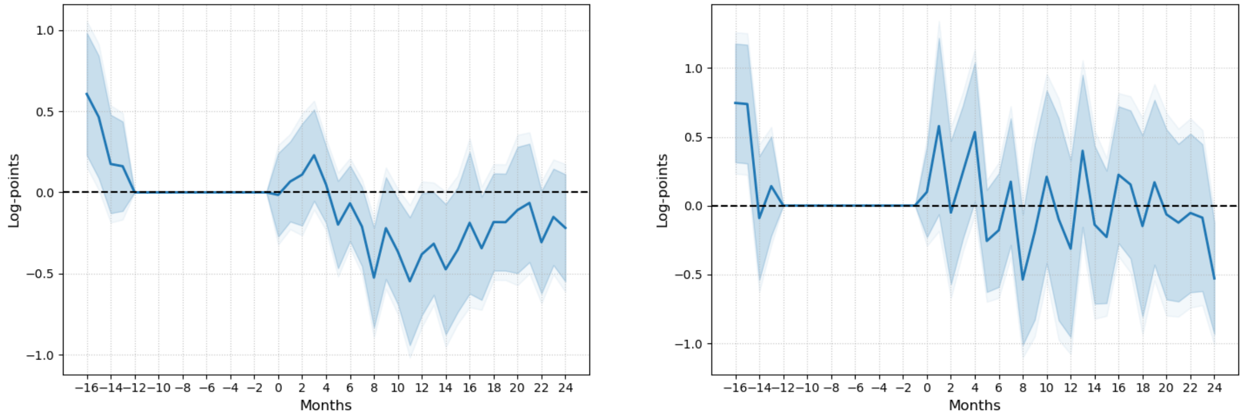

(a) Total Electricity Generation (TWh)

(b) Electricity Generation / Industrial Production

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. Conventions as in figure 4, except regressions are weighted by the lagged 6-month moving average of total electricity generation (TWh). Sample: 2002–2024.

Figure 9 analyses clean electricity generation. Total renewable energy generated decreases roughly a year after the shock by pp and stabilizing at pp at a 2-year horizon. Besides the blip in hydro after one year, this result appears driven almost entirely by the decline in solar generation which collapses by pp after 2 years. Wind, although fluctuating, appears to remain stable throughout.

Figure 9: Response of Renewable Electricity Generation by Source to an Oil Supply News Shock (TWh)

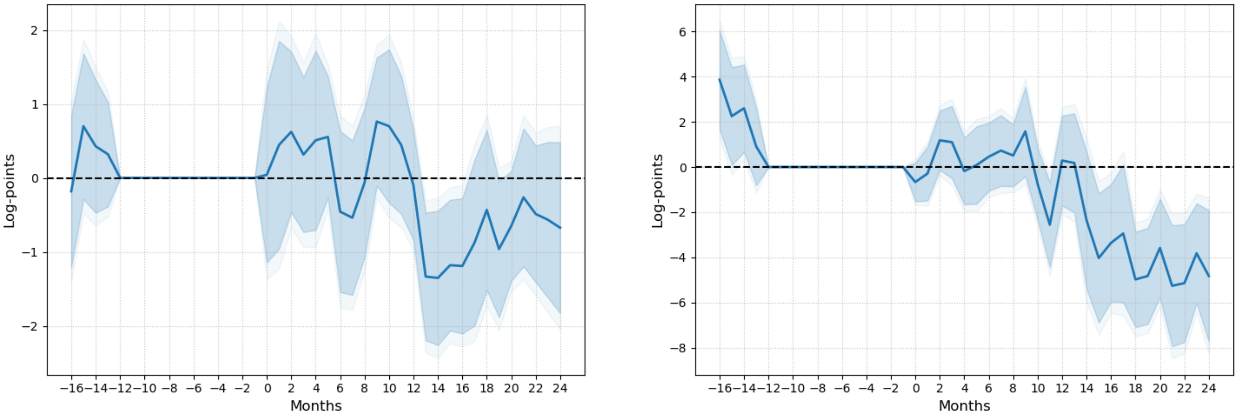

(a) Renewables x-Nuclear (b) Solar

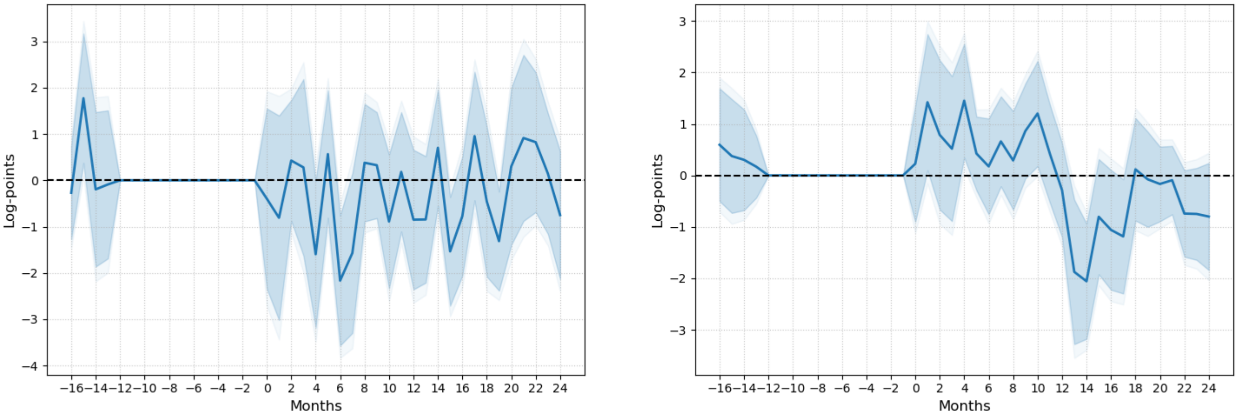

(c) Wind (d) Hydro

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. Conventions as in figure 4, except regressions are weighted by the lagged 6-month moving average of the respective generation series (TWh). Sample: 2002–2024.

The intensive-margin evidence is corroborated by the extensive margin. Figure 10 presents the impulse response functions for installed wind and solar capacity. While wind capacity does not move substantially, solar capacity decreases immediately after the oil price shock by around 1.5 pp.

Figure 10: Response of Renewable Installed Capacity to an Oil Supply News Shock

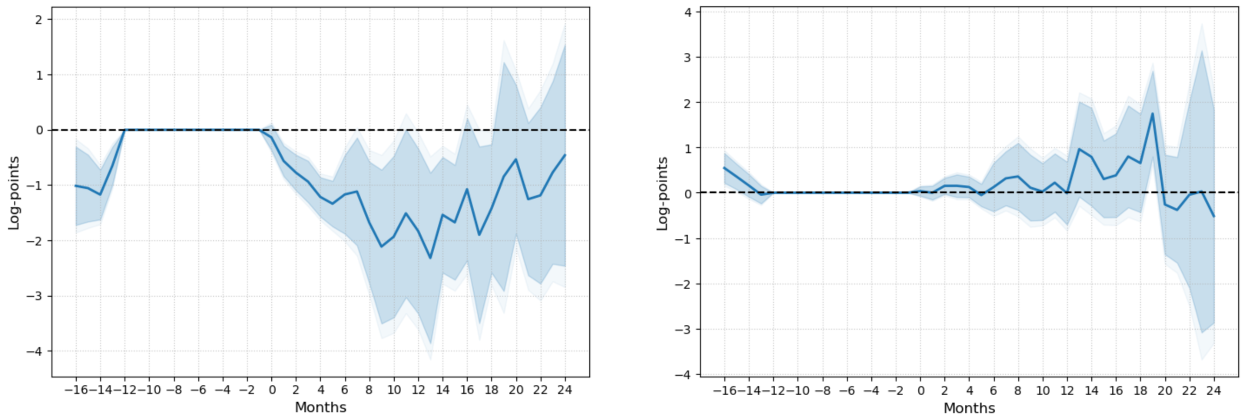

(a) Solar Installed Capacity (GW)

(b) Wind Installed Capacity (GW)

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. Conventions as in figure 4, except regressions are weighted by the lagged 6-month moving average of the respective installed capacity (GW). Capacity data cover 26 countries from 2016. Sample: 2016–2024.

While my findings regarding solar generation may appear surprising, they reflect three compounding mechanisms. First, solar power is intermittent—it generates only when irradiance is sufficient—and therefore requires fast-reacting dispatchable backup capacity, primarily natural gas peaker plants, to balance the grid during periods of low generation (Verdolini et al., 2018; Stöckl and Zerrahn, 2023). When oil shocks simultaneously disrupt natural gas supply and raise gas prices, the effective cost of deploying complementary solar capacity rises, dampening new investment. Second, the inflationary nature of the shock triggers an endogenous tightening of monetary policy (Bernanke et al., 1997), causing credit conditions to tighten—an environment that does not support capital-intensive activities such as solar installation (Bistline et al., 2023; Aghion et al., 2024; Ameli et al., 2021; Chen and Lin, 2024). Third, and separately, macroeconomic and geopolitical uncertainty rise sharply following the shock. Unlike the credit channel, elevated uncertainty is a distinct deterrent to long-lived, high fixed-cost investments, by raising the option value of waiting (Bloom, 2009). My results confirm all three channels. Figure 11 shows that the 10-year nominal interest rate, obtained from FRED, rises sharply on impact, reaching 0.16 basis points after 2 months. The term spread initially rises as long-term rates react before short-term rates catch up with monetary policy tightening, after which the spread narrows. Macroeconomic uncertainty (Jurado et al., 2015) spikes on impact and remains elevated, and geopolitical risk (Caldara and Iacoviello, 2022) rises persistently—all consistent with the suppression of renewable capacity additions documented in figure 10.

Figure 11: Financial Barriers to Clean Energy Adoption Following an Oil Supply News Shock

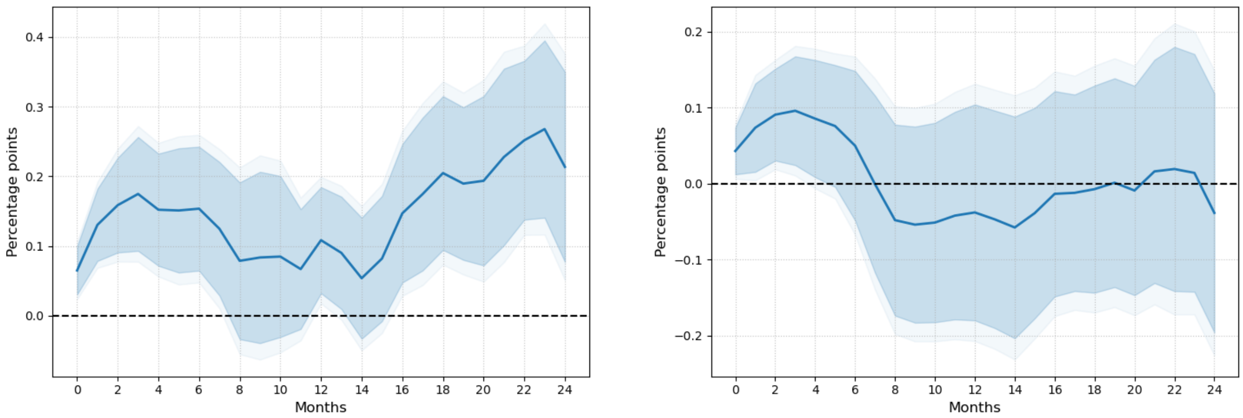

(a) 10-Year Nominal Interest Rate

(b) 10 − 1 Year Nominal Term-spread

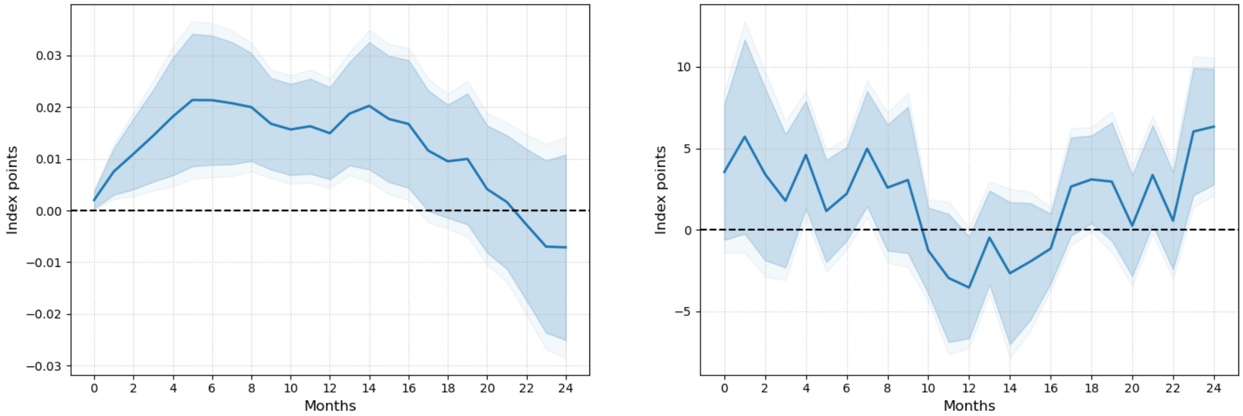

(c) Macroeconomic Uncertainty (JLN)

(d) Geopolitical Risk

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. The dependent variable in each panel is the $h + 1$ change in the respective series. Shaded bands correspond to 90% and 95% confidence intervals based on HAC standard errors. COVID-19 (2020) is excluded from the sample. Sample: January 2002 – December 2024.

Bioenergy — electricity generated from burning organic material such as wood pellets, agricultural residues, and dedicated energy crops — occupies an intermediate position between fossil fuels and clean renewables. Unlike solar and wind, it is dispatchable and emits during combustion, including CO_2 and N_2O. As a result, its behaviour following an oil price shock is informative about substitution in the dispatchable segment of the power mix.

Subsection 4.2 presents the impulse response of bioenergy generation. Following the shock, bioenergy rises persistently, reaching pp at the two-year horizon, though the effect is statistically significant only at the 10% level. This is consistent with biomass substituting for natural gas and oil in the dispatchable segment of power generation as their prices rise. Jointly with the concurrent rise in lignite production, this increase in biomass combustion contributes to the persistent elevation of atmospheric N_2O documented in section 2.

Figure 12: Response of Bioenergy Generation to an Oil Supply News Shock (TWh)

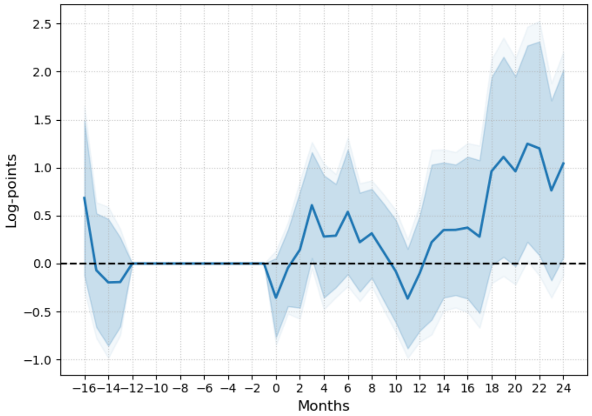

(a) Bioenergy Generation (TWh)

Notes: Impulse responses to the oil supply news shock of Känzig (2021), normalised to a 10% increase in oil prices on impact. The dependent variable is the $h + 1$ change in log monthly bioenergy generation (TWh). Shaded bands correspond to 90% and 95% confidence intervals based on Driscoll–Kraay standard errors (Driscoll and Kraay, 1998) with 12 lags. Regressions are weighted by the lagged 6-month moving average of bioenergy generation (TWh). COVID-19 (2020) is excluded. Sample: 2002–2024.

5 Conclusion

This paper documents the global environmental consequences of increases in oil prices. The aggregate picture is one of net improvement: a 10% increase in oil prices reduces atmospheric CO_2e concentration persistently. Notwithstanding, the underlying dynamics of individual greenhouse gases reveals important countervailing forces. At the fossil fuel level, the response is heterogeneous - though mostly positive. Jet fuel and kerosene undergo the deepest and most persistent demand contraction among petroleum products. Natural gas consumption falls, reflecting supply disruption rather than demand substitution. In contrast, coal exhibits a bifurcated response: steam and coking coal production decline alongside economic activity, while lignite — the cheapest and most carbon-intensive grade — rises. Bioenergy generation increases as well. A consequence of this is that N_2O's concentration in the atmosphere, a potent and long-lived greenhouse gas, rises persistently.

The clean-energy transition responds asymmetrically across sectors. At the transportation margin, the shock accelerates substitution away from internal combustion engines: ICE registrations fall persistently while non-ICE registrations rise, consistent with higher operating costs rendering cleaner vehicles relatively more attractive. At the power sector, the story reverses. Solar generation and installed capacity contract, driven by the combination of monetary-policy tightening that raises the cost of capital, rising geopolitical uncertainty, and grid-level complementarities between intermittent solar and dispatchable natural gas whose supply is itself disrupted by the shock. Total electricity generation falls by less than economic activity — indicating rising electrification intensity — but the dispatchable margin is filled by lignite and bioenergy rather than clean sources.

The emissions efficiency improvement observable in the near term — emissions per unit of industrial production decline in the first six months — is accordingly a short-lived intensive-margin response to higher energy costs, not a structural decoupling of output from emissions. It resembles the cyclical covariation between activity and carbon documented by Heutel (2012); Doda (2014) more closely than the permanent technological decoupling achievable through directed energy innovation (Hassler et al., 2021; Känzig and Williamson, 2024).

These findings carry implications for how oil market disruptions should be interpreted through a climate lens. An oil price shock delivers a partial and distorted carbon price signal: it compresses activity in oil-intensive transport but simultaneously undermines capital-intensive renewable investment and stimulates low-grade combustion alternatives at the dispatchable margin. Future versions of this work will examine how the transmission of oil shocks to the energy mix differs across countries with divergent energy mixes — a dimension that previous literature suggests is quantitatively important.

Focusing on the green transition in the long-run, future work should study how uncertainty in oil and gas production and the subsequent volatility in prices promotes the transition to green energy. This mechanism is commonly aluded to in discussions surrounding oil disruptions, but is not observed in our short-run analysis, in part because other sources of uncertainty and financial constraints hamper the installation and development of clean energy alternatives.

References

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hémous, "Climate Change, Directed Innovation, and Energy Transition: The Long-run Consequences of the Shale Gas Revolution," September 2023.

Aghion, Philippe, Antonin Bergeaud, Maarten De Ridder, and John Van Reenen, "Lost in Transition: Financial Barriers to Green Growth," March 2024.

Ameli, Nadia, Olivier Dessens, Matthew Winning, Jennifer Cronin, Hugues Chenet, Paul Drummond, Alvaro Calzadilla, Gabrial Anandarajah, and Michael Grubb, "Higher cost of finance exacerbates a climate investment trap in developing economies," *Nature Communications*, June 2021, *12* (1), 4046.

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.

_ **and James D Hamilton**, "Uncovering Disaggregated Oil Market Dynamics: A Full-Information Approach to Granular Instrumental Variables," April 2025.

_ **, Florian Huber, Thomas K. Lee, and Francesco Ravazzolo**, "Forecasting Natural Gas Prices in Real Time," *33156*, November 2024.

Bernanke, Ben S., Mark Gertler, and Mark W. Watson, "Systematic Monetary Policy and the Effects of Oil Price Shocks," 1997.

Bistline, John E. T., Neil R. Mehrota, and Catherine Wolfram, "Economics Implications of the Climate Provisions of the Inflation Reduction Act," *Brookings Papers on Economic Activity*, 2023, *2023* (2), 77–157.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA6248.

Caldara, Dario and Matteo Iacoviello, "Measuring Geopolitical Risk," *American Economic Review*, April 2022, *112* (4), 1194–1225.

Casey, Gregory, "Energy Efficiency and Directed Technical Change: Implications for Climate Change Mitigation," *The Review of Economic Studies*, January 2024, *91* (1), 192–228.

Chen, Shiu-Sheng and Tzu-Yu Lin, "Monetary policy and renewable energy production," *Energy Economics*, April 2024, *132*, 107495.

Chodorow-Reich, Gabriel, Plamen T. Nenov, and Alp Simsek, "Stock Market Wealth and the Real Economy: A Local Labor Market Approach," *American Economic Review*, May 2021, *111* (5), 1613–1657.

Climate Watch, "Historical GHG Emissions," 2021.

Climate change 2014: synthesis report

***Climate change 2014: synthesis report**, Technical Report, Intergovernmental Panel on Climate Change, Geneva, Switzerland 2015.*

Doda, Baran, *"Evidence on business cycles and CO2 emissions," Journal of Macroeconomics, June 2014, 40, 214–227.*

Driscoll, John C. and Aart C. Kraay, *"Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," The Review of Economics and Statistics, 1998, 80 (4), 549–560.*

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.*

Hamilton, James D, *"Oil and the Macroeconomy since World War II," Journal of Political Economy,* *April 1983, 91 (2), 228– 248.*

Harstad, Bård and Katinka Holtsmark, *"The Gas Trap: Outcompeting Coal versus Renewables," Journal of Political Economy. Forthcoming.*

Hassler, John, Per Krusell, and Conny Olovsson, *"Directed Technical Change as a Response to Natural Resource Scarcity," Journal of Political Economy, November 2021, 129 (11), 3039–3072.*

Heutel, Garth, *"How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks," Review of Economic Dynamics, April 2012, 15 (2), 244–264.*

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

Khan, Hashmat, Konstantinos Metaxoglou, Christopher R. Knittel, and Maya Papineau, *"Carbon emissions and business cycles," Journal of Macroeconomics, June 2019, 60, 1–19.*

Kilian, Lutz, *"Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," The American Economic Review, 2009, 99 (3), 1053–1069.*

_ **and Xiaoqing Zhou**, *"Heterogeneity in the pass-through from oil to gasoline prices: A new instrument for estimating the price elasticity of gasoline demand," Journal of Public Economics, April 2024, 232, 105099.*

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

Känzig, Diego R. and Charles Williamson, *"Unraveling the Drivers of Energy-Saving Technical Change," SSRN Electronic Journal, 2024.*

_ **, Maximilian Konradt, Lixing Wang, and Donghai Zhang**, *"Green Business Cycles," July 2025.*

Lindequist, David and Samuel Selent, *"Did shale gas green the U.S. economy?," Energy Economics, May 2025, 145, 108388.*

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.*

Newey, Whitney K. and Kenneth D. West, *"A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," Econometrica, 1987, 55 (3), 703–708.*

Ottonello, Pablo and Thomas Winberry, *"Financial Heterogeneity and the Investment Channel of Monetary Policy," Econometrica, 2020, 88 (6), 2473–2502.*

Robbie Andrew, *"Collected vehicle registration data," June 2026.*

Stöckl, Fabian and Alexander Zerrahn, *"Substituting Clean for Dirty Energy: A Bottom-Up Analysis," Journal of the Association of Environmental and Resource Economists, May 2023, 10 (3), 819–863.*

U.S. Energy Information Administration, *"Drilling often results in both oil and natural gas pro-*

duction - U.S. Energy Information Administration (EIA)," October 2013.

Verdolini, Elena, Francesco Vona, and David Popp, *"Bridging the gap: Do fast-reacting fossil technologies facilitate renewable energy diffusion?," Energy Policy, May 2018, 116, 242–256.*

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