NBER WORKING PAPER SERIES

THE UNEQUAL ECONOMIC CONSEQUENCES OF CARBON PRICING

Diego R. Känzig

Working Paper 31221
http://www.nber.org/papers/w31221

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2023, Revised September 2025

I am indebted to my advisors Paolo Surico, Hélène Rey, Florin Bilbiie and João Cocco for their invaluable guidance and support. For helpful comments and suggestions, I thank the coeditor, Emi Nakamura, three anonymous referees, my discussants Jim Stock, Johannes Stroebel, Ravi Bansal and Allan Timmermann, as well as Michele Andreolli, Juan Antolin-Diaz, Christiane Baumeister, Thomas Bourany, James Cloyne, Thomas Drechsel, Marty Eichenbaum, Martin Ellison, Rob Engle, Jesus Fernandez-Villaverde, Luis Fonseca, Luca Fornaro, Lukas Freund, Stephie Fried, Mark Gertler, Simon Gilchrist, Lars Hansen, Arshia Hashemi, Kilian Huber, Maral Kichian, Max Konradt, Joseba Martinez, Leo Melosi, Kurt Mitman, Silvia Miranda-Agrippino, Ben Moll, Elias Papaioannou, Ishan Nath, Tsveti Nenova, Luca Neri, Aleks Oskolkov, Christina Patterson, Pascal Paul, Gert Peersman, Giorgio Primiceri, Valerie Ramey, Sebastian Rast, Lucrezia Reichlin, Natalie Rickard, Esteban Rossi-Hansberg, Fabian Seyrich, Andrew Scott, Vania Stavrakeva, Jón Steinsson, Rob Vigfusson, Beatrice Weder di Mauro, Christian Wolf, Nadia Zhuravleva and participants at numerous seminars and conferences. I thank Pablo Sanchez and Luca Zanotti for excellent research assistance. I thank the ECB for the Young Economists' Prize, the IAEE for the Best Student Paper Award and the AQR Institute for the Fellowship Award. I am grateful to the Wheeler Institute for generously supporting this research. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Unequal Economic Consequences of Carbon Pricing
Diego R. Känzig
NBER Working Paper No. 31221
May 2023, Revised September 2025
JEL No. E32, E62, H23, Q54, Q58

ABSTRACT

This paper studies the economic impacts of carbon pricing. Exploiting institutional features of the European carbon market and high-frequency data, I identify carbon policy shocks and trace their dynamic effects. A restrictive carbon policy shock raises energy prices, reduces emissions, spurs green innovation, but decreases economic activity—disproportionately burdening poorer households. Not only are the poor more affected because of their higher energy spending, but they also experience larger income losses. These indirect, general-equilibrium effects via income and employment play an important role in the transmission of carbon pricing policies, accounting for about two-thirds of the aggregate consumption response.

Diego R. Känzig
Northwestern University
Department of Economics
and NBER
dkaenzig@northwestern.edu

A data appendix is available at http://www.nber.org/data-appendix/w31221
A Repository for carbon policy shocks is available at https://github.com/dkaenzig/carbonpolicyshocks

1. Introduction

The looming climate crisis has put climate change at the top of the global policy agenda. Governments around the world have started to implement carbon pricing policies to mitigate climate change, either via carbon taxes or cap and trade systems. While there is robust evidence supporting the effectiveness of carbon pricing policies in reducing emissions, the broader economic effects remain less well understood. What is the impact of carbon pricing on output, employment, and inflation and who bears the economic costs?

To answer these questions, I propose a novel approach to identify the aggregate and distributional effects of carbon pricing, exploiting institutional features of the European carbon market and high-frequency data. The European Union Emissions Trading System (EU ETS), one of the largest carbon markets globally, covers roughly 40 percent of the EU's greenhouse gas emissions. Established in phases, the market has undergone frequent regulatory updates. Using an event study approach, I collect 114 regulatory update events related to the supply of emission allowances. By measuring changes in carbon futures prices within narrow windows around these announcements, I isolate a series of carbon policy surprises. Reverse causality can be plausibly ruled out as economic conditions are known and priced before the regulatory news and are unlikely to change within the tight window considered. To address remaining concerns about predictability, I orthogonalize the surprises with respect to macroeconomic and financial data pre-dating the policy news. Using the resulting surprise series as an instrument in a semi-structural model of the European economy, I estimate the dynamic causal effects of a carbon policy shock.

I find that carbon pricing has significant effects on both emissions and the economy. A carbon policy shock tightening the carbon pricing regime causes an immediate increase in energy prices and a persistent fall in overall GHG emissions. A shock normalized to increase energy prices by 1 percent leads to a peak reduction in emissions of approximately 0.75 percent. However, these environmental gains come at an economic cost. Industrial production declines by nearly 1 percent, real GDP falls by around 0.3 percent, and the unemployment rate rises by 0.15 percentage points. Consumer prices increase by close to 0.2 percent, and stock prices fall by more than 2 percent. The emissions intensity improves, especially in the medium term. In line with this result, I document a significant uptick in low-carbon patenting as carbon pricing creates an incentive for green innovation.

The estimated magnitudes are substantially larger than what can be accounted for by the direct effect of higher energy prices alone. If energy demand is completely inelastic, the direct price effect is bounded by the energy share in expenditure—about 10 percent in Europe. In that case, a 1 percent increase in energy prices would imply a direct consumption effect of at most 0.1 percent. However, the estimated consumption response is larger, with a peak effect of about 0.3 percent. This gap suggests that indirect, general equilibrium effects—operating through prices and wages and thus income and employment—are important, accounting for roughly two-thirds of the total consumption response.

These results illustrate a trade-off between reducing emissions and the economic costs of carbon pricing. To quantify this trade-off, I estimate the marginal abatement cost implied by the observed responses to a carbon policy shock. The resulting estimate is slightly above 100 EUR per ton of CO2—highlighting that market prices may substantially understate the economy-wide costs of decarbonization.

These costs are not borne equally across society. Using a panel of advanced European economies, I document meaningful cross-country heterogeneity in responses, linked to differences in emissions intensity, the share of financially constrained households, and the strength of labor protection laws. Poorer countries also tend to be more adversely affected, though these comparisons are not very precise. Moreover, country averages may mask important within-country differences.

Using detailed household-level data from the United Kingdom (UK), I study the distributional impacts across income groups. I find that low-income households reduce their expenditure more than higher-income households in response to carbon pricing. Two factors contribute to this disparity. First, low-income households allocate a larger share of their spending to energy and thus, rising energy bills leave fewer resources for other spending. Second, they experience a sharper fall in income, as they are more likely to work in sectors that are disproportionately affected by the policy. Interestingly, these are not necessarily the most energy-intensive sectors, but rather sectors more sensitive to changes in demand—typically those producing discretionary goods and services. The overall drop in expenditure again exceeds what higher energy prices alone would predict: in monetary terms, energy spending rises only modestly, underscoring the importance of indirect effects via income.

These findings suggest that targeted fiscal policies could help alleviate the economic costs of carbon pricing. Since energy demand is relatively inelastic, such measures are unlikely to undermine emission reductions. I also find that carbon pricing significantly reduces support for climate-related policies, particularly among poorer households. Targeted compensation could therefore not only ease the economic burden but also help bolster public support for climate policy.

A comprehensive series of sensitivity checks confirms the robustness of the results across several dimensions, including the selection of event dates, identifying assumptions, estimation technique, model specification, and sample period. Importantly, the results are robust to accounting for confounding news over the event window using an heteroskedasticity-based estimator. A historical decomposition further shows that carbon policy shocks meaningfully contribute to variations in emissions over time, but do not explain the sharp decline during the global financial crisis—supporting the validity of the identification strategy.

Related literature and contribution. The traditional approach to study the environmental and economic effects of carbon pricing relies on structural models such as computable general equilibrium (CGE) or integrated assessment models (IAMs). These models highlight a trade-off between economic costs and environmental benefits but depend heavily on assumptions about technology, revenue recycling, and policy design (e.g. Goulder, 1995; Nordhaus, 1992; Golosov et al., 2014). With the expansion of carbon pricing policies worldwide, practical experience has created opportunities to complement these model-based insights with evidence directly from the data.

A key challenge for empirical work is that carbon prices are endogenous, since policymakers respond to macroeconomic conditions when setting climate policy. To address this, much of the literature has focused on firm- or facility-level effects. These studies credibly identify the *direct* effects of carbon pricing by comparing regulated to unregulated firms, with time fixed effects absorbing endogeneity in policy decisions. The evidence shows that carbon pricing reliably reduces emissions, while the effects on profits and employment are modest or insignificant (Fowlie, Holland, and Mansur, 2012; Martin, De Preux, and Wagner, 2014; Marin, Marino, and Pellegrin, 2018; Dechezleprêtre et al., 2019; Cui et al., 2021; Colmer et al., 2025, among others). Yet by construction, this approach nets out any impacts operating through market-wide price increases, such as electricity or other input prices. This limitation is important: the pass-through of carbon to electricity prices is strong (Fabra and Reguant, 2014; Hintermann, 2016), and energy price shocks are known to generate sizable macroeconomic consequences (Baumeister and Hamilton, 2019; Känzig, 2021).

Estimating the aggregate consequences of carbon prices is more challenging, since potential endogeneity cannot be addressed with time fixed effects. Recent work has sought to tackle this issue using synthetic control methods or by including rich macroeconomic controls (Andersson, 2019; Metcalf, 2019; Metcalf and Stock, 2020, 2023; Bernard and Kichian, 2021). Consistent with the firm-level evidence, this literature finds that carbon prices lower emissions but shows little evidence of adverse effects on output or employment.

I contribute to this literature by introducing a novel identification strategy that exploits high-frequency variation in carbon prices to estimate their aggregate effects. Methodologically, the approach builds on the high-frequency identification techniques widely used in monetary policy (Kuttner, 2001; Gürkaynak, Sack, and Swanson, 2005; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018*a*) and global oil markets (Känzig, 2021). In this framework, policy surprises are isolated from other macroeconomic forces by measuring asset price movements in narrow windows around policy announcements.

Several studies have used event study techniques to examine how regulatory carbon news affects carbon, energy, and stock prices (Mansanet-Bataller and Pardo, 2009; Bushnell, Chong, and Mansur, 2013; Fan et al., 2017; Meng, 2017, among others). To the best of my knowledge, this paper is the first to leverage these regulatory updates to analyze the macroeconomic effects of carbon pricing: by imposing some additional structure, I can trace the effects of carbon pricing on lower-frequency outcomes such as GDP, unemployment or consumer prices. Crucially, my approach estimates the *aggregate* effects of carbon pricing, including any general equilibrium adjustments.

Contrary to previous empirical studies, I find that carbon price changes have significant environmental *and* macroeconomic effects: emissions and economic activity fall substantially. This contrast likely reflects differences in sectoral coverage and policy design (Känzig and Konradt, 2024). The finding that ETS price changes have significant macroeconomic effects has been confirmed in subsequent work using alternative identifying assumptions (Bjørnland, Cross, and Kapfhammer, 2024; Ortubai et al., 2025).

Equipped with this identification strategy, I provide new evidence not only on the aggregate but also on the distributional consequences of carbon pricing. I find that carbon pricing in the EU has been more regressive than commonly thought, disproportionately burdening lower-income households. This stands in contrast to existing empirical studies, which tend to find more modest regressive impacts (Beznoska, Cludius, and Steiner, 2012; Ohlendorf et al., 2021). Importantly, these studies focus on the direct price impacts of carbon pricing and abstract from the indirect effects via income (Bigio et al., 2025).

Qualitatively, my findings are consistent with computable general equilibrium models in both macroeconomic (Goulder and Hafstead, 2018) and distributional impacts (Williams et al., 2015; Goulder et al., 2019), though quantitatively they lie at the upper range of model predictions. Importantly, my approach relies on much weaker structural assumptions. My findings illustrate the importance of accounting for indirect, general-equilibrium effects via prices and wages and thus income and employment. Incorporating financial frictions, household heterogeneity, and nominal rigidities is important to generate sufficiently large general equilibrium effects in macro-climate models.

Roadmap. The next section provides institutional background on the European carbon market and discusses the identification strategy. Section 3 outlines the econometric approach. Section 4 presents the aggregate effects of carbon pricing on emissions, innovation, and the macroeconomy. Section 5 studies the heterogeneous effects, using panel and household data. Section 6 concludes.

2. Institutional Background and Identification

2.1. The European carbon market

The European emissions trading system is the cornerstone of the EU's climate policy. Established in 2005, it was the world's first major carbon market and remains one of the largest globally. Covering over 11,000 heavy energy-using installations and airlines, it regulates roughly 40 percent of the EU's greenhouse gas emissions.

The market operates under a cap-and-trade principle. Unlike a carbon tax, which directly sets a price on emissions, the carbon market imposes a cap on total greenhouse gas emissions from covered installations. This cap declines over time to ensure continued emissions reductions. Companies receive emission allowances through auctions or free allocation, which they can trade, thereby establishing a market price for carbon. They may also use a limited number of international credits from eligible emission-reduction projects worldwide. Firms must monitor and report their emissions, surrendering sufficient allowances annually. Compliance is enforced with heavy fines. Since allowances are storeable, companies that reduce emissions can bank unused permits for future use or sell them at a profit (European Comission, 2020).

A brief history of the EU ETS. The development of the EU ETS was structured in different phases. Figure 1 shows the evolution of the carbon price across these phases. The first phase (2005–2007) served as a pilot to prepare for phase two,

Figure 1: The EU Carbon Price

Notes: The EU carbon price, as measured by the price of the annual EUA futures contract with closest expiry date, expressed in EUR/tCO2 or equivalent, over the different phases of the EU ETS. The red vertical lines indicate the end of a given ETS phase.

when the system needed to function efficiently to support the EU's Kyoto targets. During this initial phase, nearly all allowances were allocated freely at the national level. Due to the lack of reliable emissions data, phase one caps were set based on estimates. In 2006, the carbon price dropped sharply after it became clear that issued allowances exceeded actual emissions, eventually falling to zero, as phase one allowances could not be carried over to phase two.

The second phase (2008–2012) coincided with the first commitment period of the Kyoto Protocol, during which EU ETS countries faced binding emission targets. With verified emissions data from the pilot phase now available, the cap on allowances was reduced based on actual emissions. Free allocation decreased slightly, several countries introduced auctions, and firms were permitted to buy limited amounts of international credits. The European Commission also began expanding the system to cover more gases and sectors, including the aviation sector in 2012—though this applied only to flights within the European Economic Area. Despite these reforms, EU carbon prices remained at moderate levels. This was mainly because of the 2008 economic crisis, which caused a sharp drop in emissions. As the caps were not adjusted accordingly, a surplus of allowances accumulated, weighing down prices.

The third phase (2013–2020) introduced several key reforms. Notably, the system shifted from national caps to a single EU-wide cap, made auctioning the default method for allocating allowances with harmonized rules for free allocation, and expanded coverage to additional sectors and gases—including nitrous oxide and perfluorocarbons alongside carbon dioxide.

To address the surplus of allowances that had accumulated since the Great Recession, the European Commission postponed the auctioning of 900 million allowances in 2014—a measure known as 'back-loading'. Later, it introduced the market stability reserve, which became operational in January 2019. Designed to reduce the allowance surplus and enhance the system's resilience to major shocks, the reserve adjusted supply by withholding back-loaded and unallocated allowances from auctions during the final years of phase three.

The fourth and current phase (2021–2030) builds on earlier reforms. The legislative framework was revised in early 2018 to align the system with the EU's 2030 emission reduction targets. The annual reduction rate for total allowances increased from 1.74 percent to 2.2 percent, and the market stability reserve was strengthened to improve resilience to future shocks. Additional revisions and expansions are planned to support the EU's goal of climate neutrality by 2050 (see European Comission, 2020).

Regulatory events. As discussed above, the EU ETS has undergone significant reforms since its launch in 2005. Its institutional framework and rules have been continuously updated to address market challenges, improve efficiency, and reduce information asymmetry and distortions.

Building on the event study literature, I compile a comprehensive dataset of regulatory events in the EU ETS. These include decisions by the European Commission, votes in the European Parliament, and rulings by European courts. I focus on regulatory news related to the *supply* of emission allowances—specifically changes to the overall cap, free allocation, auctioning, and the use of international credits.

Using the official journal of the European Union as well as the European Commission Climate Action news archive, I identify 126 regulatory events between 2005 and 2019. Through a detailed narrative analysis of event coverage on Factiva, I identify a subset of events that coincided with major economic news, such as oil price shocks, the sovereign debt crisis, or Brexit. To isolate the effects of regulatory changes, I exclude these potentially confounded events, leaving 114 events for analysis (see Appendix A.1 for more information).

Only a few events pertain to setting the overall cap. In the first two phases, key events involved decisions on national allocation plans (NAPs), including Commission approvals, rejections, and court rulings on free allocation disputes. With auctioning becoming the default allocation method in phase three, regulatory news shifted toward decisions on auction timing and quantities. Starting in phase two, there were also a number of relevant events related to the use of international credits.

Finally, note that not all events are necessarily news as some regulatory changes may be anticipated by the market. My identification strategy will account for this by focusing on the unexpected component.

Carbon futures markets. EU emission allowances (EUAs) are traded in several organized markets. Each EUA grants the right to emit one ton of CO2 equivalent. Key spot markets include Bluenext in Paris, EEX in Leipzig, and Nord Pool in Oslo. There are also well established futures markets for EUAs: the EEX in Leipzig and ICE in London. In 2018, cumulative trading volume across relevant futures and spot markets reached approximately 10 billion EUAs (DEHSt, 2019). The most liquid markets for emission allowances are the futures markets. I focus on price data for the December contract from the ICE, which is the most liquid and has been found to dominate price discovery in the European carbon market (Stefan and Wellenreuther, 2020).

2.2. High-frequency identification

Carbon prices are not set in a vacuum—policymakers take economic considerations into account when deciding on climate policy. Therefore, regressing macroeconomic variables on carbon prices generally leads to biased results. To address this concern, I adopt a high-frequency identification strategy to isolate plausibly exogenous variation in carbon prices.

The institutional setting of the European carbon market provides an ideal setting for this approach. First, as discussed above, there are frequent regulatory updates that can have significant market impacts. Second, liquid futures markets provide accurate, high-frequency price data for trading allowances. This motivates the construction of a series of carbon policy surprises by examining how carbon prices respond to regulatory news. By measuring price changes within a sufficiently tight window, reverse causality can be plausibly ruled out—economic conditions are already known and priced in prior to the event, and are unlikely to change within such a short window.

I construct the carbon policy surprise series as the change in the EUA futures price on the day of a regulatory event relative to the last trading day before the event. Because carbon prices were near zero at the end of the first phase, I express the EUA price change in euros, normalized by the prevailing wholesale electricity price on the day prior to the event:

$$CPSurprise_d = \frac{F_d^{carbon} - F_{d-1}^{carbon}}{P_{d-1}^{elec}}, \tag{1}$$

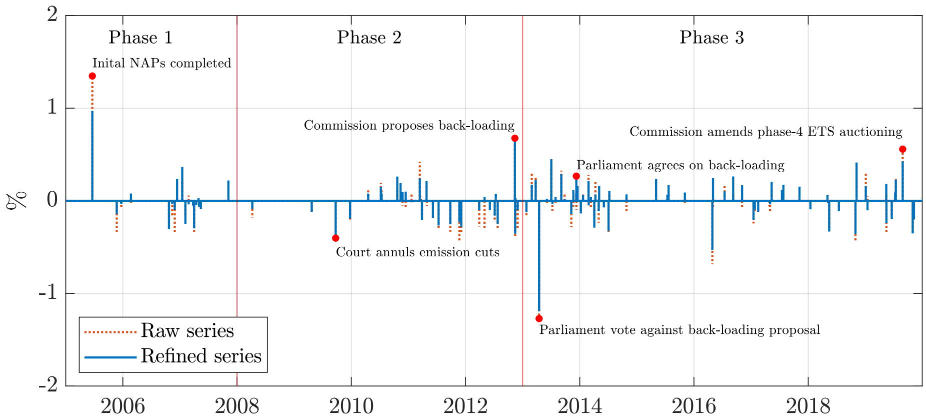

Figure 2: The Carbon Policy Surprise Series

Notes: The daily carbon policy surprise series, constructed based on (1) as EUA futures price changes around regulatory events normalized by the prevailing wholesale electricity price together with the refined surprise series, constructed as the residual of the predictive regression (2) using macro, financial, oil and climatic predictors. The red vertical lines indicate the end of a given ETS phase. A selection of notable events are depicted as red dots with annotations. Appendix Figure C.6 shows the surprises expressed as percentage changes in carbon prices.

where d indicates the event date, F_d is the settlement price of the EUA futures contract on the event day, and P_{d-1}^{elec} is the wholesale electricity price on the day before the event. This approach isolates variation in the carbon price driven by regulatory news, under the assumption that risk premia do not change systematically over the narrow event window. An alternative approach is to express the surprise series as the percentage change in the carbon price around the event. Reassuringly, this yields similar results, particularly when excluding the second half of 2007, when carbon prices approached zero. See Appendix C.1 for details.

Figure 2 presents the resulting carbon policy surprises as the red dashed line, capturing the impact of regulatory news on carbon prices relative to electicity prices. Appendix Figure C.6 shows the corresponding percentage changes.

As a first validation of the carbon policy surprise series, I examine several key regulatory events and their market impacts. The approval of the Greek NAP on June 20, 2005 marked the completion of the first round of national allocation plans. The conclusion of this phase—along with news that Greece had to drop a provision for ex-post adjustments—led to a stark increase in carbon prices. On September 23, 2009, the European Court of First Instance annulled the Commission's decision to reduce allowance allocations for Poland and Estonia. Carbon traders and brokers reacted strongly, driving prices down. On November 12, 2012, the Commission proposed back-loading 900 million allowances from the start of phase three to address the oversupply of carbon allowances. Markets responded sharply, with carbon prices increasing by nearly 9 percent. On April 16, 2013, the European Parliament voted against this proposal, mainly due to concerns over potential adverse economic impacts on a still fragile European economy. This news triggered a very stark market reaction, with prices falling by over 40 percent. On December 10, 2013, the European Parliament approved a modified version of the proposal, paving the way for higher carbon prices in the market. Finally, on August 28, 2019, the Commission adopted changes to the ETS auctioning regulation for phase four, triggering another upward movement in prices.

We have seen that all these events had sizable impacts on carbon prices. These changes are also economically meaningful when expressed relative to wholesale electricity prices, with some events implying a change in electricity prices of nearly 1.5 percent, assuming full pass-through (Fabra and Reguant, 2014).

By contrast, some events triggered little or no market reaction—either because the news was minor or the regulatory changes were anticipated and already priced in. This illustrates how the high-frequency identification approach isolates the unexpected component of policy news.

Construction choices. A crucial choice in high-frequency identification is the size of the event window. There is a trade-off between capturing the entire response to the announcement and avoiding contamination from other contemporaneous news, so-called background noise (Nakamura and Steinsson, 2018*a*). To allow markets sufficient time to react to regulatory news, I use a daily event window. Using a tighter, intraday window is complicated by the fact that exact release times of the regulatory news are mostly unavailable. To mitigate concerns about confounding news when using a daily window, I also present results from a heteroskedasticity-based approach that accounts for background noise in the surprise series. These results turn out to be very similar (see Appendix C.2).

Another key choice concerns the maturity of the futures contract. I focus on the front contract (the nearest expiry) for two reasons. First, it is the most liquid and provides the clearest price signal. Second, near-dated contracts tend to be less sensitive to risk premia (Baumeister and Kilian, 2017; Nakamura and Steinsson, 2018*a*), helping to mitigate concerns about time-varying risk. However, using contracts with longer maturities yields similar results, see Appendix Figure C.9.

A final choice concerns the selection of events. My approach aims at imposing as little judgment as possible—therefore I include all relevant events unless an obvious confounding factor is identified in the narrative analysis. The key requirement is that the event is about the supply or allocation of emission allowances. As such, it should not convey other information, such as news about the demand of allowances or economic activity more broadly. To this end, I only include specific regulatory events in the European carbon market and exclude broader developments, such as Conference of the Parties (COP) meetings or other international conferences.

In a series of sensitivity checks, I show that the results are not driven by any particular subset of events. In particular, the results remain robust when excluding events from the first trial phase or omitting event days in periods of economic distress, such as the Great Recession or the European debt crisis. I also perform a jackknife exercise to assess the influence of individual events, showing that the estimates are not driven by specific outliers (see Appendix C.1).

Predictability of carbon policy surprises. An influential literature finds that in the monetary policy context, high-frequency surprises are predictable based on publicly available macroeconomic and financial data preceding the policy announcement (Cieslak, 2018; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023*a*,*b*). This predictability challenges the interpretation of such surprises as primitive "shocks" and may bias estimates of their dynamic effects.

Are carbon policy surprises also predictable based on past macroeconomic and financial variables? To assess this, I regress the daily surprise series on relevant information available before the event:

$$CPSurprise_d = \alpha + \beta' X_{d-} + \eta_d, \tag{2}$$

where d indexes carbon policy event days, $CPSurprise_d$ denotes the carbon policy surprise series, and X_{d-} is a set of predictors known before the announcement day d, as indicated by the subscript $d-$.

As predictors, I consider a wide range of macroeconomic and financial variables. For macroeconomic and financial indicators, I follow Bauer and Swanson (2023*a*). Specifically, I include the surprise components of the latest Eurozone real GDP, unemployment rate, CPI, and PPI releases prior to each event. These surprises are defined as the difference between the actual release and Bloomberg survey expectations. I also include the log change in the Eurostoxx price index, the change in the 10-year yield, the change in the BBB bond spread, and the log change in a commodity price index—all measured over the three months leading up to the event. For oil market conditions, I add the three-month log changes in the Brent crude price, along with the log changes in European petroleum consumption and production. Finally, for climatic factors, I include the three-month change in heating degree days measured prior to the event.

Table 1 presents the results. There is limited evidence that carbon policy surprises are predictable by macroeconomic and financial variables, as none of the estimated coefficients are statistically significant at conventional levels. The same

Table 1: Predictability of Carbon Policy Surprises

Carbon policy surprise:	(a) Macro news	(b) Financials	(c) Oil market	(d) Climatic variables
Real GDP surprise	0.019	-0.074	0.141	-0.268
	(0.373)	(0.403)	(0.528)	(0.499)
Unemployment rate surprise	-0.317	-0.278	-0.158	0.056
	(0.199)	(0.200)	(0.212)	(0.215)
CPI surprise	-0.995	-1.022	-1.498	-1.566
	(0.710)	(0.689)	(0.895)	(0.855)
PPI surprise	0.114	0.080	-0.036	-0.001
	(0.126)	(0.130)	(0.187)	(0.171)
Eurostoxx (3M log change)		-0.056	-0.013	0.062
		(0.641)	(0.589)	(0.583)
10-year yield (3M change)		0.030	-0.003	0.001
		(0.048)	(0.057)	(0.056)
BBB bond spread (3M change)		0.013	0.021	0.016
		(0.053)	(0.051)	(0.052)
Commodity price index (3M log change)		0.099	0.289	0.749
		(0.453)	(0.482)	(0.484)
Brent crude (3M log change)			0.124	-0.082
			(0.376)	(0.341)
Petroleum consumption (3M log change)			-2.116	-2.553
			(1.210)	(1.183)
Petroleum production (3M log change)			-0.385	0.009
			(0.303)	(0.319)
Heating degree days (3M change)				-0.055
				(0.022)
R^2	0.049	0.055	0.120	0.177
Adj. R^2	0.014	-0.017	0.025	0.079

Notes: Estimated coefficients β, R^2 and adj. R^2 from predictive regressions (2) of carbon policy surprises. The predictors X are observed prior to the event and include: the surprise component of the most recent Eurozone real GDP, unemployment rate and CPI release in column (a); column (b) adds the log change in the Eurostoxx from 3 months before to the day before the event, the change in the Eurozone 10-year yield, the change in the BBB bond spread and the log change in the Bloomberg commodity price index over the same period; column (c) adds the log change in Brent crude, European petroleum consumption and production over the same period; column (d) adds three-month change in heating degree days. Robust standard errors in parentheses. Number of observations: 114.

holds true for commodity and oil prices. The only predictors with some explanatory power are petroleum consumption and heating degree days. This is not necessarily problematic, as it may stem from plausibly exogenous factors like unusual weather patterns. Nevertheless, the R^2 values range from 5 to 18 percent.

To account for this potential predictability, I follow the approach by Bauer and Swanson (2023*b*). Specifically, I construct a refined carbon policy surprise series as the residual from the predictive regression (2), controlling for the full information set (d). The resulting series $\widetilde{CPSurprise}_d = \eta_d$, shown as the blue line in Figure 2, closely tracks the raw series, with a correlation coefficient of 0.90.

To mitigate any identification concerns related to the predictability of the surprise series, I use the refined carbon policy surprise series as the baseline. As shown in Appendix C.1, accounting for the predictability has meaningful implications for the results, even though the differences are not statistically significant.

Aggregation and additional diagnostics. As we are ultimately interested in outcome variables observed only at lower frequencies, I aggregate the daily surprises $\widetilde{CPSurprise}_d$ to a monthly series, $\widetilde{CPSurprise}_t$, by summing over the daily surprises in a given month t. Importantly, some months include more than one regulatory event and therefore contribute multiple observations to the monthly series. In months without any regulatory events, the series takes a value of zero.

Finally, I conduct several additional diagnostic checks on the refined monthly surprise series, following Ramey (2016). First, I find no evidence that the series is serially correlated. The p-value for the Q-statistic that all autocorrelations are zero is 0.98. A series of Granger causality tests confirms that the monthly series is not forecastable by past macroeconomic variables. Lastly, I show that the surprise series is uncorrelated with other structural shock measures from the literature, including oil demand, uncertainty, financial, fiscal, and monetary policy shocks. Taken together, these findings support the validity of the carbon policy surprise series. The corresponding figures and tables can be found in Appendix B.1.

3. Econometric Approach

As illustrated above, the carbon policy surprise series exhibits many desirable properties. Nonetheless, it remains an imperfect proxy for the underlying policy shock, as it may not capture all relevant regulatory news in the carbon market and is potentially subject to measurement error (see also Stock and Watson, 2018). Therefore, I do not use it as a direct shock measure but rather as an *instrument*. Provided that the surprise series is correlated with the true carbon policy shock but uncorrelated with all other shocks, it can be used to estimate the dynamic causal effects of carbon policy shocks.

A challenge in estimating the dynamic causal effects using high-frequency surprises is the so-called power problem (Nakamura and Steinsson, 2018*a*). Over the impulse horizon, macroeconomic variables are influenced by a myriad of other shocks, while high-frequency carbon policy surprises explain only a small share of the fluctuations in energy prices—resulting in a low signal-to-noise ratio. This makes it difficult to directly estimate macroeconomic effects of high-frequency carbon policy surprises using local projections à la Jordà (2005).

To address this challenge, I rely on VAR techniques for estimation, using the external instruments approach (Stock, 2008; Stock and Watson, 2012; Mertens and Ravn, 2013).

3.1. Framework

Our goal is to jointly model the European economy and the carbon market. Let $\mathbf{y}_t$ denote a $n \times 1$ vector of monthly time series. I assume that the dynamics of $\mathbf{y}_t$ can be characterized by the following structural vector moving-average representation:

$$\mathbf{y}_t = \mathbf{B}(L)\mathbf{S}\boldsymbol{\varepsilon}_t, \tag{3}$$

where $\boldsymbol{\varepsilon}_t$ is a vector of mutually uncorrelated structural shocks driving the economy, $\mathbf{B}(L) \equiv \mathbf{I} + \mathbf{B}_1 L + \mathbf{B}_2 L^2 + \ldots$ is a matrix lag polynomial, and $\mathbf{S}$ is the structural impact matrix.

Assuming that the vector-moving average process (3) is invertible, it admits the following VAR representation:

$$\mathbf{A}(L)\mathbf{y}_t = \mathbf{S}\boldsymbol{\varepsilon}_t = \mathbf{u}_t, \tag{4}$$

where $\mathbf{u}_t$ is a $n \times 1$ vector of reduced-form innovations with variance-covariance matrix $\text{Var}(\mathbf{u}_t) = \boldsymbol{\Sigma}$ and $\mathbf{A}(L) \equiv \mathbf{I} - \mathbf{A}_1 L - \ldots$ is a matrix lag polynomial. Truncating the VAR to order p, I can estimate the model using standard techniques and recover an estimate of $\mathbf{A}(L)$.

We are interested in characterizing the causal impact of a single shock. Without loss of generality, let us denote the carbon policy shock as the first shock in the VAR, $\varepsilon_{1,t}$. Our aim is to identify the structural impact vector $\mathbf{s}_1$, which corresponds to the first column of $\mathbf{S}$.

External instrument approach. Identification using external instruments works as follows. Suppose there is an external instrument available, z_t. In the application at hand, z_t is the carbon policy surprise series. For z_t to be a valid instrument, we need

$$\mathbb{E}[z_t \varepsilon_{1,t}] = \alpha \neq 0 \tag{5}$$

$$\mathbb{E}[z_t \boldsymbol{\varepsilon}_{2:n,t}] = \mathbf{0}, \tag{6}$$

where $\varepsilon_{1,t}$ is the carbon policy shock and $\boldsymbol{\varepsilon}_{2:n,t}$ is a $(n-1) \times 1$ vector consisting of the other structural shocks. Assumption (5) is the relevance requirement and assumption (6) is the exogeneity condition. These assumptions, in combination with the invertibility requirement (4), identify $\mathbf{s}_1$ up to sign and scale:

$$\mathbf{s}_1 \propto \frac{\mathbb{E}[z_t \mathbf{u}_t]}{\mathbb{E}[z_t \mathrm{u}_{1,t}]}, \tag{7}$$

provided that $E[z_t \mathrm{u}_{1,t}] \neq 0$. To facilitate interpretation, I scale the structural impact vector such that a unit positive value of $\varepsilon_{1,t}$ has a unit positive effect on $\mathrm{y}_{1,t}$, i.e. $\mathrm{s}_{1,1} = 1$. I implement the estimator with a 2SLS procedure and estimate the coefficients above by regressing $\hat{\mathbf{u}}_t$ on $\hat{\mathrm{u}}_{1,t}$ using z_t as the instrument. To conduct inference, I employ a residual-based moving block bootstrap (Jentsch and Lunsford, 2019).

VAR assumptions. The VAR approach relies on two potentially restrictive assumptions (Nakamura and Steinsson, 2018*b*). The first is invertibility, meaning that the model incorporates all relevant information needed to recover the structural shocks of interest.[1] The second assumption concerns the dynamic structure of the VAR, specifically that a finite-order VAR adequately captures the dynamics of the data-generating process.

To assess how restrictive these assumptions are, I perform a number of sensitivity checks. In Appendix C.5, I relax the invertibility requirement and present results from an internal instruments VAR and a local projections instrumental variable approach (Plagborg-Møller and Wolf, 2021). I also evaluate the extent of lag truncation bias by using alternative estimation strategies, including simple local projections and a Bayesian VAR with more lags, see Appendix C.6. The results are robust to relaxing the assumptions underlying the baseline VAR, however, the VAR structure improves precision and allows for sharper inference.

Local projections for additional outcome variables. To analyze the effects on a wider set of outcome variables, I adopt the following approach. In a first step, I extract an estimate of the carbon policy shock series from the monthly VAR as $CPShock_t = \mathbf{s}_1' \mathbf{\Sigma}^{-1} \mathbf{u}_t$ (see Stock and Watson, 2018). Next, I estimate the effects for the additional outcome variables $y \in Y$ using simple local projections:

$$y_{t+h} = \alpha_h + \theta_h CPShock_t + \beta_{h,1} y_{t-1} + \ldots + \beta_{h,p} y_{t-p} + \xi_{t,h}, \tag{8}$$

where θ_h is the effect on variable y at horizon h. This approach follows Cloyne et al. (2023) and Drechsel (2023), and crucially relies on the VAR assumptions

[1] To be more precise, the VAR does not have to be fully invertible for identification with external instruments. As Miranda-Agrippino and Ricco (2023) show, it suffices if the shock of interest is invertible in combination with a limited lead-lag exogeneity condition.

discussed above. Importantly, the approach also allows me to estimate the effects on lower-frequency variables, such as quarterly or annual outcomes. In this case, I aggregate the shock $CPShock_t$ by summing over the respective months before running the local projections. Using the shock series directly, rather than high-frequency surprises, increases the statistical power of these projections, as the shock series is consistently observed and spans the entire sample. However, this comes at the cost of assuming invertibility. To facilitate comparison, I normalize all responses to match the peak effect on HICP energy in the baseline VAR model.

The confidence bands are based on robust standard errors. This lag-augmentation approach obviates the need to correct for serial correlation in the residuals by controlling for lags in the regression (see Montiel Olea and Plagborg-Møller, 2021). In conducting inference, I treat the carbon policy shock as observed. However, taking the estimation uncertainty in the shock into account using bootstrapping techniques yields very similar inference, see Appendix C.7.

3.2. Empirical specification

The baseline VAR model includes eight variables capturing the state of the carbon market and the European economy. For the carbon market, I use the energy component of the Harmonised Index of Consumer Prices (HICP) as well as total GHG emissions.[2] To capture economic conditions, I include the headline HICP, industrial production, and the unemployment rate. Since the economy was at the effective lower bound for most of the sample period, I use the two-year rate as the relevant monetary policy indicator (Jarociński and Karadi, 2020). Finally, I include a stock market index and the Brent crude oil price, deflated by the HICP, as financial indicators. Further details on the data and sources are provided in Appendix A.2.

The sample period runs from January 1999—when the euro was introduced—until December 2019, ending just before the outbreak of the Covid-19 pandemic. The carbon policy surprise series, however, is only available from 2005, when the EU carbon market was established. To address this discrepancy, I censor missing values in the surprise series to zero, following the approach in Noh (2019). The motivation for using a longer sample is to increase the precision of the estimates. However, restricting the sample to 2005-2019 produces similar results.[3]

[2]GHG emissions are only available at the annual frequency. I construct a monthly emissions series by temporally disaggregating the annual data using a set of relevant monthly indicators. See Appendix A.2 for details.

[3]Note that while the carbon market was only established in 2005, the EU agreed to the Kyoto protocol in 1997 and started planning on how to meet its emission targets shortly after. The directive for establishing the EU ETS came into force in October 2003 (Directive 2003/87/EC).

Following Sims, Stock, and Watson (1990), I estimate the VAR in levels. Given the short sample, I perform a small-sample bias correction, following Hall (1992). All variables enter in log-levels, except for the unemployment rate and the two-year rate, which enter in levels. The lag order is set to six and in terms of deterministics only a constant term is included. To better control for the European sovereign debt crisis, I also include a dummy variable for the period from July 2011 to March 2012. However, the results are robust to all of these choices (see Appendix C.4).

In the local projections of the additional outcome variables on the carbon policy shock, I include only lags of the outcome variable as controls: six lags for monthly series, three for quarterly variables, and one lag for annual variables. To better account for trending behavior, I also include a linear time trend.

4. The Aggregate Effects of Carbon Pricing

In this section, I study the effects of carbon policy shocks on emissions, economic activity, prices, and green innovation. The main identifying assumption behind the external instrument approach is that the instrument is correlated with the structural shock of interest but uncorrelated with all other structural shocks. For standard inference, the instrument must also be sufficiently strong. To assess this, I conduct the weak instrument test by Montiel Olea and Pflueger (2013).

The results indicate that carbon policy surprises are strong instruments, with heteroskedasticity-robust F-statistics typically exceeding conventional critical values (see Appendix Table B.3). For my baseline instrument—the refined surprise series based on the largest information set—the robust F-statistic is 16.85. Overall, this evidence suggests that there is no weak instrument problem at hand.

4.1. The impact on emissions and the macroeconomy

I now examine the macroeconomic and environmental impacts of carbon policy shocks through the lens of the baseline VAR model. Figure 3 presents the impulse responses to the identified carbon policy shock, scaled such that the HICP energy component increases by 1 percent on impact. The solid black lines denote point estimates, while the shaded areas represent confidence bands, computed from 10,000 bootstrap replications. I report 68 and 90 percent intervals, which is standard practice in the macroeconomics literature, reflecting the relatively short samples typically available. For completeness, Appendix D.1 reproduces the main results with 95 percent confidence intervals.

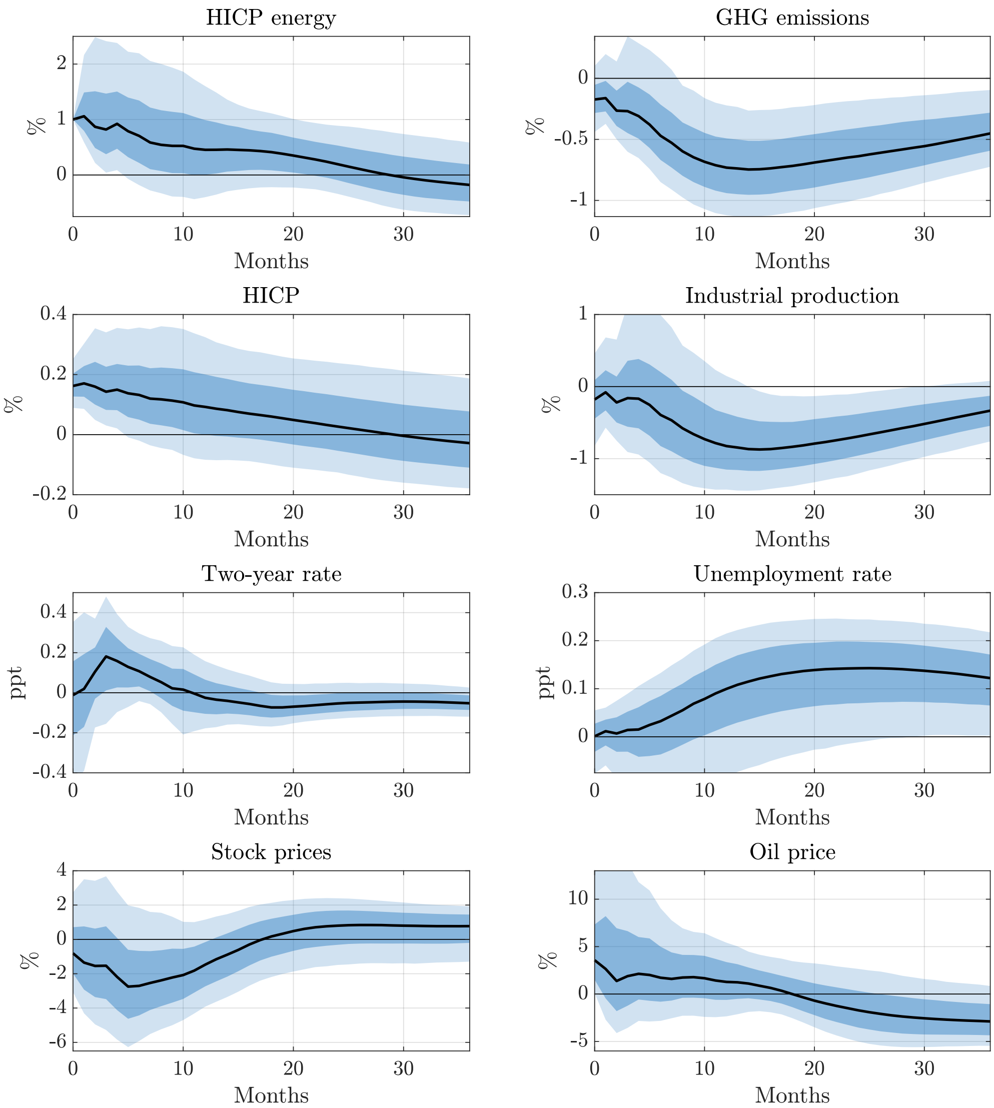

First stage regression: robust F-statistic: 16.85, R^2: 2.38%

Figure 3: Impulse Responses to a Carbon Policy Shock

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, estimated based on the VAR model (4) using the refined carbon policy surprise series from specification (d) as an instrument. Lag order: 6. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on moving-block bootstrap.

A restrictive carbon policy shock leads to a sharp, immediate increase in energy prices and a significant, persistent decline in GHG emissions. Carbon pricing effectively reduces emissions by raising the cost of emitting. I measure the price elasticity of emissions as the maximum percentage decline in emissions following a 1 percent increase in energy prices. The estimated elasticity is around 0.75, in the same ballpark as existing estimates (see e.g. Metcalf, 2019), implying potentially high decarbonization costs.

Turning to the macroeconomic variables, we can see that the fall in emissions is indeed associated with substantial macroeconomic costs. Industrial production declines significantly, by nearly 1 percent, while the unemployment rate rises by about 0.15 percentage points. Consumer prices, measured by headline HICP, increase by nearly 0.2 percent, reflecting a strong and significant pass-through to headline inflation, while pass-through to core prices is much weaker (see Appendix Figure D.3). Monetary policy does not appear to accommodate the recessionary effects, as the two-year rate tends to increase, though the response is imprecisely estimated. The stock market drops by nearly 3 percent at peak, but the response also features considerable uncertainty. While the shock raises European oil prices, the effect is modest and not statistically significant. A moderate positive impact is plausible, given that European oil producers and refineries are also covered under the emissions trading scheme.[4]

It is also worth noting that the fall in output appears to be less persistent than the reduction in emissions: three years out, emissions remain significantly below their initial level, while the industrial production response has largely recovered and is no longer significant. This suggests that the emissions intensity improves in the longer term. I will expand on this finding in Section 4.2, where I explore the effects of carbon policy shocks on green innovation.

How does carbon pricing affect the broader economy? Figure 4 shows the impulse responses of real GDP, consumption, investment, and wages, estimated using local projections on the quarterly carbon policy shock extracted from the baseline VAR. Real GDP falls significantly, by about 0.3 percent at peak. Breaking down the components, the decline in activity is primarily driven by lower consumption and investment, which fall by around 0.3 percent and 1 percent, respectively. The consumption response is particularly pronounced. In line with the rise in unemployment, we also observe a substantial decline in real wages.

Higher energy prices can affect the economy via both direct and indirect channels. They directly affect households and firms by reducing their discretionary income. Given that energy demand is inelastic, consumers and firms have less money to spend and invest after paying their energy bills (see Hamilton, 2008; Edelstein and Kilian, 2009). Energy prices also affect the economy indirectly through the general equilibrium responses of prices and wages and hence of income and employment.

The magnitudes of the estimated effects are much larger than what can be accounted for by the direct effect of higher energy prices alone. If energy demand is completely inelastic, the direct price effect is bounded by the energy share in

[4]The EU ETS covers emissions associated with exploration and drilling, production and processing, transportation, and refining of oil. This includes energy use associated with these activities and gas flaring, and may thus also affect crude oil prices. In addition, substitution away from coal-fired electricity could put further upward pressure on oil prices.

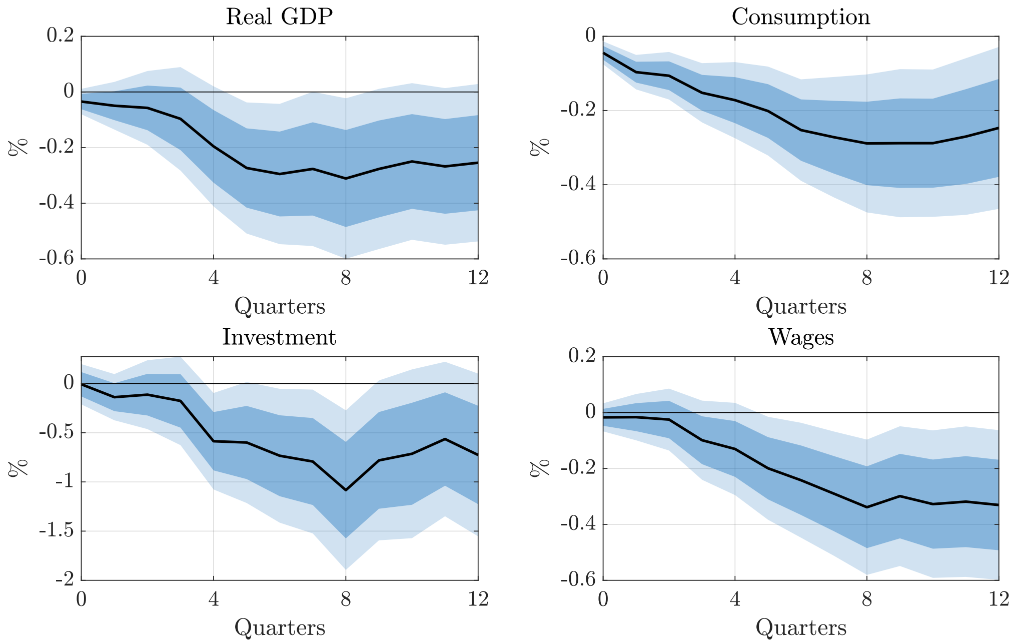

Figure 4: Responses of GDP, Consumption, Investment and Wages

Notes: Impulse responses of a selection of quarterly variables, estimated using local projections (8) of the variable of interest on the carbon policy shock from the baseline VAR. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

expenditure, which is somewhat below 10 percent in Europe. Given the shock magnitude, we would thus expect a direct impact on consumption of at most 0.1 percent ($\approx$ 1 percent $\times$ 0.1). However, the estimated consumption response is substantially larger, with a peak effect of about 0.3 percent. This suggests that indirect general equilibrium effects play a significant role in transmitting carbon policy shocks, accounting for roughly two thirds of the aggregate impact.

The additional fall in aggregate demand driven by lower employment and wages lies at the core of the indirect effect: higher carbon prices reduce households' and firms' consumption and investment, leading to a fall in output. This, in turn, puts downward pressure on employment and wages, triggering a second round of demand effects. By contrast, I find little evidence for worsening financial conditions or uncertainty as transmission channels (see Appendix Figure D.4).

To summarize, the above findings clearly illustrate a policy trade-off between reducing emissions to avert future climate damages and the economic costs of decarbonization. To explore this trade-off further, I estimate the marginal abatement cost—the cost of reducing one additional ton of CO2 equivalent—based on the empirical responses presented above. Specifically, I relate the emission reductions (in tons of CO2) to the fall in GDP (in 2019 EUR) over a 20-year horizon, conditional on a carbon policy shock at the start of my sample period in 1999. Given the short sample, impulse responses cannot be reliably estimated over such long horizons. I therefore assume that emissions stabilize at approximately –0.5 percent and output gradually returns to baseline within eight years of the shock. The GDP damages are discounted to present value terms. This back-of-the-envelope calculation gives a marginal abatement cost of 107 EUR per ton of CO2.

Interestingly, this estimate is an order of magnitude higher than the average ETS price over the sample, which was around 12 EUR per ton of CO2. The discrepancy reflects the difference between partial and general equilibrium concepts of abatement costs. The ETS price captures the marginal cost of abating emissions at the firm level within the regulated sectors, but it does not internalize broader macroeconomic spillovers. By contrast, the general equilibrium estimate incorporates second-round effects on wages, consumption, and output, which substantially raise the aggregate economic cost of abatement.

While this calculation is necessarily approximate, the order-of-magnitude gap has important implications for how marginal abatement costs are measured and used in policy debates. A common practice in cost–benefit analyses is to use allowance prices—or other static, partial-equilibrium cost estimates—as proxies for the aggregate cost of reducing emissions. My results suggest that such measures may severely understate the true economic burden once general equilibrium adjustments are taken into account.

Yet, the marginal abatement cost lies well below conventional estimates of the social cost of carbon—around $185 per ton according to Rennert et al. (2022). Therefore, a policymaker concerned with global climate damages would still choose to decarbonize. However, if the policymaker is solely focused on domestic damages, the economic case for decarbonization becomes less compelling, at least based on conventional social cost estimates.

More recent studies, however, point to potentially much higher estimates of the social cost of carbon, in excess of a $1,000 per ton of CO2 (Burke et al., 2023; Kotz, Levermann, and Wenz, 2024; Bilal and Känzig, 2024). These newer estimates completely reverse the above trade-off, as even the domestic cost of carbon in the EU becomes around 200 EUR per ton of CO2—well above the marginal abatement cost. Thus, under these estimates, unilateral, non-cooperative decarbonization policies in the EU such as the EU ETS are, in fact, cost-effective (see also Bilal and Känzig, 2025).

Another interesting finding is that the macroeconomic effects of EU ETS prices appear to be much more pronounced than those of European carbon taxes. Many European countries have introduced national carbon taxes to complement the carbon market, typically targeting sectors not covered by the EU ETS. Metcalf and Stock (2020, 2023) study the impacts of these national carbon taxes and find that, while the taxes were successful in reducing emissions, they did not lead to robust negative effects on output and employment.

A key difference is that European carbon taxes do not cover the power sector, which is included in the EU ETS and plays a central role in the macroeconomic effects I estimate. In terms of magnitudes, my results align with previous evidence on energy price shocks, such as oil supply shocks (e.g. Kilian, 2009; Baumeister and Hamilton, 2019; Känzig, 2021). Moreover, in many European countries, carbon taxes were introduced as part of broader tax reforms that often included other changes to the tax code to mitigate their impact. As I will discuss in Section 5.3, the distribution of carbon revenues plays a key role in the transmission of carbon policy shocks. Finally, EU ETS prices are far more volatile than carbon taxes, and this price uncertainty may contribute to the larger macroeconomic costs observed (see also Bilal and Stock, 2025).

In Appendix C, I perform a comprehensive series of robustness checks on the identification strategy and empirical approach used to isolate the carbon policy shock. These checks demonstrate that the results are robust along a number of dimensions including the selection of event dates, the construction of the instrument, the model specification, and the sample period. In particular, I show based on a placebo exercise that randomly-sampled dates produce impulse responses that look nothing like carbon policy shocks (see Appendix C.3).

4.2. The impact on green innovation

A key motivation for carbon pricing is to create incentives for directed technical change. Indeed, part of the vision for the EU ETS is to foster investment in clean, low-carbon technologies (European Comission, 2020). Innovation in these technologies is essential for sustaining emissions reductions without permanently lowering output.

To analyze this channel empirically, I study how patenting activity in green technologies responds to carbon policy shocks. I use data on patent applications from the Worldwide Patent Statistical Database (PATSTAT), which provides bibliographic information covering nearly the entire universe of global patent filings. Green patenting is measured using patent classification codes for climate change mitigation and adaptation technologies (Y02 and Y04s, see Haščič and Migotto, 2015), which allow for consistent and systematic identification of relevant technological advances. For further details, see Appendix A.3.

The results are presented in Figure 5. Carbon policy shocks lead to a statistically significant and economically meaningful increase in green patenting: the

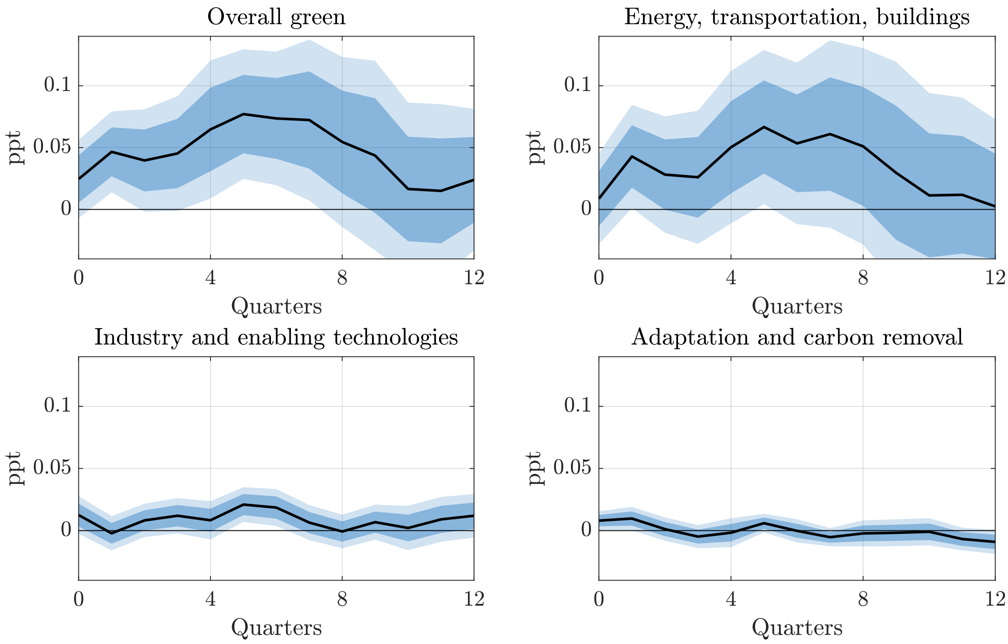

Figure 5: Patenting in Climate Change Mitigation Technologies

Notes: Impulse responses of green patenting, measured as the share of climate change mitigation patents among all biadic patent filings. The responses are estimated using local projections (8) of the patenting variables on the carbon policy shock from the baseline VAR. The figure reports responses for overall green patenting; mitigation technologies in energy generation, transportation, and buildings; mitigation in industry; and adaptation and carbon removal technologies. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

share of patents filed in climate change mitigation technologies increases by 0.08 percentage points. Given an average green patent share of around 11 percent over the sample period, this corresponds to a relative increase of about 0.75 percent. This effect is non-negligible, especially considering that the carbon policy shocks under study are relatively small. As shown in Appendix D.5, these results are robust to controlling for patent quality based on patent citations.

Which types of green patents respond the most? The most pronounced increases are observed in mitigation technologies related to energy generation, transportation, and buildings—all sectors that are relatively easy to abate. I also find a significant increase in mitigation technologies for industry. While the aggregate effect is economically smaller, this is primarily due to the lower baseline share of such patents. These results are consistent with the notion of directed technical change and suggest that carbon pricing can indeed redirect innovation toward policy-relevant areas, even over relatively short horizons.

In contrast, I find no significant response in adaptation or carbon capture technologies. These areas are typically associated with longer time horizons, greater technological uncertainty, and more complex implementation challenges.

As such, they are less likely to respond to relatively transitory changes in carbon prices. This result suggests that additional policy support—such as targeted subsidies or public R&D investment—may be needed to stimulate longer-term innovations and advance such technologies toward commercial viability.

Overall, these results align with the findings in Calel and Dechezleprêtre (2016), who use a quasi-experimental design based on installation-level inclusion criteria to show that the EU ETS increased green patenting at the firm level.

4.3. Historical importance

We have seen that carbon policy shocks can have significant effects on emissions and the economy. An equally important question is: how much of the historical variation in the variables of interest can carbon policy account for? To this end, I perform a historical decomposition exercise.

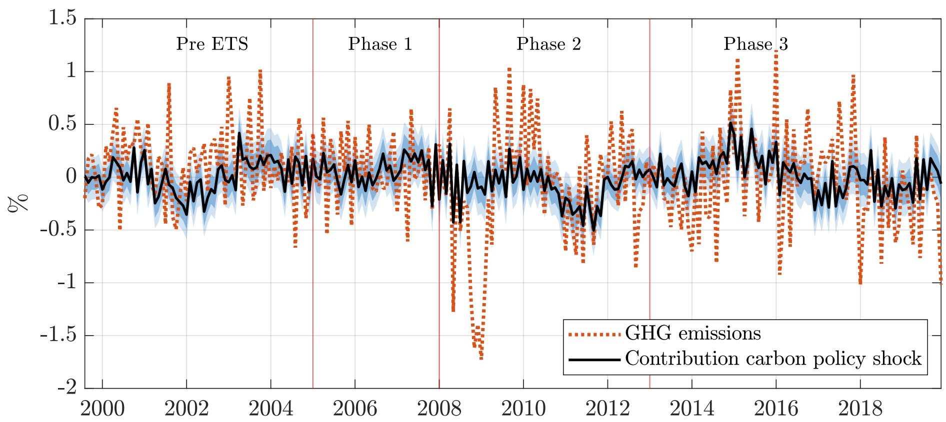

Figure 6: Historical Decomposition of GHG Emissions Growth

Notes: The figure shows the historical contribution of carbon policy shocks over the estimation sample for GHG emissions growth, estimated from the baseline VAR model (4), against the actual evolution of emissions growth. Lag order: 6. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on moving-block bootstrap.

Figure 6 shows the historical contribution of carbon policy shocks to GHG emissions growth. We can see that carbon policy shocks have contributed meaningfully to variations in GHG emissions in many episodes. Importantly, however, they cannot account for the significant fall in emissions after the global financial crisis. This suggests that the high-frequency approach is not mistakenly picking up demand-related disturbances, as the fall in emissions during the Great Recession was clearly driven by lower demand and not supply-specific factors in the European carbon market.

On average, carbon policy shocks account for about a third of the variations in emissions at horizons up to one year. Furthermore, carbon policy shocks explain a non-negligible share of the variations in energy prices and other macroeconomic (see the variance decomposition in Appendix D.3).

5. The Heterogeneous Effects of Carbon Pricing

Concerns about the unequal burden of carbon pricing have grown in Europe, especially in the context of the European Green Deal (European Comission, 2021). The situation has become even more acute since Russia's invasion of Ukraine, which led to a sharp rise in energy prices and renewed focus on energy poverty.

In light of these developments, it is crucial to better understand the distributional impact of the EU ETS. If certain groups feel left behind, this could ultimately undermine the success of climate policy. To this end, I study the heterogeneous effects of carbon pricing—first across European countries, then across households using detailed survey data from the UK. These analyses allow for a more comprehensive understanding of how carbon pricing influences economic inequality and help uncover the underlying transmission channels.

5.1. Heterogeneity across countries

How does the EU ETS affect different regions? Which countries are most severely affected by carbon policy shocks? To answer these questions, I construct a quarterly panel dataset covering advanced European countries, including the UK which was part of the EU ETS in the sample I consider. I focus on GDP impacts and examine how they vary across key country characteristics: emissions intensity, income level, share of hand-to-mouth households, and employment laws. Data on country characteristics are drawn from Almgren et al. (2022) and the EU Transaction Log. For more information on the panel dataset, see Appendix A.4.

I start by examining the average response in the panel. To this end, I estimate the following panel local projections model:

$$y_{i,t+h} = \alpha_{i,h} + \theta_h CPShock_t + \textstyle\sum_{j=1}^{p} \beta_{h,j} y_{i,t-j} + \xi_{i,t,h}, \tag{9}$$

where, $y_{i,t+h}$ is (log) real GDP in country i at time t, and $\alpha_{i,h}$ is a country fixed effect. θ_h is the dynamic causal effect of interest at horizon h. Standard errors are computed based on Driscoll and Kraay (1998).

Figure 7 presents the results. The left panel compares the average response in the panel to the aggregate response from the time series. Carbon policy shocks lead to a significant fall in real GDP on average, with the panel response closely matching the time-series evidence.

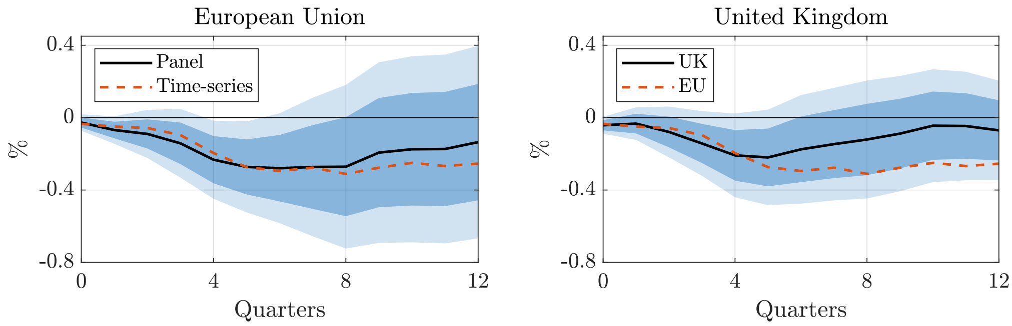

Figure 7: GDP Responses in Europe and the UK

Notes: Impulse responses of real GDP to a carbon policy shock from the baseline VAR. Left panel: average response estimated using the panel local projection (9) compared to time-series response from Section 4.1. Right panel: time-series response for the UK compared to the EU response, estimated using (8). Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid and dashed lines: point estimates. Dark and light shaded areas: 68 and 90 percent confidence bands based on Driscoll and Kraay (1998) (panel) and lag-augmentation approach (time series).

How do the impacts vary across European countries? To analyze potential heterogeneity, I interact the carbon policy shock with relevant country-specific variables in the local projections model:

$$y_{i,t+h} = \alpha_{i,h} + \delta_{t,h} + \gamma_h(CPShock_t \times d_i) + \textstyle\sum_{j=1}^{p} \beta_{h,j} y_{i,t-j} + \xi_{i,t,h} \tag{10}$$

where d_i is a dummy variable identifying countries with specific characteristics and $\delta_{t,h}$ are time fixed effects. I consider four categories: countries with a high emissions intensity (average intensity above median), low income (real GDP per capita below the 25th percentile), a high share of hand-to-mouth households (share above 0.2) and strong employment laws (employment protection index by Botero et al. (2004) above median). The panel design with interaction terms offers a key advantage: the ability to better control for potential aggregate confounders using time fixed effects.

Figure 8 shows the results. As expected, countries with a high emissions intensity are more severely affected by carbon policy shocks, displaying larger output declines. This finding is consistent with the notion that countries with more carbon-intensive economies are more exposed to changes in ETS prices. The results also point to distributional effects across countries: lower-income countries tend to experience larger output losses, though these differences, while economically meaningful, are not statistically significant. Overall, these findings are in line with the evidence in Mangiante (2024) and Känzig and Konradt (2024) on the regional impacts of carbon policy shocks.

To shed more light on the role of indirect, general equilibrium effects via in-

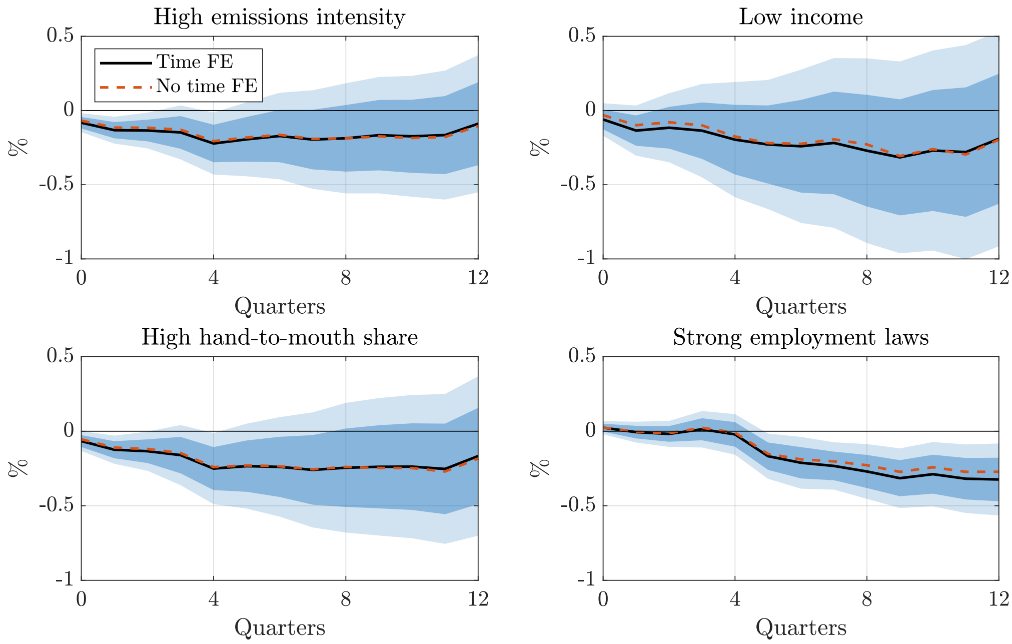

Figure 8: Heterogeneity in Country-level GDP Responses

Notes: Relative impulse responses of country-level real GDP, estimated from the panel local projection (10), with and without time fixed effects. The carbon policy shock from the baseline VAR is interacted with country characteristics: above-median emissions intensity, below-25th percentile real GDP per capita, hand-to-mouth household share above 0.2, and above-median employment protection index. The shock is normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid and dashed lines: point estimates. Dark and light shaded areas: 68 and 90 percent confidence bands based on Driscoll and Kraay (1998).

come and employment, I analyze how the impacts vary with the share of hand-to-mouth households and the degree of labor market regulation. Countries with a high share of hand-to-mouth households exhibit considerably larger output declines, though the response is only significant at a few horizons. Likewise, countries with stricter employment protection laws show greater output losses, especially at longer horizons. This delayed response is consistent with the notion that employment effects may take time to fully materialize.

Reassuringly, omitting the time fixed effects does not change the results materially. A specification that excludes them and includes the base effect of carbon policy shocks instead yields virtually identical estimates. This suggests that the identification strategy adequately accounts for potential aggregate confounders.

5.2. Distributional effects across households

The differential impacts across European countries may mask important within-country heterogeneity. I thus turn now to the heterogeneous effects across households within a given country. This requires detailed micro data on consumption expenditure and income at regular frequency over the past two decades. Unfortunately, such data does not exist for most European countries let alone at the EU level. Therefore, I focus here on the UK which is one of the few countries where such data is available.

The impact on aggregate economic activity in the UK is comparable to that at the EU level, though the responses are less persistent (see Figure 7 and Appendix Figure E.1). This helps alleviate concerns about external validity, as the distributional effects are, if anything, likely to be attenuated relative to countries with more pronounced GDP impacts.[5]

The British Living Costs and Food Survey (LCFS) is the main source of household spending data in the UK, providing high-quality, detailed information on expenditure, income, and household characteristics. It is conducted annually, with interviews carried out throughout the year and across the entire UK. I use data from the last 20 waves, covering the period 1999–2019, to construct a repeated cross-section. Each wave includes approximately 6,000 households, resulting in over 120,000 observations in total. To compute income and expenditure measures, I first convert household-level variables to per capita terms by dividing by the number of household members. I then deflate the resulting variables using the harmonized consumer price index to express them in real terms. For further details, see Appendix A.5.

Ideally, we would observe how consumption expenditure and income evolve over time at the individual level. However, the LCFS does not feature such a panel dimension. To address this, I construct a pseudo-panel using a grouping estimator, following Browning, Deaton, and Irish (1985).

A natural dimension for grouping households is their income. However, since current income may endogenously respond to the shock of interest, it cannot be used as the grouping variable. Fortunately, the LCFS collects not only current household income but also *normal* income, which serves as a proxy for permanent income.[6] Based on normal disposable household income, I group households into three pseudo-cohorts: low-, middle-, and high-income households. Following Cloyne and Surico (2017), each household is assigned to a calendar quarter based on the interview date. Within each quarter of a given year, households are classified as low-income (bottom 25 percent of the normal income distribution), middle-income (middle 50 percent), or high-income (top 25 percent).

[5]To further mitigate external validity concerns, I show that the results in other European countries are comparable, using similar survey data for Denmark and Spain, see Appendix Figure E.13.

[6]I verify that normal income does not respond significantly to the carbon policy shock, see Appendix Figure E.6. As a robustness check, I group households using an estimate of permanent income derived from a Mincerian-type regression. The results remain robust, see Appendix E.5.

Table 2: Descriptive Statistics on Households in the LCFS

	Overall	By income group		
		Low-income	Middle-income	High-income
Income and expenditure				
Normal disposable income	6,748	3,740	6,807	10,866
Total expenditure	4,458	3,025	4,444	6,238
Energy share	7.2	9.5	7.2	5.2
Non-durables (excl. energy) share	81.5	81.6	81.6	81.3
Durables share	11.2	8.9	11.2	13.5
Household characteristics				
Age	51	47	54	49
Education (share with post-comp.)	34.0	25.7	29.7	51.2
Housing tenure				
Social renters	20.8	46.9	17.4	3.7
Mortgagors	42.3	25.5	41.3	60.0
Outright owners	36.9	27.7	41.3	36.4

Notes: Descriptive statistics on quarterly household income and expenditure (in 2015 pounds), the breakdown of expenditure into energy, non-durable goods and services excl. energy, and durables (as a share of total expenditure) as well as a selection of household characteristics, both over all households and by income group. For variables in levels such as income, expenditure and age the median is shown while the shares are computed based on the mean of the corresponding variable. The expenditure shares are expressed relative to total expenditure excluding housing, and semi-durables are subsumed under the non-durable category. Age corresponds to the age of the household reference person and education is proxied by whether a household member has completed a post-compulsory education.

Individual-level variables are then aggregated using survey weights to ensure representativeness of the British population.

Table 2 presents descriptive statistics, for the entire population and by income group. I focus here on expenditure excluding housing, however, the results including housing are similar. As expected, quarterly household expenditure increases with income. Low-income households allocate a larger share of their budget to non-durable goods, while higher-income households spend more on durables. Notably, poorer households devote a significantly larger share of their expenditure to energy: nearly 10 percent for low-income households, just over 7 percent for middle-income, and around 5 percent for high-income households. This pattern implies that, to the extent energy demand is inelastic, poorer households are more exposed to increases in energy prices.

The different income groups are broadly similar in terms of age. Higher-income households, however, tend to have higher levels of education and are more likely to own their homes, either outright or with a mortgage.

Heterogeneity by household income. We can now study how household expenditure and income respond to carbon policy shocks—and, more importantly,

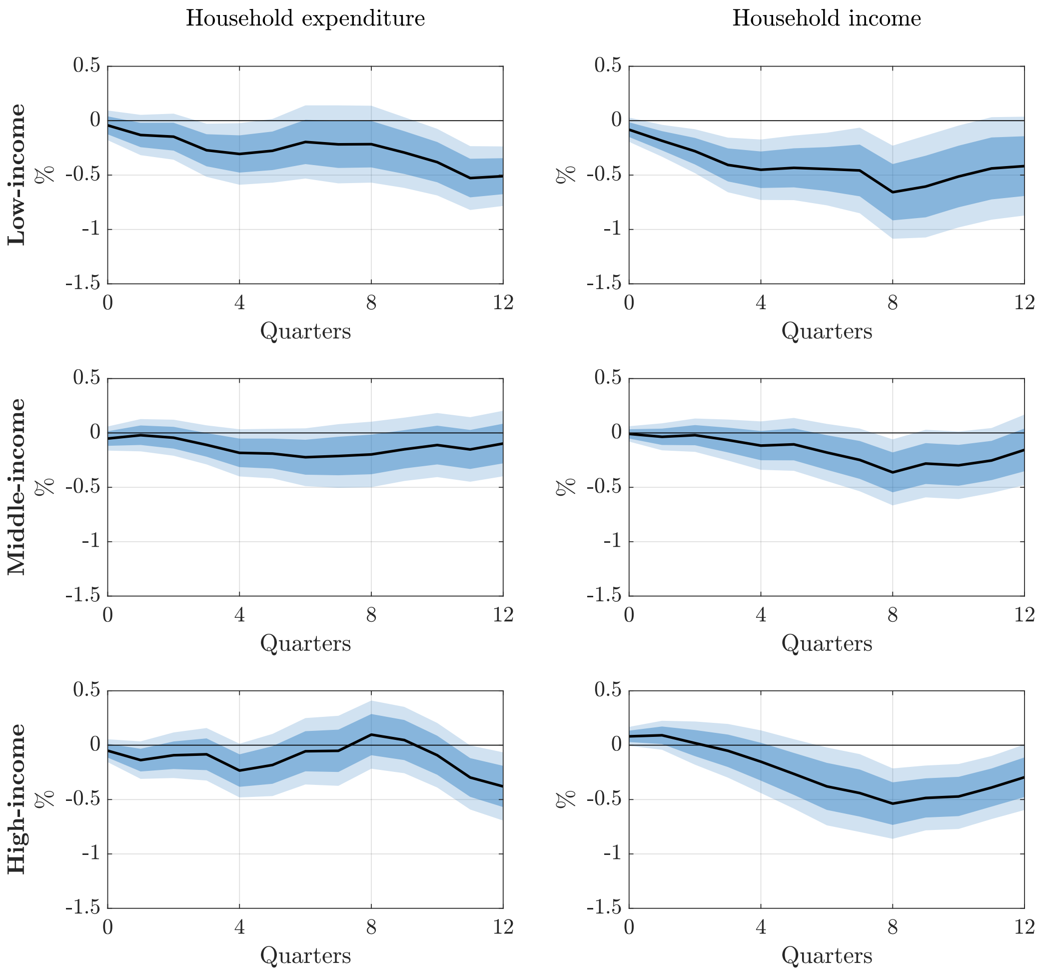

Figure 9: Household Expenditure and Income Responses by Income Group

Notes: Impulse responses of total expenditure (excluding housing) and current total household income for low-income (bottom 25 percent), middle-income (middle 50 percent) and high-income households (top 25 percent). The responses are estimated using local projections (8) of the household variables on the carbon policy shock from the baseline VAR. Households are grouped by total normal income and the responses are computed based on group medians. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

how these responses vary across groups. Figure 9 shows the responses of total expenditure and current income for the three income groups, estimated using local projections (8).[7] As before, the solid black lines represent point estimates, and the dark and light shaded areas are 68 and 90 percent confidence intervals.

The expenditure response varies meaningfully across income groups. Low-income households exhibit the strongest reaction, reducing their spending significantly and persistently. By contrast, higher-income households show smaller and

[7]To reduce the noise inherent in survey data, I smooth the expenditure and income series using a backward-looking moving average, following Cloyne, Ferreira, and Surico (2020). The results are robust when using the raw series (though responses are more jagged and less precise) or based on smooth local projections (Barnichon and Brownlees, 2019), see Appendix Figure E.5.

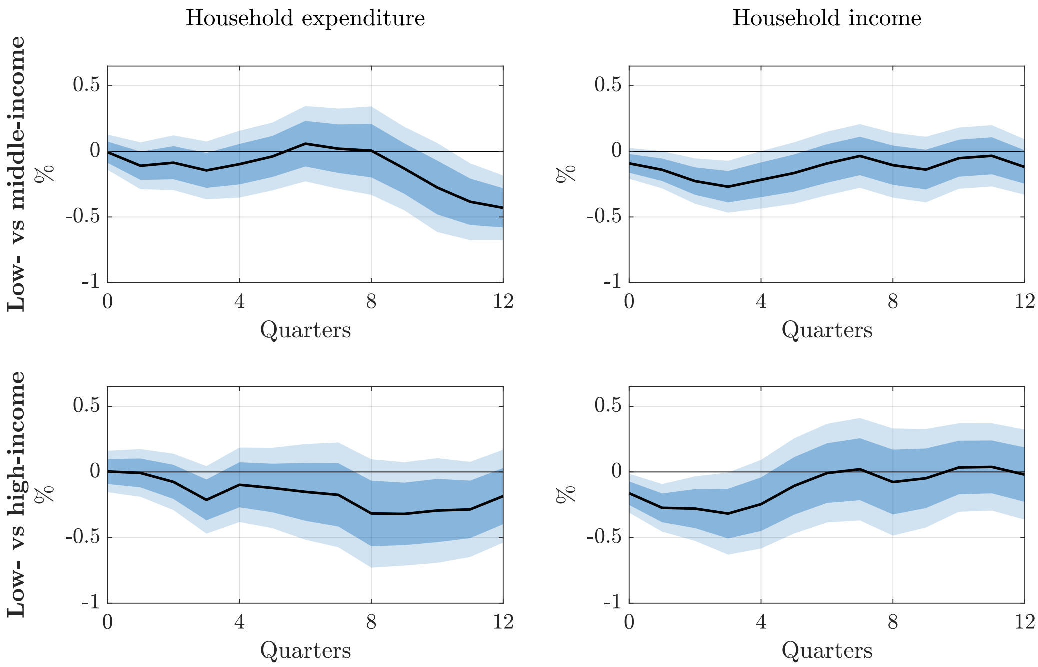

Figure 10: Group Differences in Expenditure and Income Responses

Notes: Difference in expenditure and income responses of low-income households (bottom 25 percent) versus middle- (middle 50 percent) and high-income households (top 25 percent), respectively. The responses are estimated using local projections (8) on the carbon policy shock from the baseline VAR. Households are grouped by total normal income and the responses are computed based on differences in group medians. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

more short-lived responses. For middle-income households, the expenditure response is not statistically significant at the 90 percent level, while for high-income households, it is borderline significant at a few horizons.

A similar pattern emerges on the income side. Poorer households experience the largest and most immediate decline in income. Middle-income households display a more muted response, while high-income households also see a non-negligible decline—though it takes considerably longer to materialize.

Low-income households are more affected by carbon policy shocks in two ways. First, they face a more significant increase in their energy bill, reflecting their higher baseline energy share and low demand elasticity. Second, they face a stronger fall in their income. Taken together, these findings help explain why poorer households cut their expenditure the most after a carbon policy shock.

Are these differences not only economically, but also statistically significant? Figure 10 shows the difference in the response of low- and higher-income households. While there is considerable estimation uncertainty, there is evidence that low-income households experience significantly larger declines in both income and expenditure—at some horizons, these differences are statistically significant even at the 90 percent level. One exception is the difference in the expenditure response between low- and high-income households: although economically sizable, it is only statistically significant at the 68 percent level.

At this stage, it is worth discussing a potential concern regarding selection. The assignment into income groups is not random, and other characteristics may partly explain the observed heterogeneity. To mitigate this concern, I re-group households based on alternative characteristics, including age, education, and housing tenure. However, none of these alternative groupings replicate the patterns observed for income, suggesting that I am not spuriously picking up differences in other household characteristics (see Figures E.8-E.10 in the Appendix).

Direct versus indirect effects. Carbon policy shocks have meaningful aggregate and distributional effects. The results suggest that indirect general equilibrium effects play an important role in the transmission, accounting for a large share of the aggregate impact. The richness of the household expenditure and income data allows us to further dissect the transmission mechanisms at play.

Table 3 presents the cumulative impact of a carbon policy shock over the three year impulse horizon on household expenditure—broken down into energy, non-durables excluding energy, and durables—as well as on household income, both in the aggregate and by income group. Panel A shows the cumulative response in percent (see Appendix Figure E.11 for the corresponding impulse responses).

Energy expenditure rises substantially following the shock, increasing by around 7 percent cumulatively, though the response is not statistically significant at the 90 percent level. Durable expenditure also declines markedly, but the estimates are imprecise. The most robust impacts are observed for non-durable spending and household income, which fall cumulatively by around 2 and 3 percent, respectively.

The results reinforce the previously documented heterogeneity: low-income households experience the largest responses, with non-durable expenditure falling cumulatively by 3.5 percent and income by nearly 5 percent. Higher-income households adjust their expenditure less, despite experiencing significant income declines—income falls by around 2 percent for middle-income and 3 percent for high-income households. This pattern is consistent with the notion that wealthier households are less likely to be financially constrained and are thus better able to smooth transitory income shocks. The findings also align with the panel evidence showing that countries with a higher share of hand-to-mouth households are more adversely affected by carbon policy shocks.

Table 3: Cumulative Impact on Household Expenditure and Income

	Overall	By income group		
		Low-income	Middle-income	High-income
Panel A: Percent changes				
Expenditure				
Energy	7.09 [-0.08, 14.26]	6.71 [-0.24, 13.65]	7.08 [-0.18, 14.33]	7.51 [-0.09, 15.11]
Non-durables excl. energy	-1.97 [-3.41, -0.54]	-3.46 [-5.64, -1.28]	-1.78 [-3.22, -0.34]	-0.88 [-3.14, 1.37]
Durables	-3.40 [-15.67, 8.86]	-6.34 [-12.23, -0.46]	0.37 [-4.31, 5.06]	-8.01 [-12.73, -3.29]
Income	-2.95 [-5.09, -0.80]	-4.91 [-7.55, -2.27]	-1.95 [-3.69, -0.20]	-2.99 [-4.89, -1.09]
Panel B: Changes in pounds				
Expenditure				
Energy	17.36 [0.09, 34.63]	15.39 [-0.16, 30.95]	17.49 [-0.10, 35.09]	19.06 [-0.05, 38.17]
Non-durables excl. energy	-71.19 [-123.08, -19.31]	-93.07 [-152.00, -34.13]	-71.08 [-128.59, -13.57]	-49.54 [-176.54, 77.46]
Durables	-7.39 [-34.80, 20.01]	-3.69 [-7.28, -0.10]	0.58 [-6.89, 8.06]	-27.05 [-43.60, -10.50]
Income	-189.30 [-326.85, -51.75]	-181.72 [-279.75, -83.68]	-128.61 [-243.54, -13.68]	-318.27 [-519.77, -116.76]

Notes: The cumulative impact of a carbon policy shock on expenditure and income over the three-year impulse horizon, based on responses in Figure 9 and Appendix Figure E.11. Panel A: cumulative change in percent, calculated as the present discounted value of the impulse response. Panel B: overall pound change in quarterly expenditure and income (in 2015 pounds), computed as the cumulative percent change, multiplied by the corresponding average quarterly expenditure/income. Bootstrapped 90 percent confidence intervals are reported in brackets.

Panel B converts the cumulative responses into equivalent changes in income and expenditure, expressed in pounds (GBP), over the three-year impulse horizon. Overall quarterly expenditure decreases by approximately GBP 60 while income falls by nearly GBP 190.

While the increase in energy expenditure is large in relative terms, the absolute pound change is more modest, as energy only accounts for a small share of total spending. Still, the results suggest that energy demand is relatively inelastic. As a benchmark, assuming that energy demand is completely inelastic, we would expect an increase in energy expenditure of GBP 20.5, computed as the quarterly energy spending (GBP 321) times the cumulative response of the relative price of energy (6.4 percent). This is only slightly larger than the estimated GBP 17.4, implying a short-run price elasticity of energy of around 0.2, consistent with estimates in Labandeira, Labeaga, and López-Otero (2017). Interestingly, there is little heterogeneity in the energy expenditure response, suggesting that energy demand is inelastic across households—at least in the short run.

These findings also confirm the important role of indirect effects operating through income and employment. While energy bills rise considerably, this direct channel accounts for less than a third of the overall expenditure decline ($|17.36/61.22|$). Although income falls across all groups, the drop in expenditure relative to income is largest for low-income households—supporting the idea that this group includes a disproportionate share of hand-to-mouth households.

Accounting for indirect, general equilibrium effects is important for two key reasons. First, to assess the aggregate economic impacts of carbon pricing policies. My findings contribute to the literature emphasizing the role of marginal propensity to consume (MPC) heterogeneity combined with unequal income incidence for the transmission of economic shocks (Bilbiie, 2008; Auclert, 2019; Patterson, 2023, among others). Failing to account for indirect effects—through prices, wages, and thus income and employment—risks understating the macroeconomic consequences of carbon pricing.

Second, accounting for indirect effects is important to understand the distributional impact of carbon pricing. Focusing on the direct effect via energy expenditure may underestimate the actual distributional effects. The uncovered expenditure heterogeneity, albeit subject to considerable estimation uncertainty, is notable. According to my point estimates, low-income households account for over 30 percent of the aggregate effect of carbon pricing on consumption, even though they make up for a much smaller share of consumption in normal times (around 15 percent).

What drives the income response? We have seen considerable heterogeneity in households' income responses. I now examine the underlying drivers of income incidence across household groups. There are at least two potential sources of heterogeneity. First, households differ in their labor income, for example because they work in different sectors. Second, they differ in their income composition, with some households receiving financial income in addition to labor earnings. I focus here on the first channel, which is more relevant for understanding heterogeneity at the lower end of the income distribution. I study the role of income composition in Appendix E.8.

To investigate potential heterogeneity in labor income, I study how the re-

Table 4: Sectoral Distribution of Employment

Sectors	Overall	By income group		
		Low-income	Middle-income	High-income
Energy-intensity				
High	21.6	9.8	25.6	25.8
Lower	78.4	90.2	74.4	74.2
Demand-sensitivity				
High	30.5	49.0	27.2	18.1
Lower	69.5	51.0	72.8	81.9

Notes: Statistics on the sectoral employment distribution of households in the LFS, both overall and by income group. I group sectors along two dimensions: their energy intensity and their demand sensitivity. The energy-intensive sectors include agriculture, utilities, transportation, and manufacturing. The demand-sensitive sectors include construction, wholesale and retail trade, hospitality, and entertainment and recreation.

sponses vary by sector of employment using data from the UK Labour Force Survey (LFS).[8] I consider two dimensions for grouping sectors. First, I group sectors by their energy intensity to capture the role of the conventional cost channel. Second, I group sectors based on their sensitivity to aggregate demand.[9]

Table 4 presents descriptive statistics on the sectoral distribution of households, both overall and by income group. Few low-income households work in energy-intensive sectors such as utilities or manufacturing, suggesting that sectoral energy intensity is unlikely to fully explain the heterogeneous income responses observed. Demand sensitivity on the other hand appears to be a more relevant source of heterogeneity: low-income households are disproportionately employed in sectors that are more sensitive to aggregate fluctuations—such as retail or hospitality—while the majority of higher-income households work in sectors that are less demand-sensitive.

Figure 11 shows how the median income across different sectors responds to a carbon policy shock. Interestingly, income declines are broadly similar across sectors with high and low energy intensity. In contrast, the sectoral differences by demand sensitivity are more pronounced. Households employed in demand-sensitive sectors experience the largest and most statistically significant declines in income, while those in less demand-sensitive sectors show more muted responses. These more cyclical sectors face a stronger decrease in demand after

[8]The LCFS does not include any information on the sector of employment. I therefore use annual data from the LFS, which provides detailed information on both employment sector and income. For more details on the LFS, see Appendix A.5.

[9]I measure demand sensitivity by estimating the elasticity of sectoral labor income with respect to aggregate income. Sectors producing more discretionary goods and services tend to be more demand-sensitive. See Appendix E.8 for details.

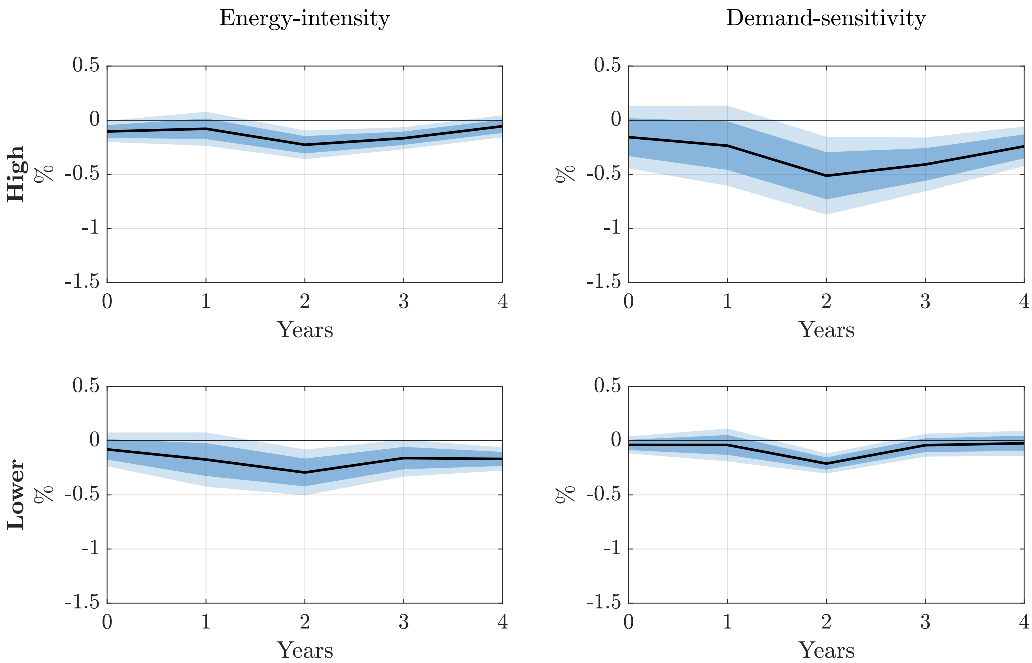

Figure 11: Income Response by Sector of Employment

Notes: Impulse responses of income (pay from main and second job) in different sectors, grouped by their energy-intensity and demand-sensitivity (see Table 4 for a detailed sectoral breakdown). The responses are estimated using local projections (8) of the household variables on the carbon policy shock from the baseline VAR. The responses are computed based on group medians and normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 1 lag of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

a carbon policy shock, also because households cut expenditure more in these sectors, and thus react by laying off workers and cutting compensation. Since low-income households are overrepresented in these sectors, they are disproportionately affected, which helps explain the strong income response for this group.

These results support the notion that carbon policy shocks transmit to the economy not only through the traditional cost channel, but also via the demand side—consistent with earlier evidence in Kilian and Park (2009) on the transmission of energy price shocks. A novel insight from my analysis is that, in the presence of household heterogeneity, the demand channel may become even more important. This finding contributes to a growing literature on Keynesian supply shocks (see e.g., Guerrieri et al., 2022; Cesa-Bianchi and Ferrero, 2021).

Alternative channels. Thus far, we focused on the direct effect via energy prices and the indirect effects through income. While other channels may also be at work, several reasons suggest they are unlikely to play a major role in the transmission of carbon policy shocks. First, carbon pricing may affect the prices of other goods through substitution effects, which could influence household budgets. However, as shown in Section 4.1, the response of core consumer prices is muted and only marginally significant, suggesting this channel is limited in scope. Second, carbon policy may influence durable expenditure—either through uncertainty or precautionary motives, or due to reductions in durables that are complementary with energy use (see also Edelstein and Kilian, 2009). Yet, the overall response of durable spending is quantitatively small, suggesting these channels are limited as well. Finally, households might adjust their saving behavior in response to rising interest rates. However, this channel is likely more relevant for high-income households, who account for a smaller share of the aggregate consumption response.

5.3. Policy implications

We have seen that that the economic costs of carbon pricing tend to be unevenly distributed across society: poorer households are more affected—both directly and indirectly—while higher-income households are less impacted.

These findings suggest that targeted fiscal policies can help alleviate the economic burden of climate change mitigation and ease the trade-off between reducing emissions and maintaining economic activity. Importantly, since energy demand is inelastic, targeted transfers need not undermine the environmental effectiveness of carbon pricing.

Such policies could be implemented by recycling a portion of the ETS auction revenues. Notably, the current EU ETS does not feature such a direct redistribution mechanism. Auction revenues are almost entirely allocated to climate- and energy-related purposes—both domestically and internationally—but not to alleviate the economic burden on households.[10]

The above intuition is confirmed in a New Keynesian model with a climate block à la Golosov et al. (2014) and heterogeneity in households' energy expenditure shares, income incidence and MPCs. Calibrated to match key empirical macro and micro moments, the model suggests that redistributing carbon revenues to high MPC households can reduce aggregate consumption losses by around 40 percent at the expense of a comparatively modest change in emission reductions. The model also illustrates that household heterogeneity plays an important role in the transmission of carbon policy and helps reconcile the empirical responses. Appendix F presents the model and its implications in detail.

Another important argument for cushioning the distributional impact of carbon pricing is that a successful climate transition requires broad public support.

[10]There are only some indirect solidarity measures in place, such as through the Cohesion Fund or the Just Transition Fund.

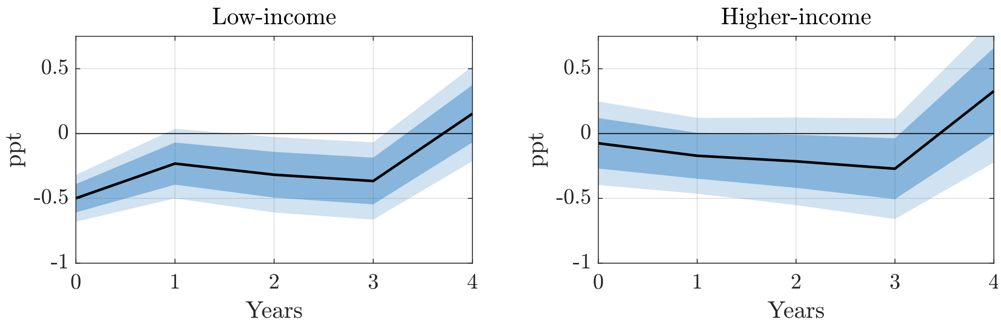

Figure 12: Effect on Attitudes Towards Climate Policy

Notes: Impulse responses of public attitude towards climate policy for low-income (bottom 25 percent) and higher-income (top 75 percent) groups. The responses are estimated using local projections (8) of the household variables on the carbon policy shock from the baseline VAR. Public attitude towards climate policy is proxied by the share of households in the BSA survey that express support for environmentally-motivated fuel taxes. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 1 lag of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

If certain groups feel left behind, this can undermine the effectiveness of climate policy—as illustrated, for example, by the gilets jaunes movement in France, which began as a protest against higher fuel taxes (see also Knittel, 2014).

I find some support for this hypothesis using data from the British social attitudes (BSA) survey. The annual survey captures public opinion on a wide range of topics and serves as an important barometer of attitudes in the UK. To proxy support for climate policy, I draw on a question about the approval of environmentally motivated fuel taxes (see Appendix E.10 for more information).

Figure 12 shows how the approval rate for environmentally motivated tax policies responds to a carbon policy shock across income groups. While the response among higher-income households is not statistically significant, support among low-income households declines significantly. Recall that these households are also the most adversely affected by carbon policy shocks. These results suggest that compensating the most affected groups could help strengthen public support for climate change mitigation—consistent with recent findings in Anderson, Marinescu, and Shor (2019) and Dechezleprêtre et al. (2025).

6. Conclusion

This paper studies the impacts of carbon pricing on the environment and the economy, exploiting a new identification strategy and data from the European carbon market. A shock tightening the carbon pricing regime leads to a sharp increase in energy prices, a persistent reduction in greenhouse gas emissions, and an uptick in green innovation. However, these gains come at the cost of temporarily lower economic activity and higher inflation. The estimated magnitudes are economically meaningful: a 1 percent increase in energy prices reduces emissions by 0.75 percent, and output and consumption by around 0.3 percent at peak. The consumption response is much larger than what can be accounted for by the direct effect via energy prices alone. This suggests that indirect, general equilibrium channels—operating through income and employment—play a key role, accounting for roughly two-thirds of the total consumption effect.

These findings point to a policy trade-off between environmental goals and the economic costs of decarbonization. Based on the observed responses, I estimate a marginal abatement cost of around 100 EUR per ton of CO2. Importantly, these costs are not borne equally across society. Poorer households reduce their consumption significantly while richer households are less affected. My results suggest that targeted transfers to the most affected groups could help mitigate the economic burden of decarbonization while preserving its environmental impact. In future work, it would be interesting to explore how climate and fiscal policy should be coordinated to support a successful and equitable transition to a low-carbon economy.

References

Almgren, Mattias, José-Elías Gallegos, John Kramer, and Ricardo Lima. 2022. "Monetary policy and liquidity constraints: Evidence from the euro area." *American Economic Journal: Macroeconomics*, 14(4): 309–340.

Anderson, Soren T., Ioana Marinescu, and Boris Shor. 2019. "Can Pigou at the Polls Stop Us Melting the Poles?"

Andersson, Julius J. 2019. "Carbon Taxes and CO2 Emissions: Sweden as a Case Study." *American Economic Journal: Economic Policy*, 11(4): 1–30.

Auclert, Adrien. 2019. "Monetary policy and the redistribution channel." *American Economic Review*, 109(6): 2333–67.

Barnichon, Regis and Christian Brownlees. 2019. "Impulse response estimation by smooth local projections." *Review of Economics and Statistics*, 101(3): 522–530.

Bauer, Michael D. and Eric T. Swanson. 2023*a*. "An Alternative Explanation for the "Fed Information Effect"." *American Economic Review*, 113(3): 664–700.

Bauer, Michael D. and Eric T. Swanson. 2023*b*. "A reassessment of monetary policy surprises and high-frequency identification." *NBER Macroeconomics Annual*, 37(1): 87–155.

Baumeister, Christiane and James D. Hamilton. 2019. "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks." *American Economic Review*, 109(5): 1873–1910.

Baumeister, Christiane and Lutz Kilian. 2017. "A general approach to recovering market expectations from futures prices with an application to crude oil."

Bernard, Jean-Thomas and Maral Kichian. 2021. "The Impact of a Revenue-Neutral Carbon Tax on GDP Dynamics: The Case of British Columbia." *The Energy Journal*, 42(3).

Beznoska, Martin, Johanna Cludius, and Viktor Steiner. 2012. "The incidence of the European Union Emissions Trading System and the role of revenue recycling: Empirical evidence from combined industry-and household-level data."

Bigio, Saki, Diego R. Känzig, Pablo Sanchez, and Conor Walsh. 2025. "Carbon Pricing and Inequality: A Normative Perspective."

Bilal, Adrien and Diego R. Känzig. 2024. "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature." National Bureau of Economic Research.

Bilal, Adrien and Diego R. Känzig. 2025. "Does Unilateral Decarbonization Pay For Itself?" National Bureau of Economic Research.

Bilal, Adrien and James H. Stock. 2025. "Macroeconomics and Climate Change."

Bilbiie, Florin O. 2008. "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic." *Journal of Economic Theory*, 140(1): 162–196.

Bjørnland, Hilde Christiane, Jamie Cross, and Felix Kapfhammer. 2024. *The drivers of emission reductions in the European carbon market.* Working Paper.

Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. "The regulation of labor." *The Quarterly Journal of Economics*, 119(4): 1339–1382.

Browning, Martin, Angus Deaton, and Margaret Irish. 1985. "A profitable approach to labor supply and commodity demands over the life-cycle." *Econometrica*, 503–543.

Burke, Marshall, Mustafa Zahid, Noah Diffenbaugh, and Solomon M Hsiang. 2023. "Quantifying climate change loss and damage consistent with a social cost of greenhouse gases." National Bureau of Economic Research.

Bushnell, James B., Howard Chong, and Erin T. Mansur. 2013. "Profiting from regulation: Evidence from the European carbon market." *American Economic Journal: Economic Policy*, 5(4): 78–106.

Calel, Raphael and Antoine Dechezleprêtre. 2016. "Environmental policy and directed technological change: evidence from the European carbon market." *Review of Economics and Statistics*, 98(1): 173–191.

Cesa-Bianchi, Ambrogio and Andrea Ferrero. 2021. "The transmission of Keynesian supply shocks."

Chow, Gregory C. and An-loh Lin. 1971. "Best Linear Unbiased Interpolation, Distribution, and Extrapolation of Time Series by Related Series." *The Review of Economics and Statistics*, 53(4): 372–375.

Cieslak, Anna. 2018. "Short-rate expectations and unexpected returns in treasury bonds." *The Review of Financial Studies*, 31(9): 3265–3306.

Cloyne, James and Paolo Surico. 2017. "Household debt and the dynamic effects of income tax changes." *The Review of Economic Studies*, 84(1): 45–81.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary policy when households have debt: new evidence on the transmission mechanism." *The Review of Economic Studies*, 87(1): 102–129.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico. 2023. "Monetary policy, corporate finance, and investment." *Journal of the European Economic Association*, 21(6): 2586–2634.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J. Wagner. 2025. "Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading system." *Review of Economic Studies*, 92(3): 1625–1660.

Cui, Jingbo, Chunhua Wang, Junjie Zhang, and Yang Zheng. 2021. "The effectiveness of China's regional carbon market pilots in reducing firm emissions." *Proceedings of the National Academy of Sciences*, 118(52): e2109912118.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2025. "Fighting climate change: International attitudes toward climate policies." *American Economic Review*, 115(4): 1258–1300.

Dechezleprêtre, Antoine, Tomasz Koźluk, Tobias Kruse, Daniel Nachtigall, and Alain De Serres. 2019. "Do environmental and economic performance go together? A review of micro-level empirical evidence from the past decade or so." *International Review of Environmental and Resource Economics*, 13(1-2): 1–118.

DEHSt, German Emissions Trading Authority. 2019. "German Auctioning of Emission Allowances: Periodical Report: Annual Report 2018."

Drechsel, Thomas. 2023. "Earnings-Based Borrowing Constraints and Macroeconomic Fluctuations." *American Economic Journal: Macroeconomics*, 15(2): 1–34.

Driscoll, John C. and Aart C. Kraay. 1998. "Consistent covariance matrix estimation with spatially dependent panel data." *Review of Economics and Statistics*, 80(4): 549–560.

Edelstein, Paul and Lutz Kilian. 2009. "How sensitive are consumer expenditures to retail energy prices?" *Journal of Monetary Economics*, 56(6): 766–779.

European Comission. 2020. "EU Emissions Trading System (EU ETS)." *https://ec.europa.eu/clima/policies/ets_en*, [Online; accessed 02-Dec-2020].

European Comission. 2021. "Fit for 55: delivering the EU's 2030 Climate Target on the way to climate neutrality." *https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021DC0550*, [Online; accessed 30-Aug-2021].

Fabra, Natalia and Mar Reguant. 2014. "Pass-through of emissions costs in electricity markets." *American Economic Review*, 104(9): 2872–99.

Fan, Ying, Jun-Jun Jia, Xin Wang, and Jin-Hua Xu. 2017. "What policy adjustments in the EU ETS truly affected the carbon prices?" *Energy Policy*, 103: 145–164.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. "What do emissions markets deliver and to whom? Evidence from Southern California's NOx trading program." *American Economic Review*, 102(2): 965–993.

Gertler, Mark and Peter Karadi. 2015. "Monetary policy surprises, credit costs, and economic activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. "Optimal taxes on fossil fuel in general equilibrium." *Econometrica*, 82(1): 41–88.

Goulder, Lawrence H. 1995. "Environmental taxation and the double dividend: a reader's guide." *International tax and public finance*, 2: 157–183.

Goulder, Lawrence H. and Marc A.C. Hafstead. 2018. *Confronting the Climate Challenge.* Columbia University Press.

Goulder, Lawrence H., Marc A.C. Hafstead, GyuRim Kim, and Xianling Long. 2019. "Impacts of a carbon tax across US household income groups: What are the equity-efficiency trade-offs?" *Journal of Public Economics*, 175: 44–64.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning. 2022. "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review*, 112(5): 1437–74.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements." *International Journal of Central Banking*, 1: 55–93.

Hall, Peter. 1992. *The bootstrap and Edgeworth expansion.* Springer Science & Business Media.

Hamilton, James D. 2008. "Oil and the macroeconomy." *The New Palgrave Dictionary of Economics*.

Haščič, Ivan and Mauro Migotto. 2015. "Measuring environmental innovation using patent data."

Hintermann, Beat. 2016. "Pass-through of CO2 emission costs to hourly electricity prices in Germany." *Journal of the Association of Environmental and Resource Economists*, 3(4): 857–891.

Jarociński, Marek and Peter Karadi. 2020. "Deconstructing monetary policy surprises—the role of information shocks." *American Economic Journal: Macroeconomics*, 12(2): 1–43.

Jentsch, Carsten and Kurt G. Lunsford. 2019. "The dynamic effects of personal and corporate income tax changes in the United States: Comment." *American Economic Review*, 109(7): 2655–78.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review*, 95(1): 161–182.

Känzig, Diego R. and Maximilian Konradt. 2024. "Climate Policy and the Economy: Evidence from Europe's Carbon Pricing Initiatives." *IMF Economic Review*, 72(3): 1081–1124.

Kilian, Lutz. 2009. "Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market." *American Economic Review*, 99(3): 1053–69.

Kilian, Lutz and Cheolbeom Park. 2009. "The impact of oil price shocks on the US stock market." *International Economic Review*, 50(4): 1267–1287.

Knittel, Christopher R. 2014. "The Political Economy of Gasoline Taxes: Lessons from the Oil Embargo." *Tax Policy and the Economy*, 28(1): 97–131.

Kotz, Maximilian, Anders Levermann, and Leonie Wenz. 2024. "The economic commitment of climate change." *Nature*, 628(8008): 551–557.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, 47(3): 523–544.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review*, 111(4): 1092–1125.

Labandeira, Xavier, José M Labeaga, and Xiral López-Otero. 2017. "A meta-analysis on the price elasticity of energy demand." *Energy policy*, 102: 549–568.

Mangiante, Giacomo. 2024. "The geographic effects of carbon pricing." *European Economic Review*, 168: 104825.

Mansanet-Bataller, Maria and Angel Pardo. 2009. "Impacts of regulatory announcements on CO2 prices." *The Journal of Energy Markets*, 2(2): 1–33.

Marin, Giovanni, Marianna Marino, and Claudia Pellegrin. 2018. "The impact of the European Emission Trading Scheme on multiple measures of economic performance." *Environmental and Resource Economics*, 71(2): 551–582.

Martin, Ralf, Laure B. De Preux, and Ulrich J. Wagner. 2014. "The impact of a carbon tax on manufacturing: Evidence from microdata." *Journal of Public Economics*, 117: 1–14.

Meng, Kyle C. 2017. "Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy." *American Economic Review*, 107(3): 748–84.

Mertens, Karel and Morten O. Ravn. 2013. "The dynamic effects of personal and corporate income tax changes in the United States." *American Economic Review*, 103(4): 1212–47.

Metcalf, Gilbert E. 2019. "On the economics of a carbon tax for the United States." *Brookings Papers on Economic Activity*, 2019(1): 405–484.

Metcalf, Gilbert E. and James H. Stock. 2020. "Measuring the Macroeconomic Impact of Carbon Taxes." *AEA Papers and Proceedings*, 110: 101–06.

Metcalf, Gilbert E. and James H. Stock. 2023. "The Macroeconomic Impact of Europe's Carbon Taxes." *American Economic Journal: Macroeconomics*, 15(3): 265–86.

Miranda-Agrippino, Silvia and Giovanni Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, 13(3): 74–107.

Miranda-Agrippino, Silvia and Giovanni Ricco. 2023. "Identification with external instruments in structural VARs." *Journal of Monetary Economics*, 135: 1–19.

Montiel Olea, José Luis and Carolin Pflueger. 2013. "A robust test for weak instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Montiel Olea, José Luis and Mikkel Plagborg-Møller. 2021. "Local projection inference is simpler and more robust than you think." *Econometrica*, 89(4): 1789–1823.

Nakamura, Emi and Jón Steinsson. 2018*a*. "High-frequency identification of monetary non-neutrality: The information effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Nakamura, Emi and Jón Steinsson. 2018*b*. "Identification in macroeconomics." *Journal of Economic Perspectives*, 32(3): 59–86.

Noh, Eul. 2019. "Impulse-response analysis with proxy variables."

Nordhaus, William D. 1992. "An optimal transition path for controlling greenhouse gases." *Science*, 258(5086): 1315–1319.

Ohlendorf, Nils, Michael Jakob, Jan Christoph Minx, Carsten Schröder, and Jan Christoph

Steckel. 2021. “Distributional impacts of carbon pricing: A meta-analysis.” *Environmental and Resource Economics*, 78(1): 1–42.

Ortubai, Joxe Mari Barrutiabengoa, Rodrigo Enrique Falbo Piacentini, Agustın Garcıa Serrador, and Juan F. Rubio-Ramırez. 2025. “Unraveling the impact of a carbon price shock on macroeconomic variables: a Narrative Sign Restrictions approach.”

Patterson, Christina. 2023. “The matching multiplier and the amplification of recessions.” *American Economic Review*, 113(4): 982–1012.

Plagborg-Møller, Mikkel and Christian K. Wolf. 2021. “Local projections and VARs estimate the same impulse responses.” *Econometrica*, 89(2): 955–980.

Quilis, Enrique M. 2020. “Temporal disaggregation.” https://www.mathworks.com/matlabcentral/fileexchange/69800-temporal-disaggregation, [Online; accessed 06-Dec-2020].

Ramey, Valerie A. 2016. “Macroeconomic shocks and their propagation.” In *Handbook of Macroeconomics*. Vol. 2, 71–162. Elsevier.

Rennert, Kevin, Frank Errickson, Brian C. Prest, Lisa Rennels, Richard G Newell, William Pizer, Cora Kingdon, Jordan Wingenroth, Roger Cooke, Bryan Parthum, et al. 2022. “Comprehensive evidence implies a higher social cost of CO2.” *Nature*, 610(7933): 687–692.

Sims, Christopher A., James H. Stock, and Mark W. Watson. 1990. “Inference in linear time series models with some unit roots.” *Econometrica*, 113–144.

Stefan, Martin and Claudia Wellenreuther. 2020. “London vs. Leipzig: Price discovery of carbon futures during Phase III of the ETS.” *Economics Letters*, 188: 108990.

Stock, James H. 2008. “What’s new in Econometrics: Time Series, lecture 7.” *NBER Summer Institute Short Course Lectures*.

Stock, James H. and Mark W. Watson. 2012. “Disentangling the channels of the 2007-2009 recession.” *Brookings Papers on Economic Activity*.

Stock, James H. and Mark W. Watson. 2018. “Identification and estimation of dynamic causal effects in macroeconomics using external instruments.” *The Economic Journal*, 128(610): 917–948.

Williams, Roberton C., Hal Gordon, Dallas Burtraw, Jared C. Carbone, and Richard D. Morgenstern. 2015. “The initial incidence of a carbon tax across income groups.” *National Tax Journal*, 68(1): 195–213.

Online Appendix

The Unequal Economic Consequences of Carbon Pricing

Diego R. Känzig[1]
Northwestern University

Contents

[1]Contact: Diego R. Känzig, Northwestern University, 2211 Campus Drive, Evanston IL-60202. E-mail: dkaenzig@northwestern.edu. Web: diegokaenzig.com.

A. Data

A.1. Regulatory events

In this Appendix, I provide a detailed list of all the regulatory events used in the paper. To collect the events, I relied on a number of different sources. After 2010, most of the relevant news can be found on the European Commission Climate Action news archive: `https://climate.ec.europa.eu/news-your-voice/news_en`. Before that, I used information from the official journal of the European Union: `https://eur-lex.europa.eu/homepage.html`. Finally, the decisions on the NAPs in the first two phases are taken from Mansanet-Bataller and Pardo (2009). Table A.1 lists all the events.

To mitigate the risk of including events that are potentially confounded by other macroeconomic news, in particular due to news related to demand or the oil market, I performed a detailed narrative analysis of the regulatory event dates. Specifically, I studied the reporting from leading newswires on Factiva, searching for news related to the European carbon market around each of the collected events. Based on this news coverage, I was able to identify a number of events that were likely confounded by other economic news, related to the oil market, the sovereign debt crisis, or Brexit. I mark these events in the fourth column of Table A.1 and provide the exact reason for the exclusion in the last column.

Table A.1: Regulatory Update Events

	Date	Event description	Type	Exclude	Reason
1	25/05/2005	Italian phase I NAP approved	Free alloc.	x	Bullish market data driving oil prices up
2	20/06/2005	Greek phase I NAP approved	Free alloc.		
3	23/11/2005	Court judgement on proposed amendment to NAP, UK vs Commission	Free alloc.		
4	22/12/2005	Further guidance on allocation plans for the 2008–2012 trading period	Cap		
5	22/02/2006	European commission rejects UK NAP amendment	Free alloc.		
6	23/10/2006	Stavros Dimas delivered the signal to tighten the cap of phase II	Cap		
7	13/11/2006	Decision avoiding double counting of emission reductions for projects under the Kyoto Protocol	Intl. credits		
8	29/11/2006	Commission decision on the NAP of several member states	Free alloc.		
9	14/12/2006	Decision determining the respective emission levels of the community and each member state	Cap		
10	16/01/2007	European Commission ruled that Belgium and the Netherlands must make additional emission cuts	Free alloc.		
11	05/02/2007	Slovenia phase II NAP approved	Free alloc.		
12	26/02/2007	Spain phase II NAP approved	Free alloc.		
13	26/03/2007	Phase II NAPs of Poland, France and Czech Republic approved	Free alloc.		
14	02/04/2007	European Commission orders cuts to Austrian phase II NAP	Free alloc.		
15	16/04/2007	European Commission orders cuts to Hungarian phase II NAP	Free alloc.		
16	30/04/2007	Court order on German NAP, EnBW AG vs Commission	Free alloc.		
17	04/05/2007	Estonian phase II NAP approved	Free alloc.		
18	15/05/2007	Italian phase II NAP approved	Free alloc.		
19	07/11/2007	Court judgement on German NAP, Germany vs Commission	Free alloc.		
20	08/04/2008	Court order on German NAP, Saint-Gobain Glass GmbH vs Commission	Free alloc.		
21	23/04/2009	Directive 2009/29/EC amending Directive 2003/87/EC to improve and extend the EU ETS	Cap		

	Date	Event description	Type	Exclude	Reason
22	23/09/2009	Court judgement on NAP, Poland vs Commission	Free alloc.		
23	24/12/2009	Decision determining sectors and subsectors which have a significant risk of carbon leakage	Free alloc.		
24	19/04/2010	Commission accepts Polish NAP for 2008-2012	Free alloc.		
25	09/07/2010	Commission takes first step toward determining cap on emission allowances for 2013	Cap		
26	14/07/2010	Member states back Commission proposed rules for auctioning of allowances	Auction		
27	22/10/2010	Cap on emission allowances for 2013 adopted	Cap		
28	12/11/2010	Commission formally adopted the regulation on auctioning	Auction		
29	25/11/2010	Commission presents a proposal to restrict the use of credits from industrial gas projects	Intl. credits		
30	15/12/2010	Climate Change Committee supported the proposal on how to allocate emissions rights	Free alloc.		
31	21/01/2011	Member states voted to support the ban on the use of certain industrial gas credits	Intl. credits		
32	15/03/2011	Commission proposed that 120 million allowances to be auctioned in 2012	Auction		
33	22/03/2011	Court judgement on NAP, Latvia vs Commission	Free alloc.		
34	29/03/2011	Decision on transitional free allocation of allowances to the power sector	Free alloc.		
35	27/04/2011	Decision 2011/278/EU on transitional Union-wide rules for harmonized free allocation of allowances	Free alloc.		
36	29/04/2011	Commission rejects Estonia's revised NAP for 2008-2012	Free alloc.		
37	07/06/2011	Commission adopts ban on the use of industrial gas credits	Intl. credits		
38	13/07/2011	Member states agree to auction 120 million phase III allowances in 2012	Auction		
39	26/09/2011	Commission sets the rules for allocation of free emissions allowances to airlines	Free alloc.		
40	14/11/2011	Clarification on the use of international credits in the third trading phase	Intl. credits	x	EUA prices fell as recession fears grew
41	23/11/2011	Regulation 1210/2011 determining the volume of allowances to be auctioned prior to 2013	Auction		
42	25/11/2011	Update on preparatory steps for auctioning of phase 3 allowances	Auction	x	EUA prices collapsed due to recession fears
43	05/12/2011	Commission decision on revised Estonian NAP for 2008-2012	Free alloc.		
44	29/03/2012	Court judgments on NAPs for Estonia and Poland	Free alloc.		
45	02/05/2012	Commission publishes guidelines for review of GHG inventories in view of setting national limits for 2013-20	Cap		
46	23/05/2012	Commission clears temporary free allowances for power plants in Cyprus, Estonia and Lithuania	Free alloc.	x	EUA prices fell due to Grexit fears
47	05/06/2012	Commission publishes guidelines on State aid measures in the context of the post-2012 trading scheme	Free alloc.		
48	06/07/2012	Commission clears temporary free allowances for power plants in Bulgaria, Czech Republic and Romania	Free alloc.		
49	13/07/2012	Commission rules on temporary free allowances for power plants in Poland	Free alloc.		
50	25/07/2012	Commission proposed to backload certain allowances from 2013-2015 to the end of phase III	Auction		
51	12/11/2012	Commission submits amendment to back-load 900 million allowances to the years 2019-2020	Auction		
52	14/11/2012	Commission presents options to reform the ETS to address growing supply-demand imbalance	Cap		
53	16/11/2012	Auctions for 2012 aviation allowances put on hold	Auction		
54	30/11/2012	Commission rules on temporary free allowances for power plants in Hungary	Free alloc.		
55	25/01/2013	Update on free allocation of allowances in 2013	Free alloc.		
56	28/02/2013	Free allocation of 2013 aviation allowances postponed	Free alloc.		
57	25/03/2013	Auctions of aviation allowances not to resume before June	Auction		
58	16/04/2013	The European Parliament voted against the Commission's back-loading proposal	Auction		
59	05/06/2013	Commission submits proposal for international credit entitlements for 2013 to 2020	Intl. credits		
60	03/07/2013	The European Parliament voted for the carbon market back-loading proposal	Auction		
61	10/07/2013	Member states approve addition of sectors to the carbon leakage list for 2014	Free alloc.		
62	30/07/2013	Update on industrial free allocation for phase III	Free alloc.		
63	05/09/2013	Commission finalized decision on industrial free allocation for phase three	Free alloc.		
64	26/09/2013	Update on number of aviation allowances to be auctioned in 2012	Auction		
65	08/11/2013	Member states endorsed negotiations on the back-loading proposal	Auction		
66	21/11/2013	Commission submitted non-paper on back-loading to the EU Climate Change Committee	Auction		
67	10/12/2013	European Parliament voted for the back-loading proposal	Auction		
68	11/12/2013	Climate Change Committee makes progress on implementation of the back-loading propsal	Auction		
69	18/12/2013	Commission gives green light for a first set of member states to allocate allowances for calendar year 2013	Free alloc.		

	Date	Event description	Type	Exclude	Reason
70	08/01/2014	Climate Change Committee agrees back-loading	Auction		
71	22/01/2014	Commission proposed to establish a market stability reserve for phase V	Cap		
72	26/02/2014	Commission gives green light for free allocation by all member states	Free alloc.		
73	27/02/2014	Back-loading: 2014 auction volume reduced by 400 million allowances	Auction		
74	13/03/2014	Commission approves first batch of international credit entitlement tables	Intl. credits		
75	28/03/2014	Commission approves second batch of international credit entitlement tables	Intl. credits	x	EUA prices fell ahead of bearish emissions data
76	04/04/2014	Update on approval of international credit entitlement tables	Intl. credits		
77	11/04/2014	Commission approves four more international credit entitlement tables	Intl. credits		
78	23/04/2014	Commission approves final international credit entitlement tables	Intl. credits		
79	02/05/2014	Commission published the number of international credits exchanged	Intl. credits		
80	05/05/2014	Commission submits proposed carbon leakage list for 2015-2019	Free alloc.		
81	04/06/2014	Auctioning of aviation allowances to restart in September	Auction		
82	04/07/2014	Commission published the first update on the allocation of allowances from the New Entrants' Reserve	Free alloc.		
83	09/07/2014	Climate Change Committee agrees proposed carbon leakage list for the period 2015-2019	Free alloc.		
84	27/10/2014	Commission adopts the carbon leakage list for the period 2015-2019	Free alloc.		
85	04/11/2014	Updated information on exchange and international credit use	Intl. credits	x	EUA prices fell due to weak economic outlook
86	04/05/2015	Updated information on exchange and international credit use	Intl. credits		
87	15/07/2015	Proposal to revise the EU emissions trading system for the period after 2020	Cap		
88	23/07/2015	Commission publishes status update for New Entrants' Reserve and allocation reductions	Free alloc.		
89	04/11/2015	Updated information on exchange and international credit use	Intl. credits		
90	15/01/2016	Commission publishes status update for New Entrants' Reserve	Free alloc.	x	EUA prices plunge amid weakening demand
91	28/04/2016	Court judgment on free allocation in the EU ETS for the period 2013-2020	Free alloc.		
92	02/05/2016	Updated information on exchange and international credit use	Intl. credits		
93	23/06/2016	Following court judgement, commission to modify cross-sectoral correction factor for 2018-2020	Free alloc.	x	Brexit vote
94	15/07/2016	Commission published a status update on the allocation of allowances from the New Entrants' Reserve 2013-2020	Free alloc.		
95	08/09/2016	Court judgment on free allocation in the EU ETS for the period 2013-2020	Free alloc.		
96	04/11/2016	Updated information on exchange and international credit use	Intl. credits		
97	16/01/2017	Commission publishes status update for New Entrants' Reserve	Free alloc.		
98	24/01/2017	Commission adopts Decision to implement Court ruling on the cross-sectoral correction factor	Free alloc.		
99	15/02/2017	European Parliament voted in support of the revision of the ETS Directive for the period after 2021	Cap		
100	27/04/2017	Climate Change Committee approves technical changes to auction rules	Auction		
101	02/05/2017	Updated information on exchange and international credit use	Intl. credits		
102	12/05/2017	Commission publishes first surplus indicator for ETS Market Stability Reserve	Auction		
103	17/07/2017	Commission publishes status update for New Entrants' Reserve	Free alloc.		
104	26/07/2017	Court judgment again confirms benchmarks for free allocation of ETS allowances for 2013-2020	Free alloc.		
105	06/11/2017	Updated information on exchange and international credit use	Intl. credits		
106	15/01/2018	Commission publishes status update for New Entrants' Reserve	Free alloc.		
107	04/05/2018	Updated information on exchange and international credit use	Intl. credits		
108	08/05/2018	Commission Notice on the preliminary carbon leakage list for phase IV (2021-2030)	Free alloc.		
109	15/05/2018	ETS Market Stability Reserve will start by reducing auction volume by almost 265 million allowances	Auction		
110	16/07/2018	Commission publishes status update for New Entrants' Reserve	Free alloc.		
111	30/10/2018	Commission adopts amendment to ETS auctioning regulation	Auction		
112	06/11/2018	Updated information on exchange and international credit use	Intl. credits		
113	05/12/2018	Poland's 2019 auctions to include some allowances not used for power sector modernisation	Auction	x	Uncertainty in oil markets due to OPEC meeting
114	04/01/2019	Amendment to the ETS auctioning regulation	Auction		
115	15/01/2019	Commission publishes status update for New Entrants' Reserve	Free alloc.		
116	15/02/2019	Adoption of the Delegated Decision on the carbon leakage list for 2021-2030	Free alloc.	x	EUA prices fell due to bearish economic data and worries ahead of Brexit deadline
117	23/04/2019	Iceland, Liechtenstein and Norway to start auctions on the common auction platform soon	Auction	x	Oil price jump after US Iranian sanctions
118	15/05/2019	ETS Market Stability Reserve to reduce auction volume by almost 400 million allowances	Auction		
119	16/05/2019	Revised 2019 auction calendars including EEA EFTA volumes published	Auction		

	Date	Event description	Type	Exclude	Reason
120	12/06/2019	Poland's 2020 auction volume to include allowances not used for power sector modernisation	Auction	x	Oil prices fell due to weak economic outlook and trade policy uncertainty
121	19/06/2019	Updated information on exchange and international credit use	Intl. credits		
122	11/07/2019	2020 and revised 2019 auction calendars of the common auction platform published	Auction		
123	15/07/2019	Commission publishes status update for New Entrants' Reserve	Free alloc.		
124	28/08/2019	Commission amends ETS auctioning regulation for phase 4	Auction		
125	31/10/2019	Commission adopts the Regulation on adjustments to free allocation due to activity level changes	Free alloc.		
126	08/11/2019	Auctioning regulation amendment for phase 4 of the EU ETS published and to enter into force	Auction		

A.2. Macro data

In this Appendix, I provide details on the financial and macroeconomic data used in the paper. Table A.2 provides information on the source of the data, its construction and coverage.

Table A.2: Description, Sources, and Coverage of Macro Data

Variable	Description	Source	Sample
Instrument			
LEXC.01 (PS)	EUA futures front contract (settlement price)	Datastream	22/04/2005-31/12/2019
ELECWAVG	Wholesale electricity price, constructed as weighted average of EEX, APX, Nordpool, Powernext, OMEL, GME, and the EPX spot price, converted to EUR/tCO2 using GHG emissions intensity of electricity generation	Datastream/European Environment Agency/own calculations	22/04/2005-31/12/2019
Baseline variables			
EKESCPENF	HICP energy (EA-19)	Datastream	1999M1-2019M12
GHGTOTAL	Total GHG emissions excluding LULUCF and including international aviation (EU)	Eurostat/own calculations	1999M1-2019M12
EKCPHARMF	HICP all items (EA-19)	Datastream	1999M1-2019M12
EKIPTOT.G	Industrial production excl. construction (EA-19)	Datastream	1999M1-2019M12
EMECB2Y.	Two-year government bond yield	Datastream	1999M1-2019M12
EKESUNEMO	Unemployment rate (EA-19)	Datastream	1999M1-2019M12
DJSTOXX	Euro STOXX	Datastream	1999M1-2019M12
DCOILBRENTEU	Brent Crude price	FRED	1999M1-2019M12
Additional variables			
Other carbon futures	LEXC.0h (PS), for h in $(2,3,4,5)$	Datastream	22/04/2005-31/12/2019
BAMLHE00EHYIOAS	ICE BofA euro high yield index option-adj. spread	FRED	1999M1-2019M12
VSTOXX	Euro STOXX 50 volatility	stoxx.com	1999M1-2019M12
EKGDP...D	Real GDP (EA-19)	Datastream	1999Q1-2019Q4
EKESENMZD	Final consumption expenditure (EA-19)	Datastream	1999Q1-2019Q4
EKGFCF..D	Gross fixed capital formation (EA-19)	Datastream	1999Q1-2019Q4
EMESJSABB	Wages and salaries: all activities	Datastream	1999Q1-2019Q4

As discussed in the main text, emissions data from Eurostat is unfortunately not available at the monthly frequency. For a large part of our sample, data on emissions are only available at the annual frequency. Only starting from 2010, Eurostat also started to publish quarterly figures.

As GHG emissions are a key outcome of interest, we would like to include a measure of emissions in our baseline model. To this end, I construct a monthly emissions series by temporally disaggregating the annual emissions data using a selection of relevant monthly indicators, following the method by Chow and Lin (1971).[2] As the relevant monthly indicators I use petroleum and other fuels consumption, industrial production and harmonized consumer prices (energy and headline). This choice is motivated by a validation exercise based on the quarterly Eurostat emission figures. However, the results are relatively robust to the exact indicators included.[3]

Figure A.1 shows the resulting interpolated monthly GHG emissions series together with a naive series that is obtained by uniformly distributing the annual total. The interpolated series looks meaningful and behaves as we would expect. A nice feature is also that by construction, the interpolated monthly series will aggregate to the annual total.

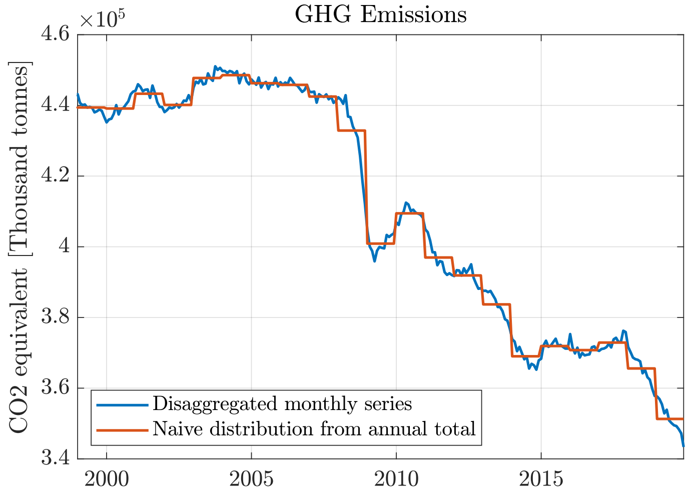

Figure A.1: The Monthly GHG Emissions Series

Notes: This figure shows the disaggregated monthly GHG emissions series, constructed with the Chow and Lin (1971) approach using petroleum and other fuels consumption, industrial production and harmonized consumer prices (energy and headline) as the relevant monthly indicators, together with a naive series that is obtained by uniformly distributing the annual total.

The transformed series used in the baseline VAR are depicted in Figure A.2. We can see a big spike in the two-year rate associated with the sovereign debt crisis. As this may potentially confound the two-year rate as a measure of the monetary stance, I account for this spike using a dummy variable. The results are robust to using a shadow rate instead.

[2] I implement the approach using the Quilis's (2020) Matlab code suite.

[3] I make the disaggregated monthly emissions data available on the following Github repository: https://github.com/dkaenzig/monthlyGHGemissions.

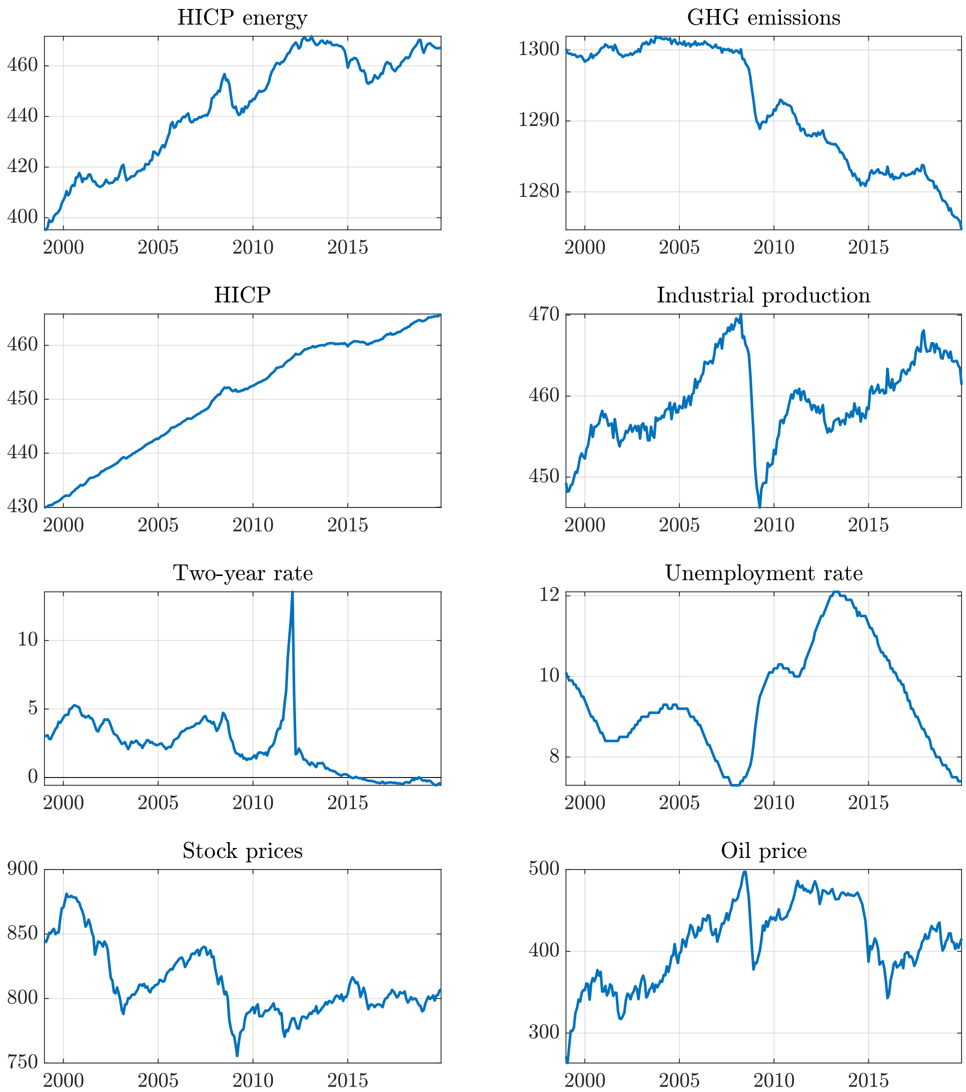

Figure A.2: Transformed Data Series

A.3. Patent data

For the patent data, I rely on the World Patent Statistical Database (PATSTAT), which encompasses bibliographic information for close to the universe of patents globally. I use the autumn 2023 edition (version 5.22) with data for 2005-2019.

I follow the literature (e.g. Hémous et al., 2025) and focus on patent families, i.e. patents representing the same innovation filed at different patent offices. For each patent family I use the original application date to capture the time of the innovation and assign nationality based on the respective filing office. I restrict the analysis to biadic patents, i.e. patents that are filed in at least two jurisdictions (e.g. at the USPTO and the EPO), to screen out low-quality patents.

To measure green innovation, I rely on the International Patent Classification (IPC) and Cooperative Patent Classification (CPC) codes. Specifically, I use the

OECD (Haščič and Migotto, 2015) subclasses Y02 and Y04S of the C/IPC, which include climate change mitigation and adaptation technologies. I classify patent families with multiple C/IPC codes as green if any of the respective codes falls into the relevant subclasses.

To study which green patents are responding the most, I consider three subgroups of Y02/Y04S: (i) mitigation in energy generation, transportation, and buildings (Y02E, Y02T, Y02B, Y02W, Y04S), (ii) mitigation in industry/manufacturing (Y02P, Y02D), (iii) adaptation or carbon capture technologies (Y02A, Y02C). Table A.3 shows that the majority of green patent filings in my sample is associated with energy generation, transportation, and buildings.

Table A.3: Descriptive Statistics on Green Patenting

Patent type	Share
Overall green	11.17
Green narrow definition	9.78
Energy, transportation, buildings	87.02
Industry and enabling technologies	22.38
Adaptation and carbon removal	10.63

Notes: Descriptive statistics on green patenting, including the share of overall green patents relative to all patents filed, the share of green patents according to the narrower definition of Acemoglu et al. (2023), as well as the sub shares within green patents (energy, transportation and buildings; industry and enabling technologies; adaptation and carbon removal—all expressed relative to overall green patents filed). Note that these shares do not sum to 100 because patents can belong to multiple groups.

A.4. Panel data

For the analysis in Section 5.1, I construct a quarterly panel of 15 advanced European countries that were part of the EU ETS, spanning the period from 1999Q1 to 2019Q4. Specifically, the panel includes Austria, Belgium, Denmark, Finland, France, Germany, Greece, Italy, Ireland, the Netherlands, Norway, Portugal, Spain, Sweden, and the United Kingdom. The choice of countries is guided primarily by data availability and the desire to focus on a more homogeneous set of advanced economies. Real GDP data is from Eurostat, while country characteristics are drawn from Almgren et al. (2022) and the EU Transaction Log. Table A.4 provides more information on the variables used as well as their source.

Table A.4: Panel Data Description

Variable	Description	Source
CLV10_MEUR	Real Gross Domestic Product (Millions of Chained 2010 Euros, Seasonally Adjusted)	Eurostat
rgdppc_na	Real GDP per capita (sample average)	PWT
emission_intensity	GHG emissions intensity (total GHG emissions over GDP, sample average)	Känzig and Konradt (2024)
h2m	Share of hand-to-mouth households	Almgren et al. (2022)
Employment_laws_index	Employment protection index by Botero et al. (2004)	Almgren et al. (2022)

A.5. Household data

In this Appendix, I provide detailed information on the micro data used in Section 5.2 of the paper. I use data from a selection of different surveys, which are discussed in detail below.

A.5.1. LCFS

The living costs and food survey (LCFS) data can be obtained from the UK Data Service. I use the waves from 1999-2001 of the Family Expenditure Survey, the 2001-2007 waves from the Expenditure and Food Survey and the 2008-2019 waves from the LCFS, which superseded the previous two surveys. Note that within this sample, the reporting frequency changed two times first from financial year to calendar year and then back again to the financial year format. The waves are adjusted to consistently reflect the calendar year prior to creating the pooled cross-section. Most variables of interest are available in the derived household datasets. The age at which full-time education was completed, as well as current wages, is aggregated from the personal derived datasets.

As the main measure of expenditure, I use total expenditure excluding housing (p550tp-p536tp). For current income, I use current total disposable income, calculated by subtracting income taxes and NI contributions from the gross income (p352p-p392p-p388p-p029hp). I group the households by their normal disposable income (p389p). For earnings, I use wages net of taxes (aggregate p004p to the household level, subtract current taxes and add back taxes on financial income p068h). For financial income, I use p324p, which includes interest income, dividends and rents. For age, I use the age of the household reference person, p396p. Education is proxied by the highest age a person in the household has completed a full-time education (a010 aggregated to the household level). The housing tenure status is recorded in variable a121.

For energy expenditure, I use expenditure on fuel, light and power (p537t). Constructing measures of non-durable, services and durable expenditure is not trivial in the LCFS data, as the broader available expenditure categories do not allow a clean split, e.g. personal goods and services (p544t) is a mix of non-durable goods and services while household goods (p542t) includes both non-durable and durable goods. To construct clean measures of non-durables, services and durables expenditure, I split these broader subcategories into non-durable, services and durable parts by grouping the items in a particular subcategory accordingly, following closely the COICOP guidelines. A further challenge in doing so is that the code names for disaggregated expenditure items changed when the FES became the EFS in 2001. In Table A.5, I detail how the non-durable, services and durable expenditure measures are constructed. At the item level, I provide both, the relevant codes in the FES and the EFS/LCFS. Note that semi-durables are subsumed under non-durables, and services do not include housing.

Table A.5: Expenditure Classification in LCFS

Category	Subcategories	Items
Non-durables	Fuel, light power (p537t)	
	Food, alcoholic drinks, tobacco (p538t, p539t, p540t)	
	Clothing and footwear (p541t)	
	Non-durable household goods (subset of p542t)	*LCFS codes:* c52111t, c52112t, c53311t, c55214t, c56111t, c56112t, c56121t, c56123t, c93114t, c93313t, c93411t, c95311t, c95411t, cc1311t *FES codes:* d070104t, d070105t, d070211t, d070209t, d070401t, d070402t, d070302t, d070601t, d120304t, d070501t
	Non-durable personal goods (subset of p544t)	*LCFS codes:* c61112t, c61211t, c61311t, c61313t, cc1312t, cc1313t, cc1314t, cc1315t, cc1316t, cc1317t, cc3211t, cc3222t, cc3223t, cc3224t *FES codes:* d090402t, d090102t, d090501t, d090101t, d090103t, d090104t, d090105t, d090301t, d090202t, d090302t, d090303t
	Non-durable motoring expenditure (subset of p545t)	*LCFS codes:* c72114t, c72211t, c72212t, c72213t *FES codes:* d100405t, d100301t, d100302t, d100303t
	Non-durable leisure goods (subset of p547t)	*LCFS codes:* c91126t, c91411t, c91412t, c91413t, c91414t, c93111t, c93113t, c93311t, c95111t, c95211t, c95212t *FES codes:* d120114t, d120108t, d120110t, d120109t, d120401t, d120113t, d070703t, d120303t, d120301t, d120302t
	Miscellaneous non-durable goods (subset of p549t)	*LCFS codes:* ck5511c, cc3221t *FES codes:* d070801t, d140601c, d090701t
Services	Household services (p543t)	
	Fares and other travel costs (p546t)	
	Leisure services (p548t)	
	Service part of household goods (subset of p542t)	*LCFS codes:* c53312t, c53313t, c53314t, c93511t, cc5213t *FES codes:* d070212t, d070213t
	Personal services (subset of p544t)	*LCFS codes:* c61111t, c61312t, c62111t, c62112t, c62113t, c62114t, c62211t, c62212t, c62311t, c62321t, c62322t, c62331t, c63111t, cc1111t *FES codes:* d090401t, d090502t, d090403t, d090404t, d090601t

Category	Subcategories	Items
	Service part of motoring expenditure (subset of p545t)	*LCFS codes:* b187-b179, b188, b249, b250, b252, c72313t, c72314t, c72411t, c72412t, c72413t, ck3112t, c72311c, c72312c, cc5411c *FES codes:* b187-b179, b188, b249, b250, b252, d100403t, d100406t, d100407t, d100404t, d100408t, d100201c, d100204c, d100401c
	Leisure services (subset of p547t)	*LCFS codes:* c91511t, c93112t, c94238t, c94239t, c94246t *FES codes:* d120111t, d120112t
	Miscellaneous services (subset of p549t)	*LCFS codes:* b237, b238, ck5315c, ck5213t, ck5214t *FES codes:* b237, b238, d140402, d140406c
Durables	Durable household goods (subset of p542t)	*LCFS codes:* b270, b271, c51111c, c51211c, c51212t, c51113t, c51114t, c53111t, c53121t, c53122t, c53131t, c53132t, c53133t, c53141t, c53151t, c53161t, c53171t, c53211t, c54111t, c54121t, c54131t, c54132t, c55111t, c55112t, c55213t, c56122t, c93212t, c93312t, c93412t, cc1211t *FES codes:* b270, b271, d070101c, d070102c, d070103t, d070304t, d070704t, d070203t, d070202t, d070204t, d070207t, d070208t, d070201t, d070206t, d070303t, d070301t, d070205t, d070701t, d070305t, d070306t, d070702t, d070602t
	Durable personal goods (subset of p544t)	*LCFS codes:* cc3111t *FES codes:* d090201t
	Durable motoring expenditure (subset of p544t)	*LCFS codes:* b244, b2441, b245, b2451, b247, c31315t, c71112t, c71122t, c71212t, c92114t, c92116t, c71111c, c71121c, c71211c, c92113c, c92115c, c72111t, c72112t, c72113t, c91112t *FES codes:* b244, b245, b247, d100105t, d100106t, d100107t, d100101c, d100102c, d100104c, d100203t, d100202t, d100205t
	Durable leisure goods (subset of p547t)	*LCFS codes:* c91124t, c82111t, c82112t, c82113t, c91111t, c91113t, c91121t, c91122t, c91123t, c91125t, c91211t, c91311t, c92211t, c92221t, c93211t *FES codes:* d120104t, d080202t, d080205t, d080207t, d120105t, d120101t, d120102t, d120103t, d120115t, d120402t, d120106t, d120107t, d120201t

Regarding the sample, I apply the following restrictions. I drop households that have a household reference person younger than 18 or older than 90 years. Furthermore, I drop households with a negative normal disposable income. To account for some (unrealistically) high or low values of consumption, for each quarter and income group, I drop the top and bottom 1% of observations for total expenditure.

A.5.2. LFS

To get information on the sector of employment, I use data from the UK Labour Force Survey (LFS). The LFS studies the employment circumstances of the UK population. It is the largest household study in the UK and provides the official measures of employment and unemployment. Apart from detailed information on employment, it also contains a wide range of related topics such as occupation, training, hours of work and personal characteristics of household members aged 16 years and over. The data can be obtained from the UK Data Service. I use the quarterly waves from 1999-2019 to construct a pooled cross-section. For the employment sector, I use the variable indsect, which describes the industry sector in the main job based on the SIC 2003 classification. To proxy income, I use the net pay from the main and second job (netwk and netwk2).

A.5.3. BSA

To proxy public attitudes towards climate policy, I use data from the British social attitudes (BSA) survey. The data can also be obtained from the UK Data Service. I use the waves from 1999-2019 to construct a pooled cross-section. To construct the income groups, I use the income quartiles that are provided from 2010 onwards (hhincq). For the years before, I use the household income variable (hhincome) to construct the quartiles. The survey contains many questions on the attitudes towards climate change, the environment and climate/environmental policy, but unfortunately most variables are not part of the main set of questions that are asked in every year. One exception concerns a question about taxes for car owners (cartaxhi), in particular it asks whether you agree with the following statement: “For the sake of the environment, car users should pay higher taxes”, which was fielded for all years up to 2017. Thus, I use the proportion of households agreeing with this statement as a proxy for the public attitude towards climate policy.

B. Instrument Diagnostics

In this Appendix, I examine several properties of the carbon policy surprise series, related to instrument validity and strength.

B.1. Validity checks

As discussed in the paper, I perform a number of additional validity checks on the surprise series. In particular, I investigate the autocorrelation and forecastability of the surprise series as well as the relation to other shocks from the literature.

Figure B.1 displays the monthly carbon policy surprise series, aggregated by summing over the surprises in a given month.

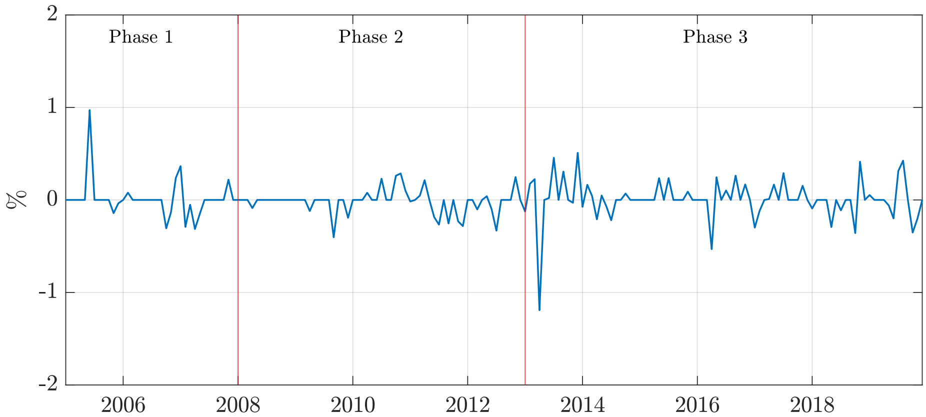

Figure B.1: Monthly Carbon Policy Surprise Series

Notes: The figure shows the refined carbon policy surprise series, aggregated to the monthly frequency by summing over the surprises in a given month.

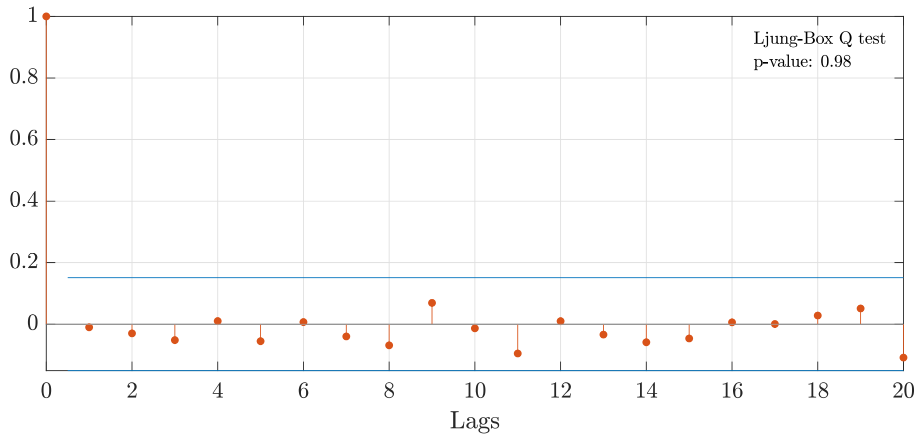

Figure B.2: The Autocorrelation Function of the Carbon Policy Surprise Series

The series looks well behaved, with approximately as many positive as negative surprises. The series is also not serially correlated, as can be seen from the autocorrelation function in Figure B.2.

I also perform a series of Granger causality tests. Table B.1 shows that the monthly series is not forecastable by past macroeconomic variables.

Table B.1: Granger Causality Tests

Variable	p-value
Instrument	0.9724
EUA price	0.7377
HICP energy	0.2063
HICP	0.7327
Industrial production	0.7199
Two-year rate	0.9568
Unemployment rate	0.6601
Stock prices	0.7170
REER	0.3049
Oil price	0.9825
Joint	0.9795

Notes: p-values of a series of Granger causality tests of the monthly carbon policy surprise series using a selection of macroeconomic and financial variables.

Table B.2: Correlation with Other Shock Measures

Shock	Source	ρ	p-value	n	Sample
Monthly measures					
Global oil market					
	Kilian (2009) (updated)	-0.06	0.46	164	2005M05-2018M12
	Caldara, Cavallo, and Iacoviello (2019)	-0.01	0.89	128	2005M05-2015M12
	Baumeister and Hamilton (2019)	-0.09	0.24	176	2005M05-2019M12
	Känzig (2021) (updated)	0.04	0.62	176	2005M05-2019M12
Global demand	Kilian (2009) (updated)	0.04	0.61	164	2005M05-2018M12
	Baumeister and Hamilton (2019)	-0.02	0.78	176	2005M05-2019M12
Oil-specific demand	Kilian (2009) (updated)	0.02	0.76	164	2005M05-2018M12
Consumption demand	Baumeister and Hamilton (2019)	0.05	0.49	176	2005M05-2019M12
Inventory demand	Baumeister and Hamilton (2019)	-0.05	0.53	176	2005M05-2019M12
Monetary policy					
Monetary policy shock	Jarociński and Karadi (2020)	0.05	0.52	140	2005M05-2016M12
Central bank info	Jarociński and Karadi (2020)	0.06	0.51	140	2005M05-2016M12
Financial & uncertainty					
Financial conditions	BBB spread residual	-0.02	0.75	176	2005M05-2019M12
Financial uncertainty	VIX residual (Bloom, 2009)	-0.00	0.97	176	2005M05-2019M12
	VSTOXX residual	0.00	0.97	176	2005M05-2019M12
Policy uncertainty	Global EPU (Baker, Bloom, and Davis, 2016)	0.03	0.72	176	2005M05-2019M12
Quarterly measures					
Fiscal policy	Euro area (Alloza, Burriel, and Pérez, 2019)	0.09	0.58	43	2005Q2-2015Q4
	Germany	0.22	0.15	43	2005Q2-2015Q4
	France	0.03	0.86	43	2005Q2-2015Q4
	Italy	0.01	0.97	43	2005Q2-2015Q4
	Spain	-0.02	0.90	43	2005Q2-2015Q4

Notes: Correlation coefficients of the carbon policy surprise series with a wide range of different shock measures from the literature, including global oil market shocks, monetary policy, financial and uncertainty shocks. ρ is the Pearson correlation coefficient, the p-value corresponds to the test whether the correlation is different from zero and n is the sample size.

Finally, I study how the series correlates with other shock series from the literature and find that it is not correlated with other structural shock measures, including oil demand, uncertainty, financial, fiscal and monetary policy shocks, see Table B.2.[4]

B.2. First stage

Having provided suggestive evidence on the validity of the carbon policy surprise series, I now evaluate instrument strength. Table B.3 presents the results of the weak instruments test for a selection of instruments. I include the raw carbon policy surprise series, normalized by electricity prices or expressed as percentage changes in carbon prices, along with refined surprises purged using different information sets. The results indicate that carbon policy surprises serve as strong instruments, with heteroskedasticity-robust F-statistics typically exceeding the conventional threshold of 10.

Table B.3: Tests on instrument strength

Instrument	Raw		Refined baseline			
	Baseline	$\Delta\log(F)$	(a) Macro	(b) Financials	(c) Oil	(d) Climate
Coefficient	0.610	0.026	0.589	0.590	0.667	0.891
Robust F-stat	12.54	12.85	9.62	9.53	11.69	16.85
R^2	1.40	1.29	1.23	1.22	1.46	2.38
Adj. R^2	0.99	0.89	0.82	0.82	1.05	1.98

Notes: Results of first-stage regressions of the energy price residual $\hat{u}_{1,t}$ on a selection of different instruments: the raw baseline surprise series based on (1), the variant based on the percentage change in the carbon price as well as refined baseline surprise using different information sets: (a) controls for macro news, (b) adding financials, (c) adding oil market variables, (d) adding climatic variables, based on (2). Reported are coefficient estimates, robust F-statistics allowing for heteroskedasticity, the R^2 and adjusted R^2. The number of observations is 246.

[4] I thank Mario Alloza for kindly sharing their fiscal policy shock series.

C. Sensitivity Analysis

In this Appendix, I perform a number of robustness checks on the identification strategy and the empirical specification used to isolate the carbon policy shock, as discussed in Sections 2-3 of the paper. Throughout, I report the point estimate as the solid black line and 68 and 90 percent confidence bands as dark and light shaded areas, respectively.

C.1. Instrument construction

Selection of relevant events. A crucial choice in the high-frequency event study approach concerns the selection of relevant events. For the exclusion restriction to be satisfied, the events should only release information about the supply of emission allowances and not about other factors such as macroeconomic or geopolitical news. To this end, I have not included more high level events such as the Paris agreement or other COP meetings but limited the analysis to specific events in the European carbon market.

As discussed in Section 2.2, I also performed a narrative analysis to identify events that are potentially confounded. Figure C.1 shows how excluding these events affects the results. Accounting for these potentially confounded events makes some difference, particularly for the oil price and the emissions responses, even though the differences are not statistically significant. To mitigate concerns about potentially confounded events, I exclude the flagged events in the baseline analysis.

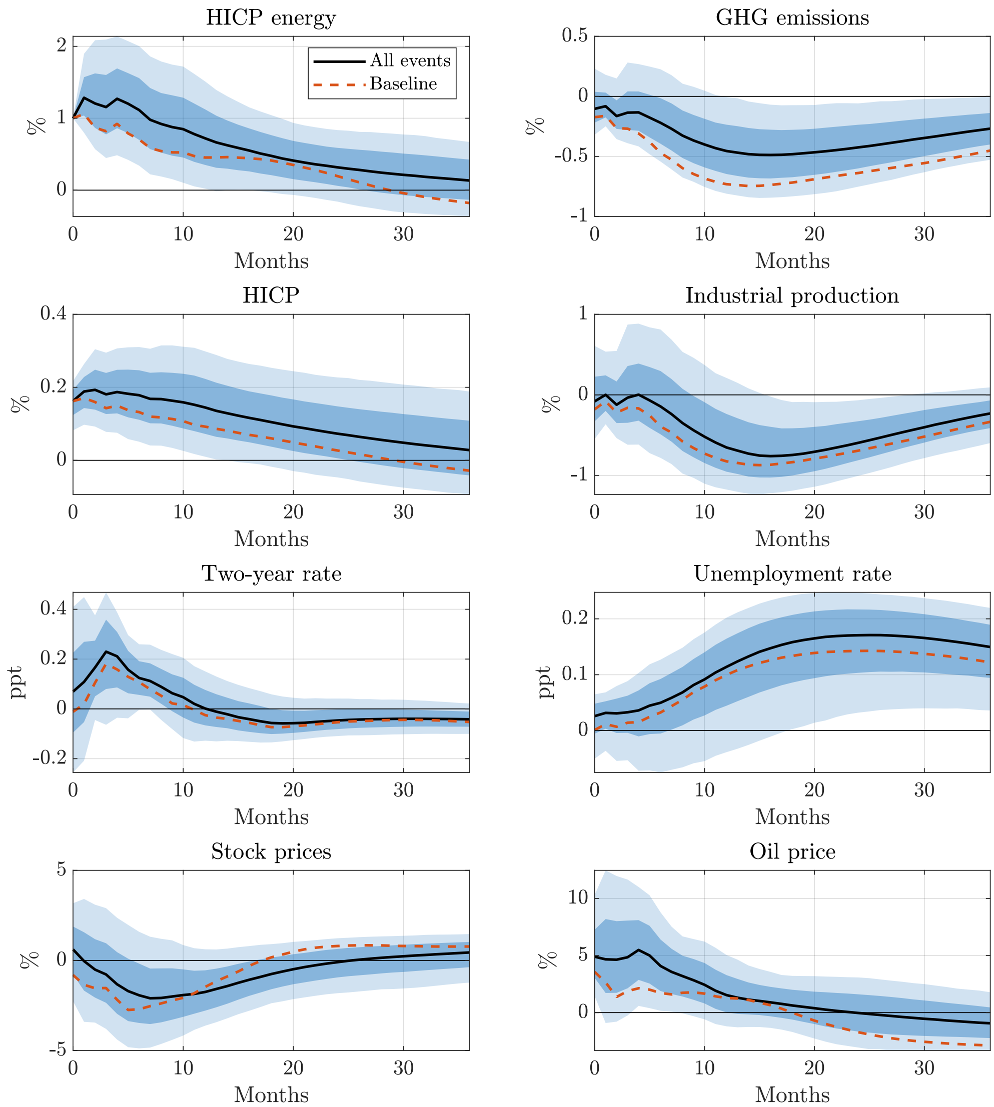

First stage regression: F-statistic: 22.36, R^2: 3.38%

Figure C.1: Excluding Potentially Confounded Events

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

In a series of additional sensitivity checks, I show that the results are not driven by a particular subset of events. To ensure that the results are not driven by surprises in times of economic distress, I exclude (i) events during the Great Recession, (ii) events during the European sovereign debt crisis, and (iii) the 2014-16 oil shock. As an additional check, I also exclude events from the the first, trial phase of the EU ETS, where markets were not as liquid and regulations not as strict.

The results are shown in Figure C.2. The results are very robust to the selection of events. The responses based on these alternative instruments all lie within the 68 percent confidence bands of the baseline model. This suggests that the results are not driven by some specific influential periods such as the sovereign debt crisis or the 2014 oil shock.

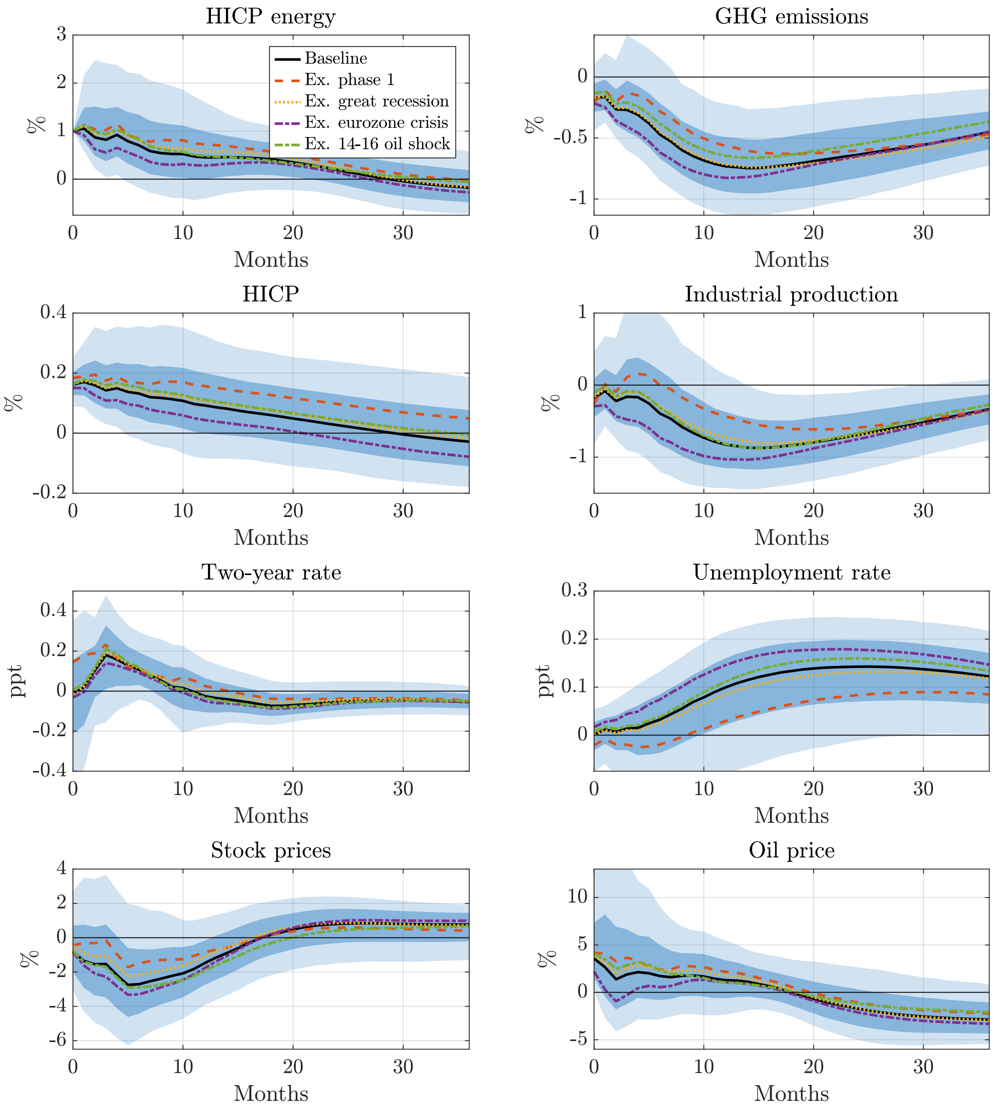

Figure C.2: Excluding Certain Periods from Instrument Construction

To analyze the role of influential events more systematically, I perform a jackknife exercise. Specifically, I censor one value of the carbon policy surprise series at a time to zero, and re-run the external instruments VAR. Figure C.3 shows our baseline response in black, together with the responses from the jackknife exercise in gray. The estimated responses are not driven any single extreme carbon policy surprise. When censoring certain surprises the impacts can get even bigger, while when dropping others the effects can be somewhat attenuated. In all cases, the responses always lay squarely within the 68 percent confidence bands of the baseline model.

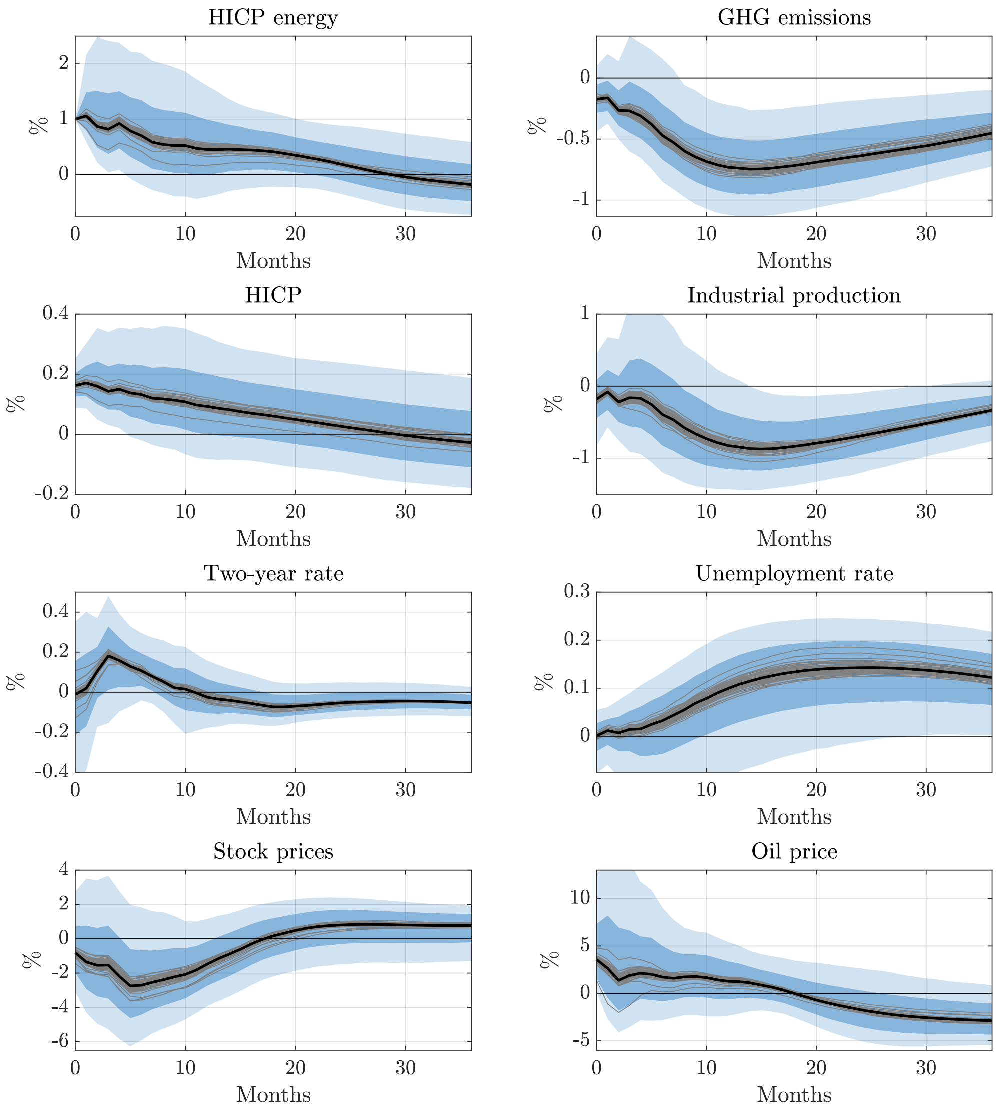

Figure C.3: Results from Jackknife Exercise

Finally, I also perform some robustness with respect to the event type. Specifically, I (i) exclude events concerning the cap as they tend to be broader in nature, (ii) exclude events about international credits, which affect the supply of allowances more indirectly, and (iii) exclude events about data updates which affect the supply of allowances indirectly, e.g. by triggering the market stability reserve. Figure C.4 presents the results. Reassuringly, the results turn out to be robust across the different type of events, and excluding the events described above does not change the responses meaningfully compared to the baseline model.

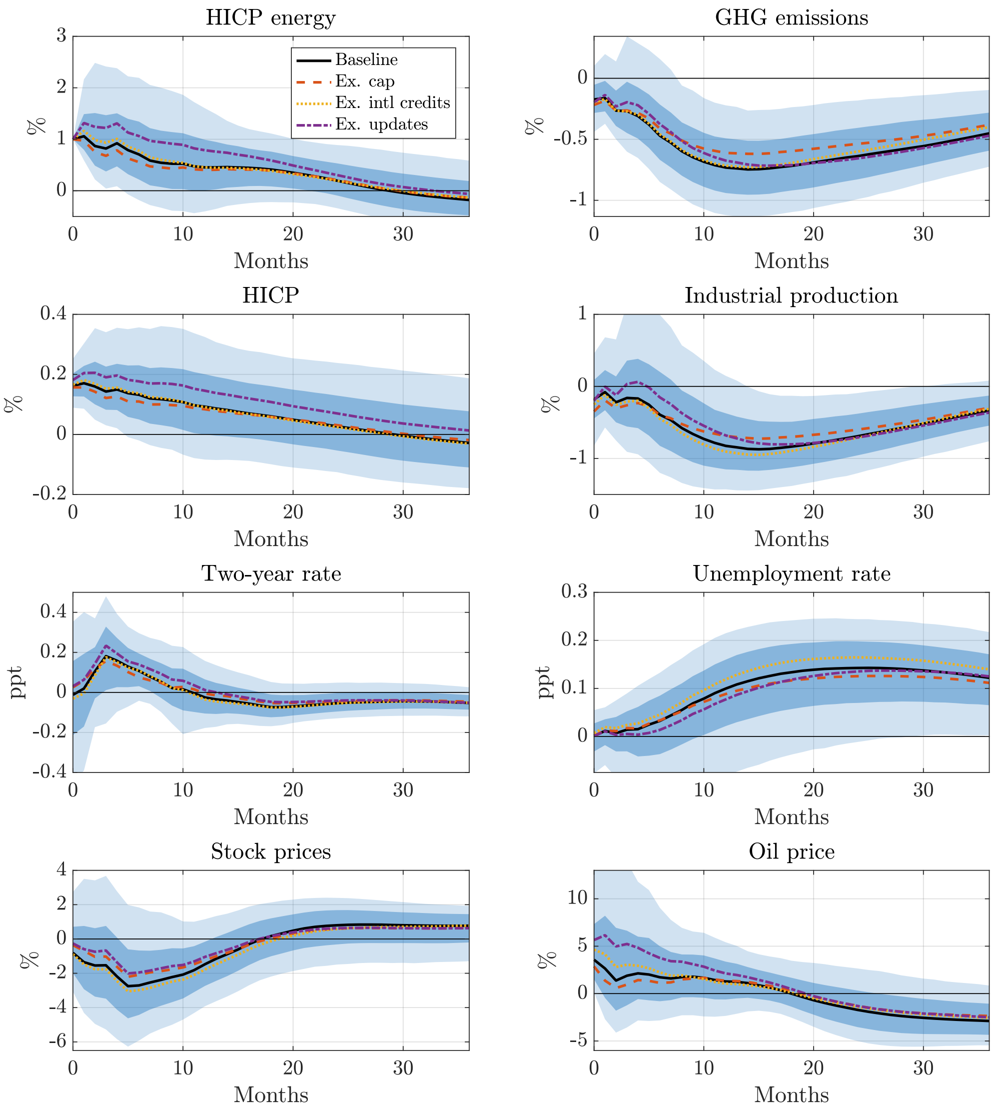

Figure C.4: Sensitivity With Respect to Event Types

Removing predictability. In Section 2.2, we have seen that there is some potential predictability in the carbon policy surprise series, driven by oil market specific and climate related variables. Figure C.5 shows how accounting for this predictability affects the results.

Accounting for the potential predictability makes some difference, particularly for the oil price response, even though the differences are not statistically significant. To mitigate any identification concerns related to the predictability of the surprise series, I use the refined carbon policy surprise series as the baseline.

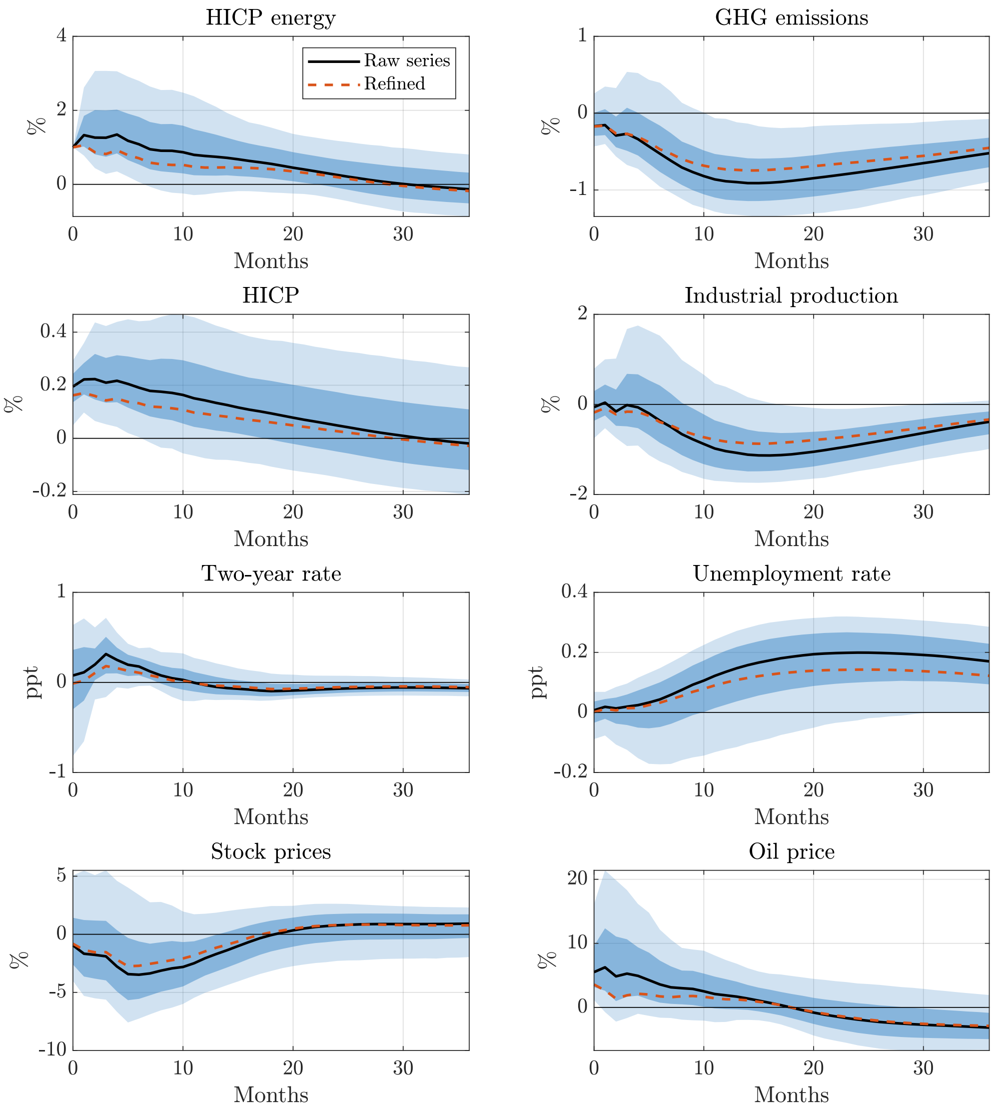

First stage regression: F-statistic: 12.54, R^2: 1.40%

Figure C.5: The Role of Removing Predictability

Alternative instrument. As discussed in Section 2.2, I measure the carbon policy surprises as the change in the EUA futures price on the day of the regulatory event relative to the prevailing wholesale electricity price on the day before the event.[5]

The main advantage of this approach is that it directly gives a notion of how economically relevant a carbon policy surprise is. In particular, it gives less weight to large percentage changes in carbon prices that occurred in an environment where carbon prices were very low. An alternative approach is to simply measure the surprise as the percentage change in the carbon price on event days.

[5]To mitigate the influence of extreme observations in the wholesale electricity price, I use an average of the price over the last 5 trading days before the event.

To fix ideas, the carbon policy surprise is in this case computed as follows:

$$CPSurprise_{t,d} = \log(F^{carbon}_{t,d}) - \log(F^{carbon}_{t,d-1})$$

where d and t indicate the day and the month of the event, respectively and $F_{t,d}$ is the settlement price of the EUA futures contract. To account for the fact that percentage surprises are not really meaningful in late 2007 when carbon prices were approaching zero, I exclude these events from the analysis. Finally, I also purge the surprise series from macro, financial, oil market and climate related news to remove any potential remaining predictability. The daily surprises are shown in Figure C.6.

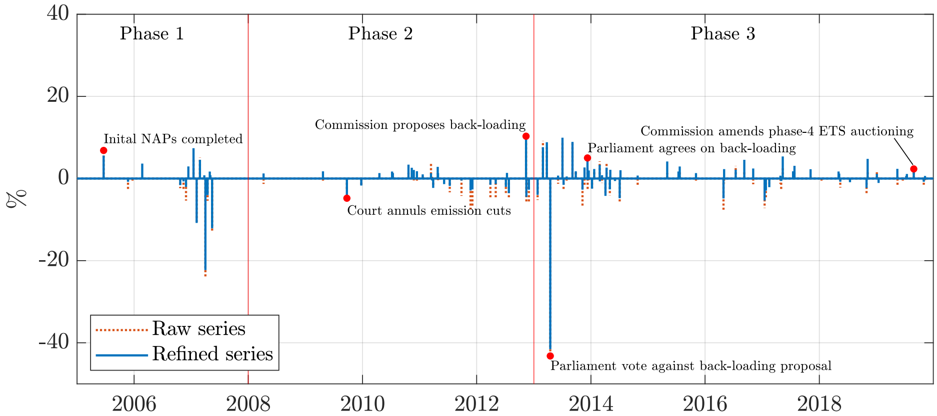

Figure C.6: Alternative Carbon Policy Surprise Series, Daily

Notes: The daily carbon policy surprise series, measured as the percentage change in carbon prices around event days, compared to the baseline carbon policy surprise series, which is expressed relative to the prevailing wholesale electricity price before the event.

Figure C.7 shows the corresponding aggregated monthly carbon policy surprise series, compared to the baseline surprise series. The two series are fairly similar, particularly during phase two and the beginning of the third phase. The correlation coefficient between the two series stands at around 0.76.

Figure C.8 presents the impulse responses to a carbon policy shock using the alternative instrument, together with the baseline responses. Using the alternative instrument produces similar results, however, the responses are somewhat less precisely estimated. This illustrates the importance of measuring the surprise series in economically relevant terms.

Figure C.7: Alternative Carbon Policy Surprise Series, Monthly

Notes: The monthly carbon policy surprise series, measured as the percentage change in carbon prices around event days, compared to the baseline carbon policy surprise series, which is expressed relative to the prevailing wholesale electricity price before the event.

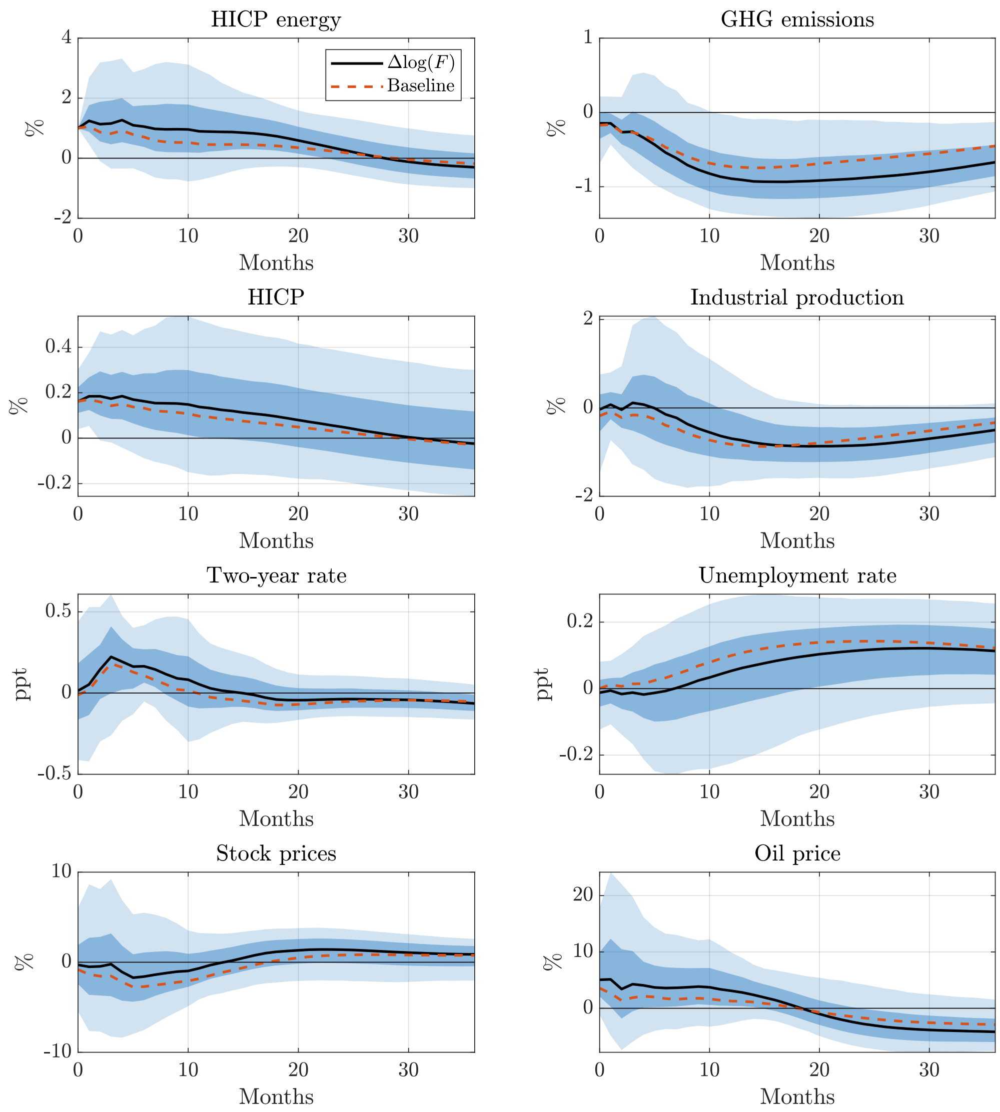

First stage regression: F-statistic: 12.24, R^2: 1.49%

Figure C.8: Responses Based on Alternative Surprise Series

Futures contracts. EUA futures are traded on both annual and quarterly cycles. Annual contracts expire in December and are available multiple years out, while quarterly contracts expire at the end of each quarter. As a baseline, I use the current December contract, which is the main contract used and typically the most liquid. However, at the start of the year, its distant expiry may raise concerns about risk and term premia. As an alternative, I use the front quarterly contract—the one with the nearest expiry. Figure C.9 shows that results are virtually identical across annual and quarterly futures.

I also examine how the results change when using contracts with longer maturities. Results from one-year, two-year, and four-year contracts are similar, further suggesting that risk premia do not significantly affect the findings.

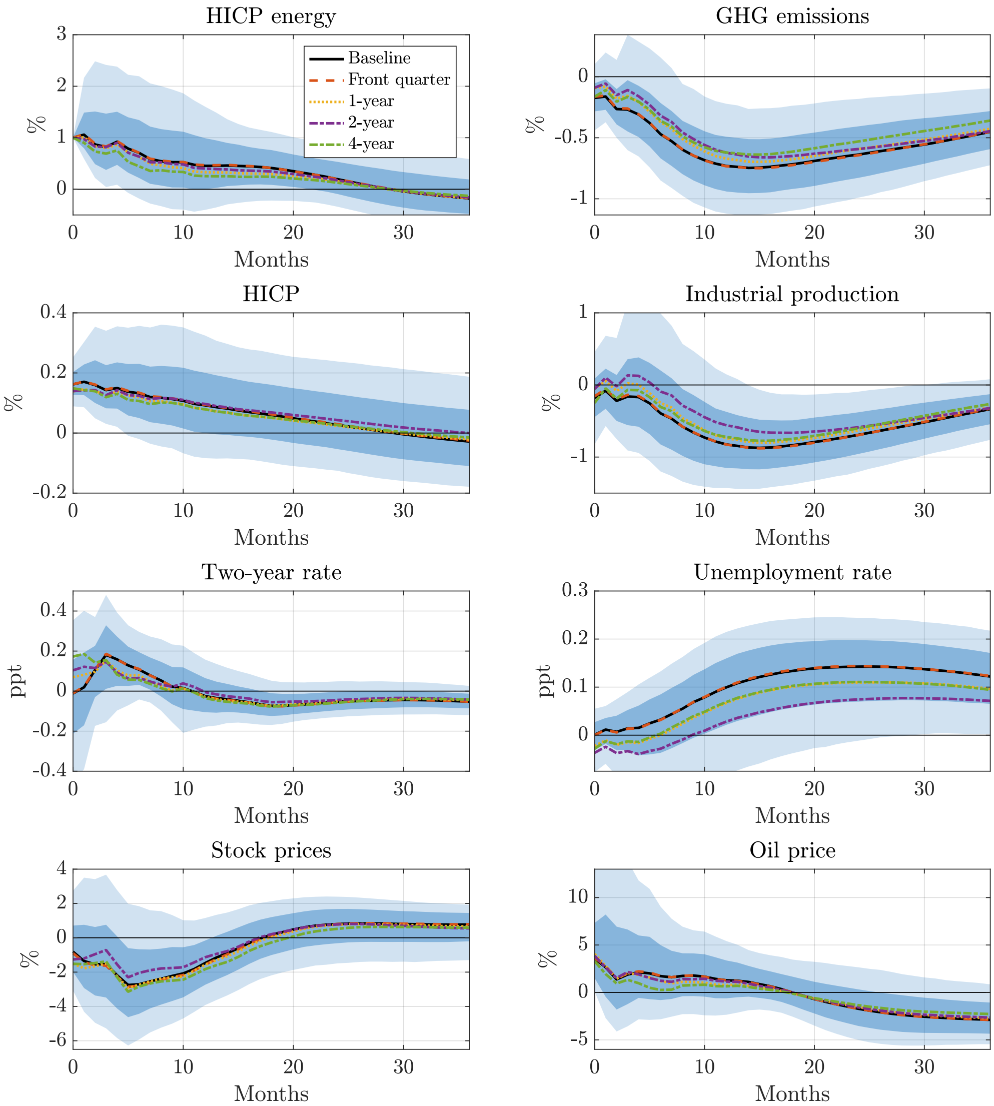

Figure C.9: Using Different Futures Contracts

C.2. Confounding news

Another important choice in high-frequency identification concerns the size of the event window. As discussed in Section 2.2, there is a trade-off between capturing the entire response to the policy news and background noise, i.e. the threat of other news confounding the response. Common window choices range from 30-minutes to multiple days. Unfortunately, the exact release times are unavailable for the majority of the policy events considered, making it infeasible to use an intraday window. Therefore, I use a daily window to compute the policy surprises.

To mitigate concerns about other news confounding the carbon policy surprise series, I employ an alternative identification strategy exploiting the heteroskedasticity in the data (Rigobon, 2003; Nakamura and Steinsson, 2018*a*). The idea is to clean out the background noise in the surprise series by comparing movements in carbon prices during policy event windows to other equally long and otherwise similar event windows that do not contain a regulatory update event.

The identification strategy works as follows. Suppose that the movements in the EUA futures z_t we observe in the data are governed by both carbon policy and other shocks:

$$z_t = \varepsilon_{1,t} + \sum_{j\neq 1} \varepsilon_{j,t} + v_t,$$

where $\varepsilon_{j,t}$ are other shocks affecting carbon futures and $v_t \sim iidN(0, \sigma_v^2)$ captures measurement error such as microstructure noise. Because z_t is also affected by other shocks, it is no longer a valid external instrument. However, we can still identify the structural impact vector by exploiting the heteroskedasticity in the data.

The identifying assumption is that the variance of carbon policy shocks increases at the time of regulatory update events while the variance of all other shocks is unchanged. Define $R1$ as a sample of regulatory events in the EU ETS and $R2$ as a sample of trading days that do not contain an regulatory event but are comparable on other dimensions. $R1$ can be thought of as the treatment and $R2$ as the control sample. The identifying assumptions can then be written as

$$\begin{aligned} \sigma^2_{\varepsilon_1,R1} &> \sigma^2_{\varepsilon_1,R2} \\ \sigma^2_{\varepsilon_j,R1} &= \sigma^2_{\varepsilon_j,R2}, \quad \text{for } j = 2, \dots, n. \\ \sigma^2_{v,R1} &= \sigma^2_{v,R2}. \end{aligned} \tag{11}$$

Under these assumptions, the structural impact vector is given by

$$\mathbf{s}_1 = \frac{\mathbb{E}_{R1}[z_t \mathbf{u}_t] - \mathbb{E}_{R2}[z_t \mathbf{u}_t]}{\mathbb{E}_{R1}[z_t^2] - \mathbb{E}_{R2}[z_t^2]}. \tag{12}$$

As shown by Rigobon and Sack (2004), we can also obtain this estimator through an IV approach, using $\tilde{\mathbf{z}} = (\mathbf{z}'_{R1}, -\mathbf{z}'_{R2})'$ as an instrument in a regression of the reduced-form innovations on $\mathbf{z} = (\mathbf{z}'_{R1}, \mathbf{z}'_{R2})'$.

To implement this approach, I construct a sample of control days. In particular, I use the changes in carbon futures prices on the same weekday and week in the months prior a given regulatory event. For the samples of event and control days, I compute surprises as the percentage change in EUA futures, purged from information available prior to the day. I use percentage changes in carbon prices as opposed to the change relative to the wholesale electricity price to capture the variance in carbon prices more directly.

The assumption that the variance is much larger on event days than on a sample of controls days is indeed supported by the data. Figure C.10 shows the carbon policy surprise series together with the control series. We can see that the policy surprise series is much more volatile than the control series, and a Brown–Forsythe test for the equality of group variances confirms that this difference is also statistically significant.

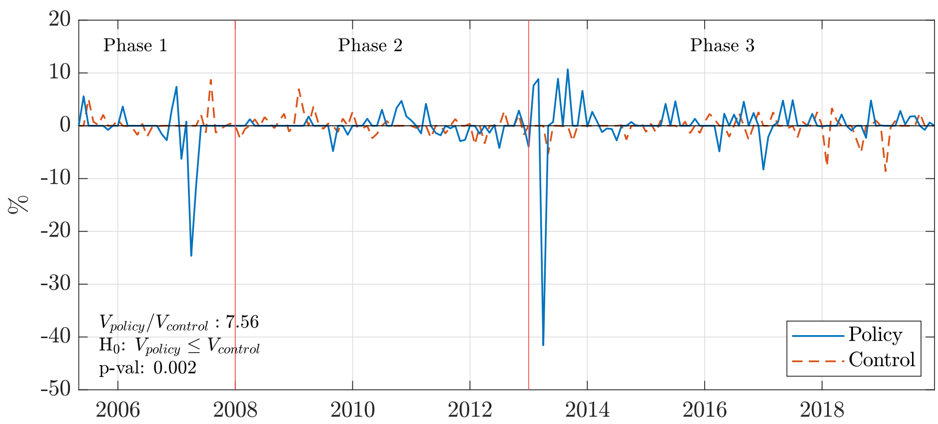

Figure C.10: The Alternative Carbon Policy and the Control Series

Notes: The alternative carbon policy surprise series based on the percentage change in carbon prices together with the surprise series constructed on a selection of control days.

Figure C.11 shows the impulse responses estimated from this alternative approach. The results turn out to be consistent with the baseline results from the external instrument approach, even though the responses are a bit less precisely estimated. These results suggest that the bias induced by background noise is likely negligible in the present application. However, part of the statistical strength under the external instrument approach appears to come from the stronger identifying assumptions.

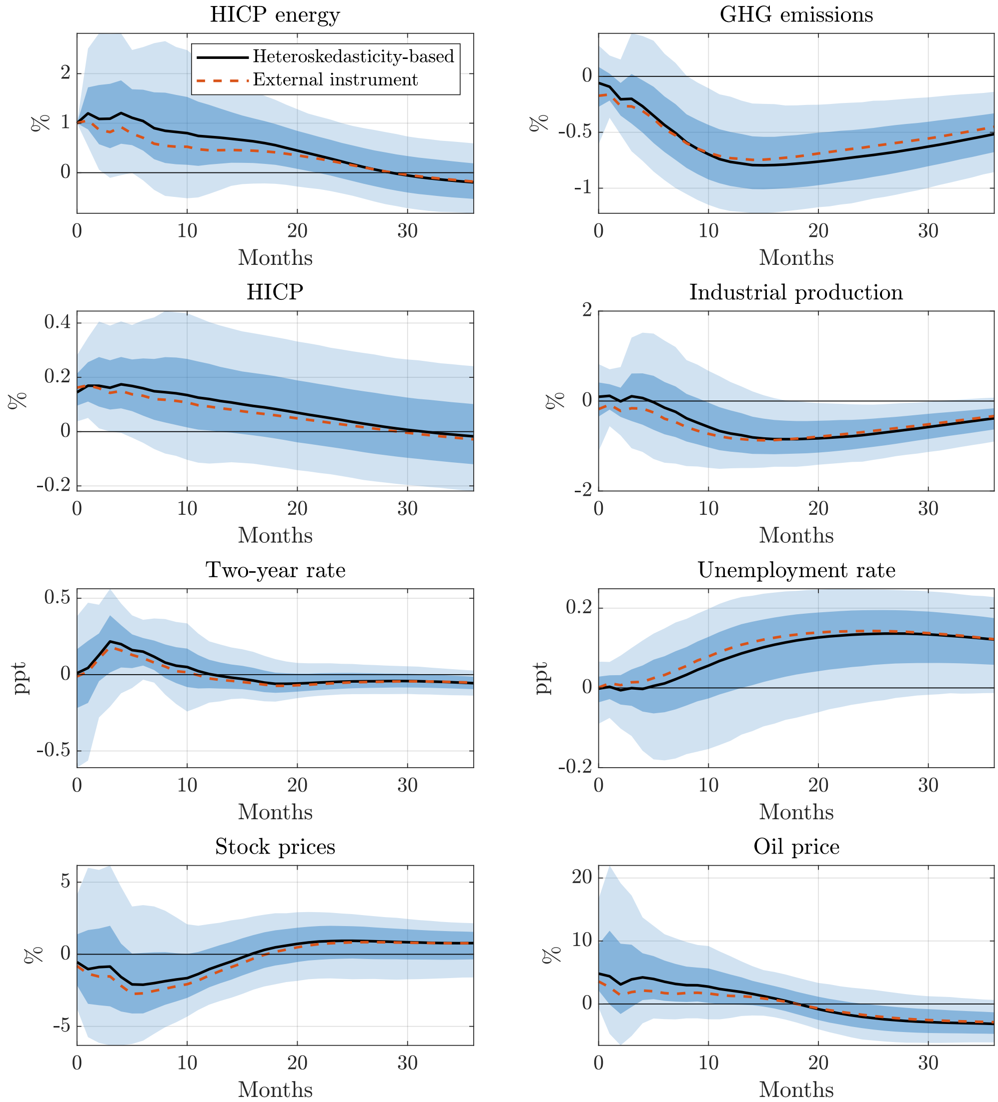

Figure C.11: Heteroskedasticity-based Identification

C.3. Placebo exercise

In Appendix C.2, I exploit the heteroskedasticity between event and non-event days explicitly for identification. In this appendix, I show that randomly-drawn placebo dates do not lead to systematic changes in energy prices, emissions and economic activity.

The exercise proceeds as follows. I first draw 114 placebo dates from the sample of non-event dates. Next, I construct a placebo surprise series, measuring changes in carbon prices on these dates, and remove any predictability using my extended set of predictive variables (d). I then aggregate the refined surprises into a monthly placebo series. Using this series as an external instrument in my baseline VAR model, I estimate impulse responses. I repeat these steps a 1,000 times.

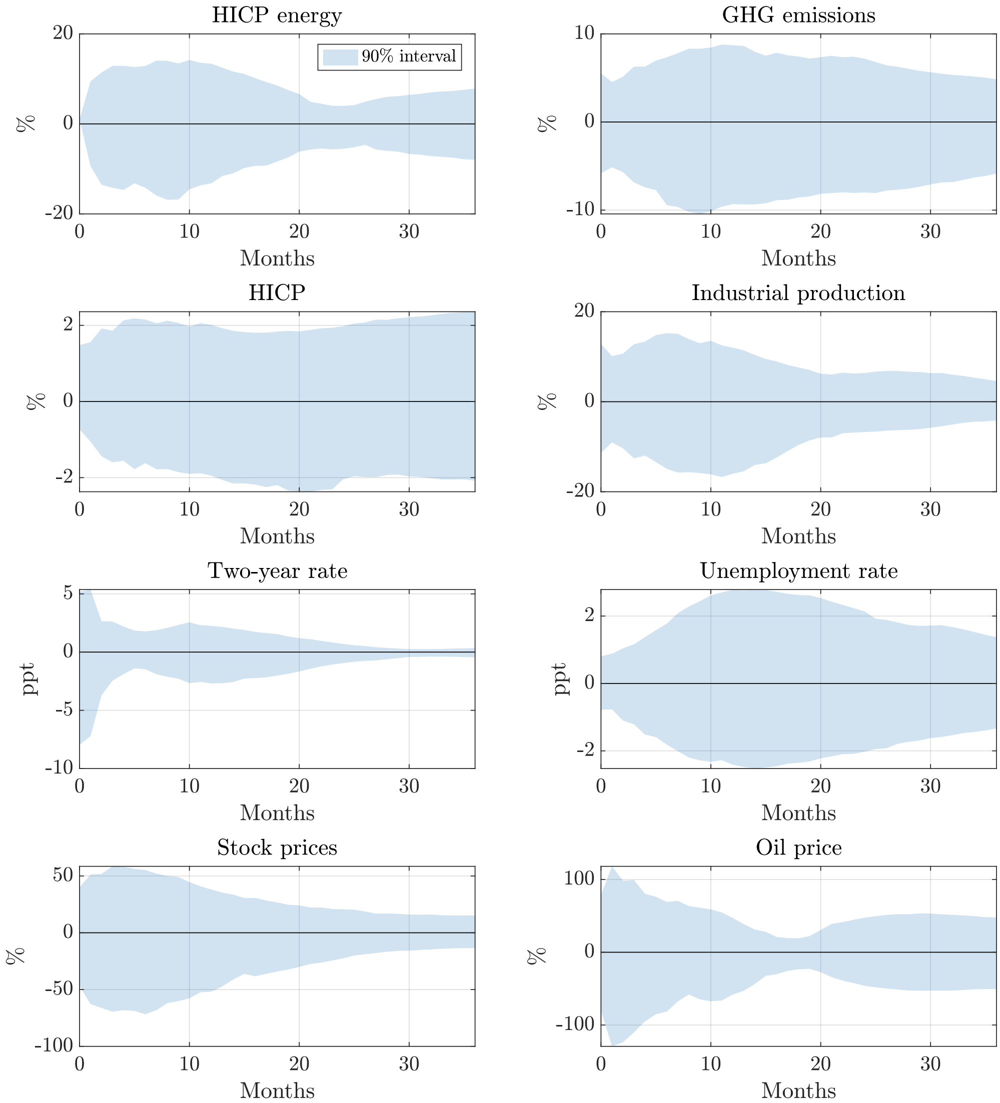

Figure C.12: Placebo Responses

Figure C.12 shows the 5th and 95th percentiles of the placebo responses. News on placebo dates are not associated with systematic effects on energy prices, emissions and economic and financial variables. This is not merely due to offsetting positive and negative responses—all responses are normalized to increase the HICP energy by 1 percent on impact. The width of the interval further underscores this point. Overall, this evidence suggests that the external instrument VAR approach is not picking up spurious correlations in the data.

Another way to assess this is by examining the first stage. Random placebo dates should capture a mix of different shocks, and as such, they should not produce a systematic response in energy prices. Consequently, the placebo surprises should be weak instruments. This intuition is confirmed by the distribution of F-statistics from the first-stage regressions, where placebos serve as instruments for the HICP energy residual.

Figure C.13 shows the histogram of the (robust) F-statistics. Most values cluster near zero and virtually all of the mass is below the relevant threshold of 10. The contrast with the F-statistic for the actual carbon policy surprise series is even more pronounced. These results confirm that the placebos are weak instruments, explaining the wide range of possible effects they produce.

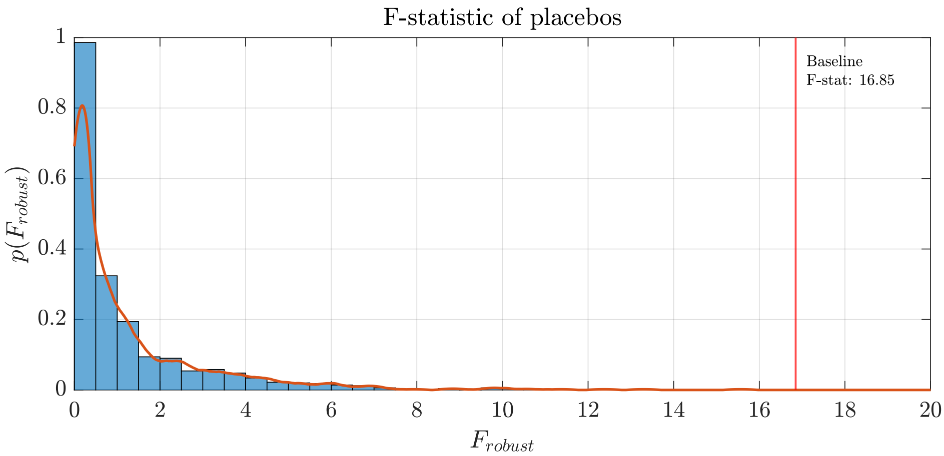

Figure C.13: Placebo F-Statistics

C.4. Sample and specification choices

An important robustness check concerns the estimation sample. Recall, the baseline sample goes back to 1999, which is longer than the instrument sample that only starts in 2005. The main motivation for using the longer sample is to increase the precision of the estimates. As a robustness check, I restrict the overall sample to the 2005-2019 period. The responses are shown in Figure C.14. Overall, the results are similar to the ones using the longer sample. However, the responses of emissions and industrial production turn out to be less pronounced. Some responses also turn out to be less stable, pointing to difficulties in estimating the model dynamics on the relatively short sample. As a second check, I stop the estimation sample before the oil shock in mid-2014. From Figure C.15, we can see that the responses turn out to be very similar to the baseline case. However, the responses overall are less precisely estimated, again illustrating the challenges of estimating the model on an even shorter sample.

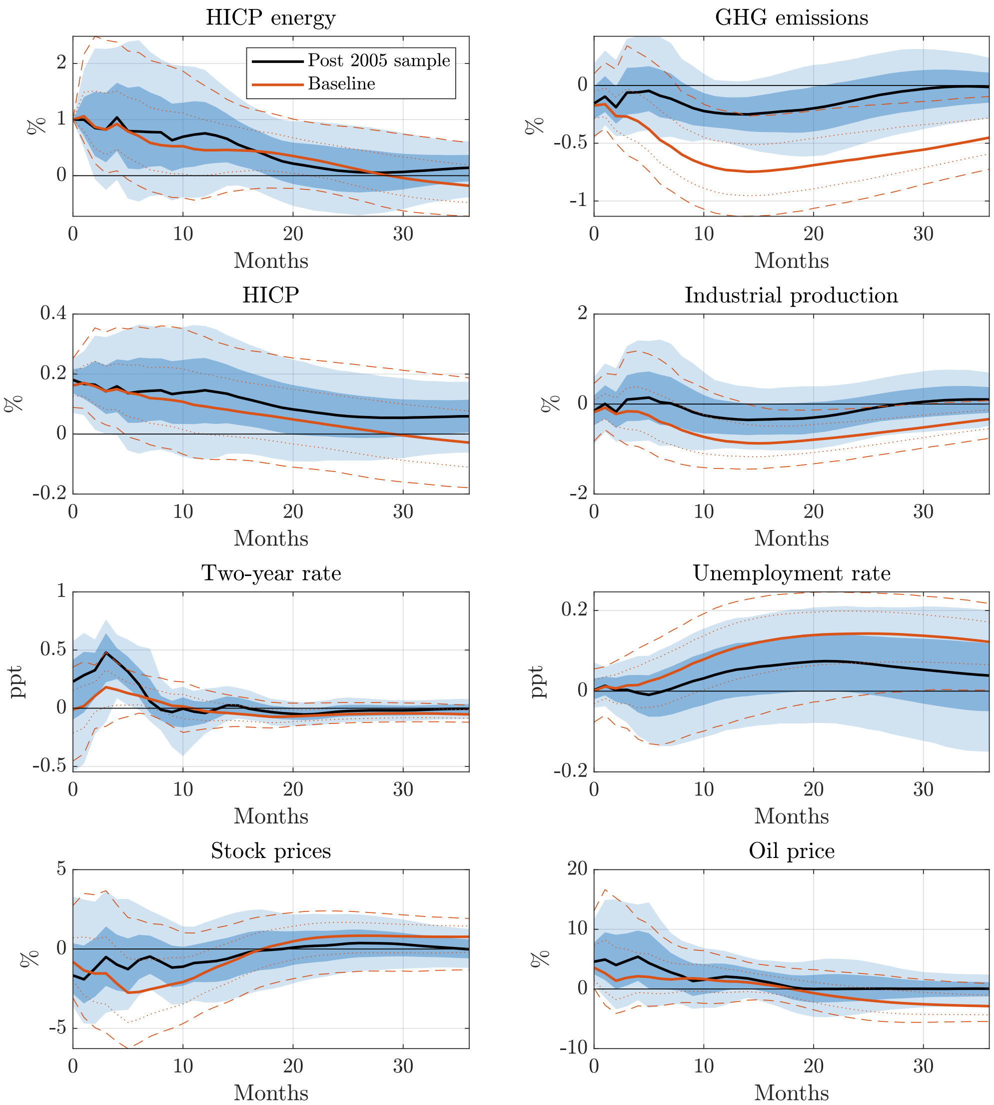

First stage regression: F-statistic: 10.95, R^2: 3.98%

Figure C.14: Results Using 2005-2019 Sample

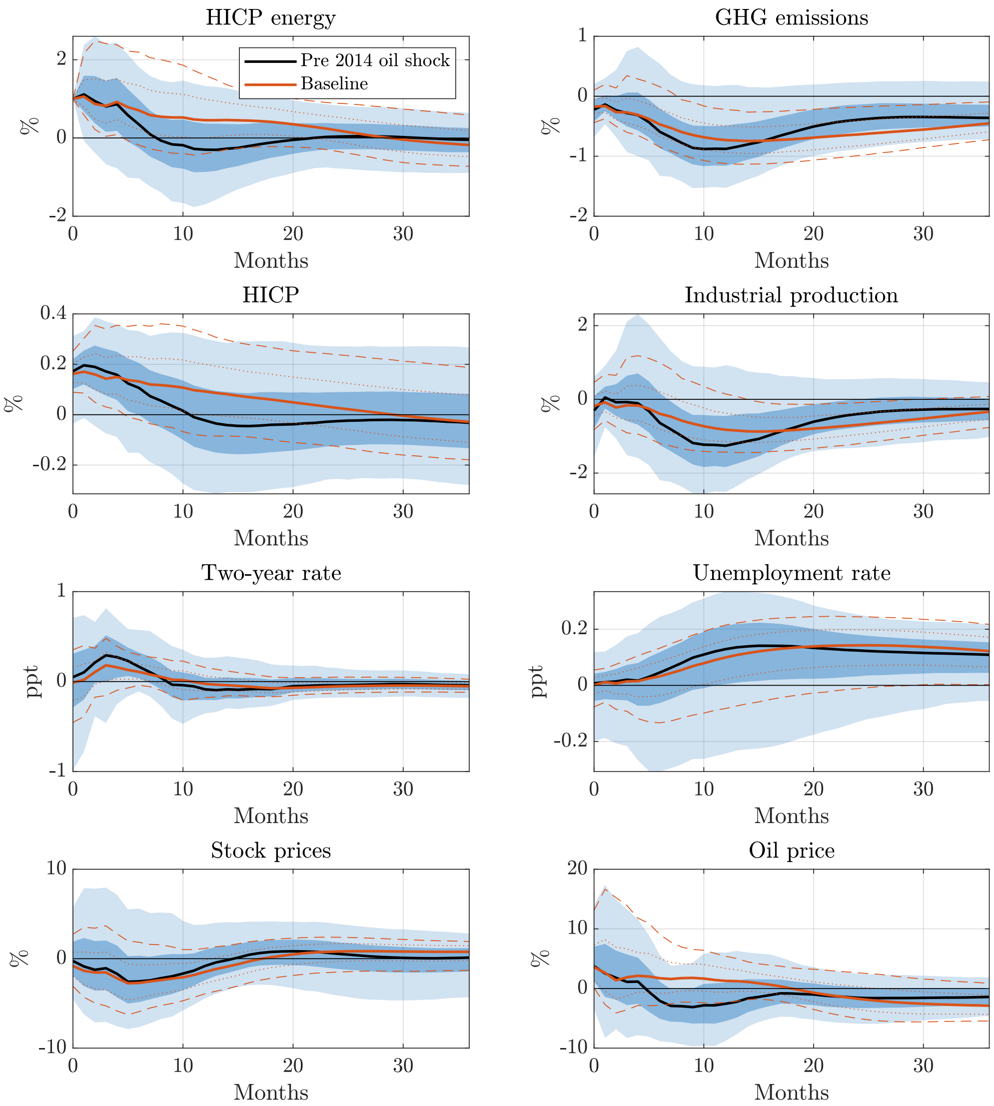

First stage regression: F-statistic: 10.94, R^2: 2.83%

Figure C.15: Results Using 1999-2014 Sample

I also perform a number of other sensitivity checks on the specification of the model. The baseline VAR includes 8 variables, which is relatively large given the short sample. As a robustness test, I use a 6-variable model, excluding stock and oil prices. As can be seen from Figure C.16, the results from this smaller model turn out to be consistent with the larger baseline model. The results are also robust to the lag order (Figure C.17 shows the responses using 3 or 9 lags) and the choice of deterministics (Figure C.18 shows responses of a model with a linear trend and a model excluding the dummy for the sovereign debt crisis).

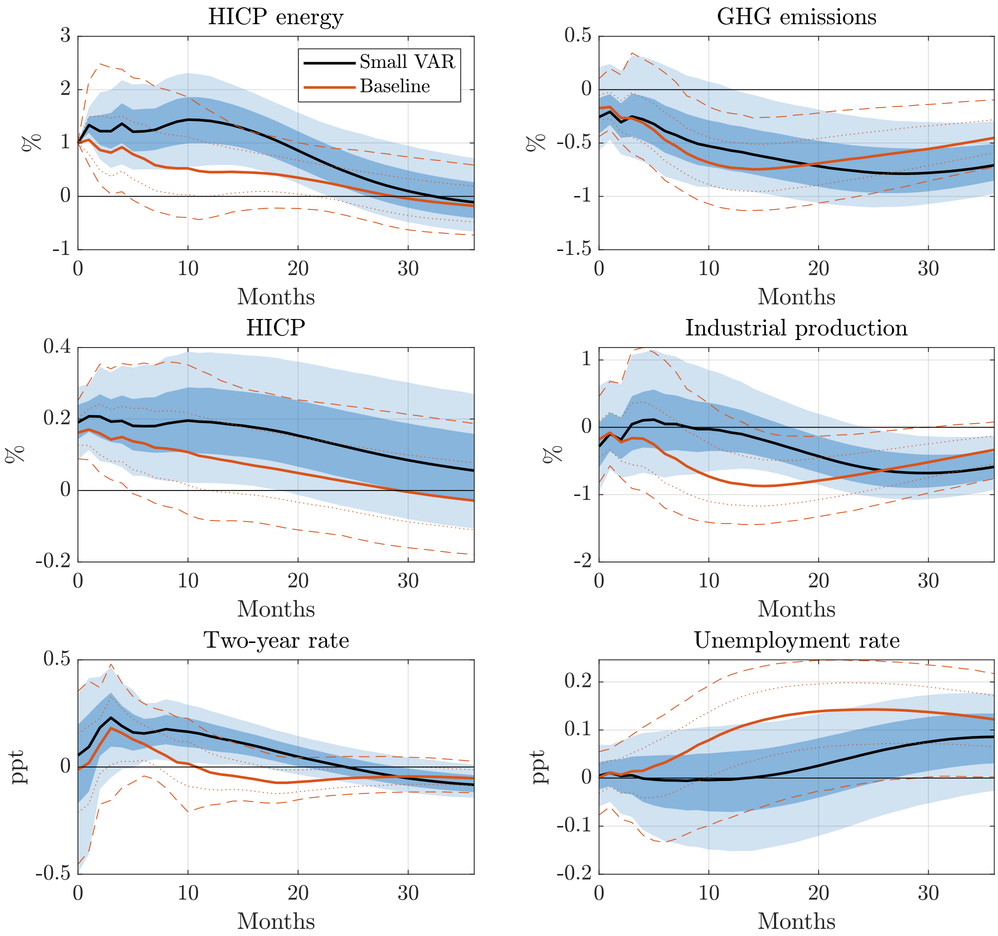

First stage regression: F-statistic: 6.49, R^2: 1.12%

Figure C.16: Responses from Smaller VAR

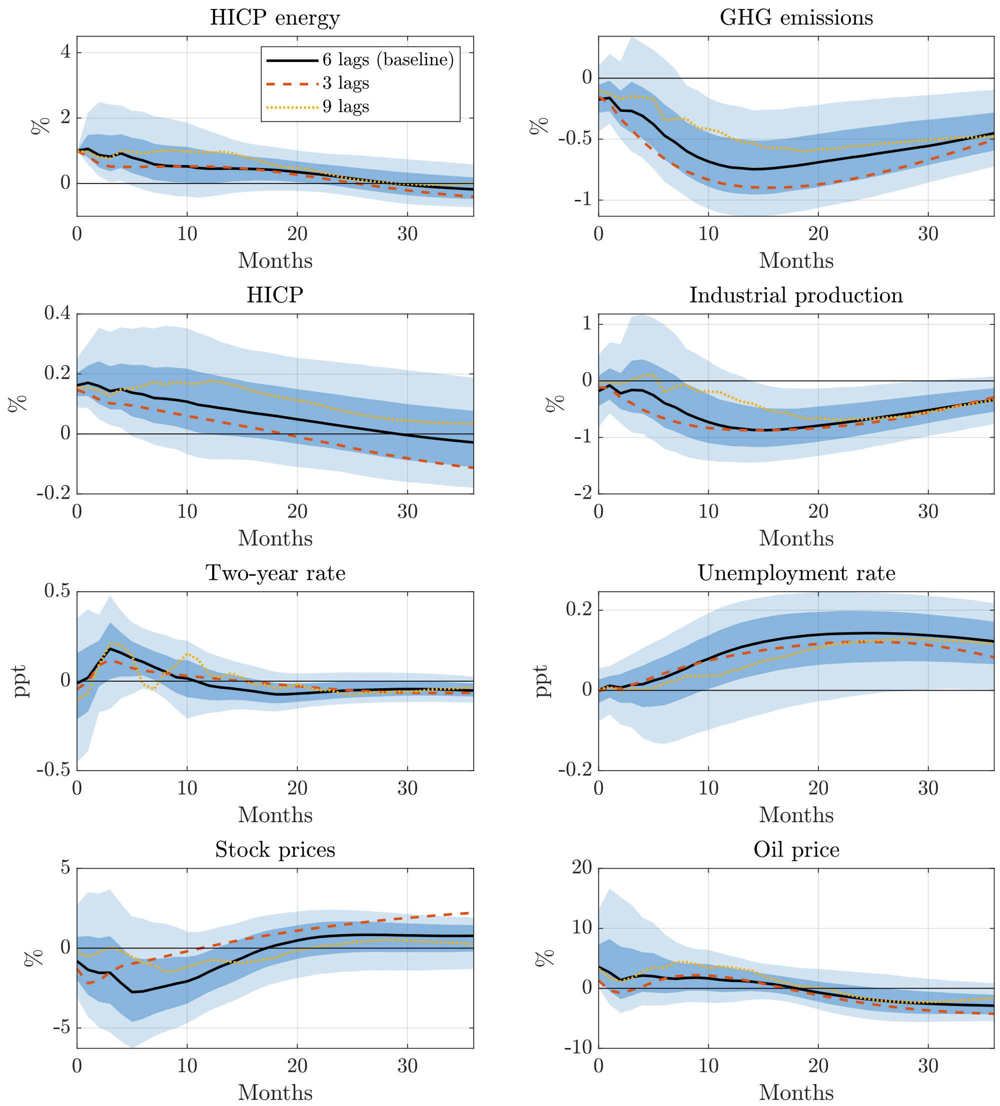

Figure C.17: Sensitivity to Lag Order

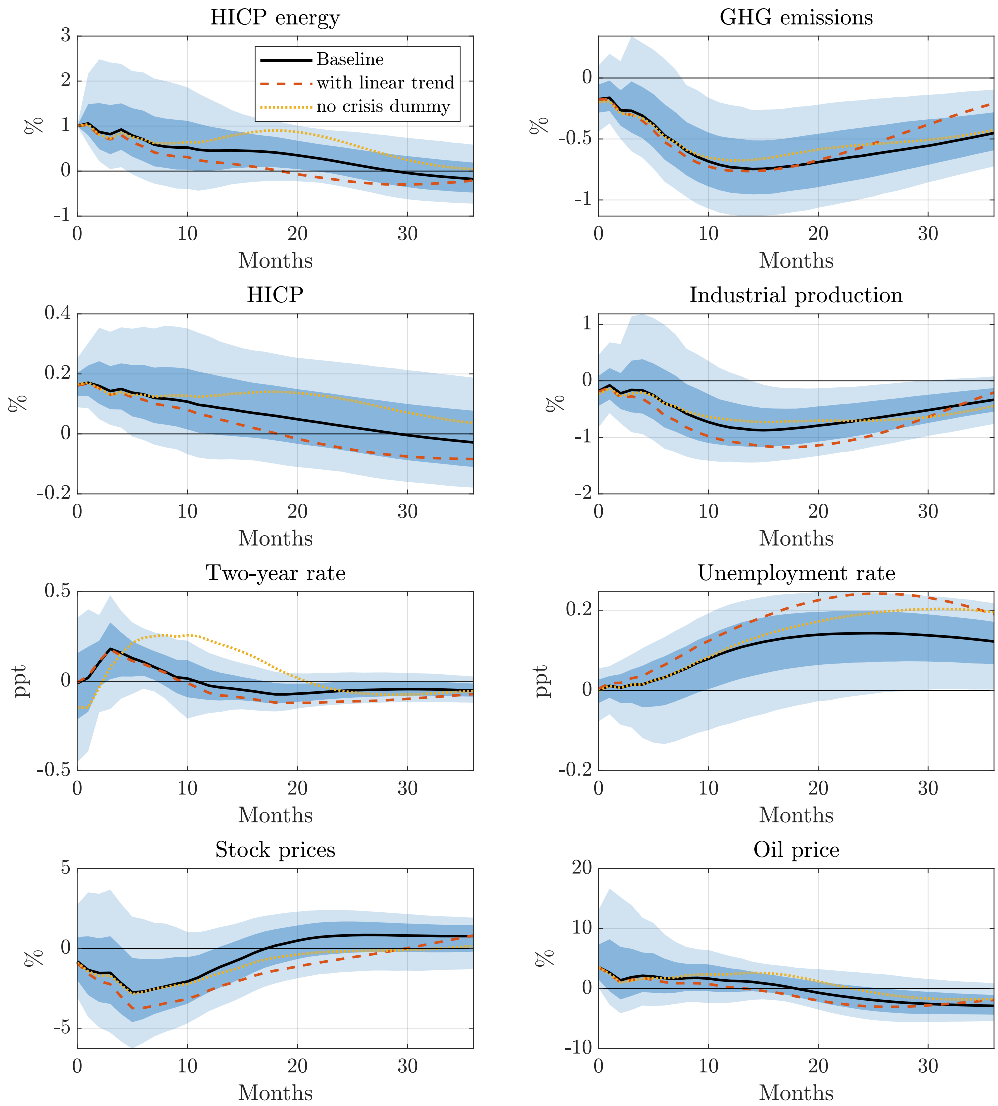

Figure C.18: Sensitivity to Deterministic Variables

C.5. Relaxing the invertibility requirement

A key advantage of the external instruments approach lies in its efficiency. However, this comes at the cost of assuming (partial) invertibility. If the invertibility assumption is not satisfied, this can lead to biased results (Li, Plagborg-Møller, and Wolf, 2024). To mitigate concerns about invertibility, I perform two additional exercises.

Internal instruments approach. First, I present results using the internal instruments approach (Ramey, 2011; Plagborg-Møller and Wolf, 2021), which involves placing the instrument first in a recursive VAR. This method is robust to non-invertibility.

Figure C.19 compares impulse responses from the internal instruments VAR with the external instrument baseline. Since the system is now larger, I reduce the number of lags to three, improving stability without materially affecting the results. The responses remain qualitatively and quantitatively similar, though the estimated response of the two-year rate is somewhat less stable. Notably, the internal instrument responses are much less precisely estimated, as indicated by wider confidence bands. Overall, these findings suggest that the results are robust to relaxing the invertibility assumption, but assuming invertibility improves precision.

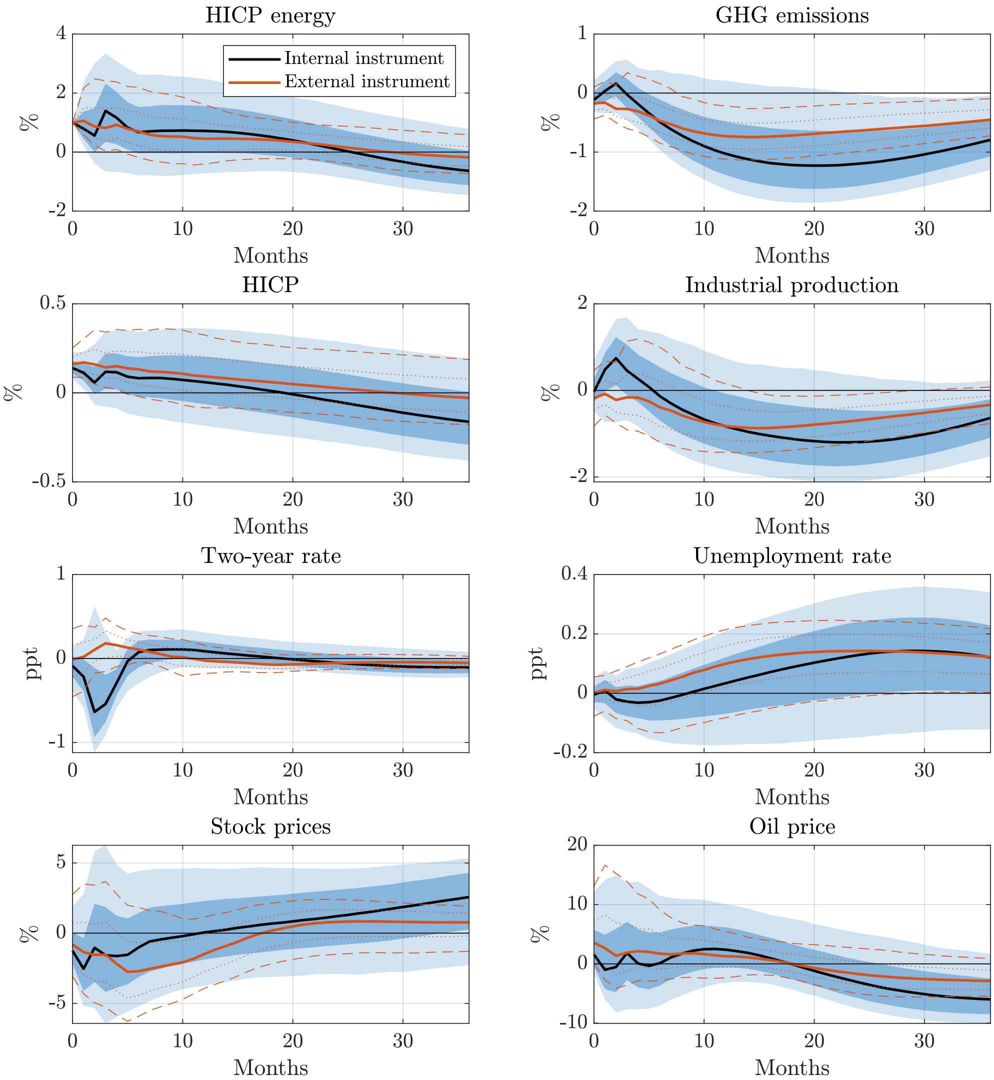

Figure C.19: Internal Versus External Instrument VAR

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact. The solid dark and red lines are the point estimates for the internal instrument and the external instrument VAR, respectively, and the shaded areas / dashed lines are 68 and 90 percent confidence bands.

Local projection-instrumental variable approach. As an invertibility-robust alternative, I estimate the impulse responses using a local projections instrumental variable (LP-IV) approach à la Jordà, Schularick, and Taylor (2015) and Ramey and Zubairy (2018). To fix ideas, the dynamic causal effects, θ_h^i, can be estimated from the following set of regressions:

$$y_{i,t+h} = \beta_h^i + \theta_h^i y_{1,t} + \boldsymbol{\beta}_h^{i\prime} \mathbf{x}_{t-1} + \xi_{i,t,h}, \tag{13}$$

using z_t as an instrument for $y_{1,t}$. Here, $y_{i,t+h}$ is the outcome variable of interest, $y_{1,t}$ is the endogenous regressor, $\mathbf{x}_{t-1}$ is a vector of controls, $\xi_{i,t,h}$ is a potentially serially correlated error term, and h is the impulse response horizon. I use the same controls as in the internal instruments VAR. For inference, I follow again the lag-augmentation approach from Montiel Olea and Plagborg-Møller (2021).

As the impacts of carbon policy are potentially quite persistent, we want to look at the dynamic causal effects relatively far out. This is challenging in the LP-IV framework because of a power problem (Nakamura and Steinsson, 2018*a*): the macroeconomic outcomes several quarters or years out are affected by a myriad of other shocks, rendering the signal-to noise ratio from the relatively small carbon policy shocks too low to credibly identify the effects of interest. Additionally, each increase in the impulse horizon h reduces the degrees of freedom, further complicating estimation given the short sample. To address these limitations, I restrict the impulse horizon in the LP-IV regressions to 12 months.

Figure C.20 compares the responses from the LP-IV approach to those from the internal instrument VAR. Both methods rely on the same invertibility-robust identifying restrictions but differ in their estimation techniques. The results are broadly consistent, particularly for horizons up to one year. At longer horizons the differences tend to be larger, however, the responses are also much less precisely estimated.

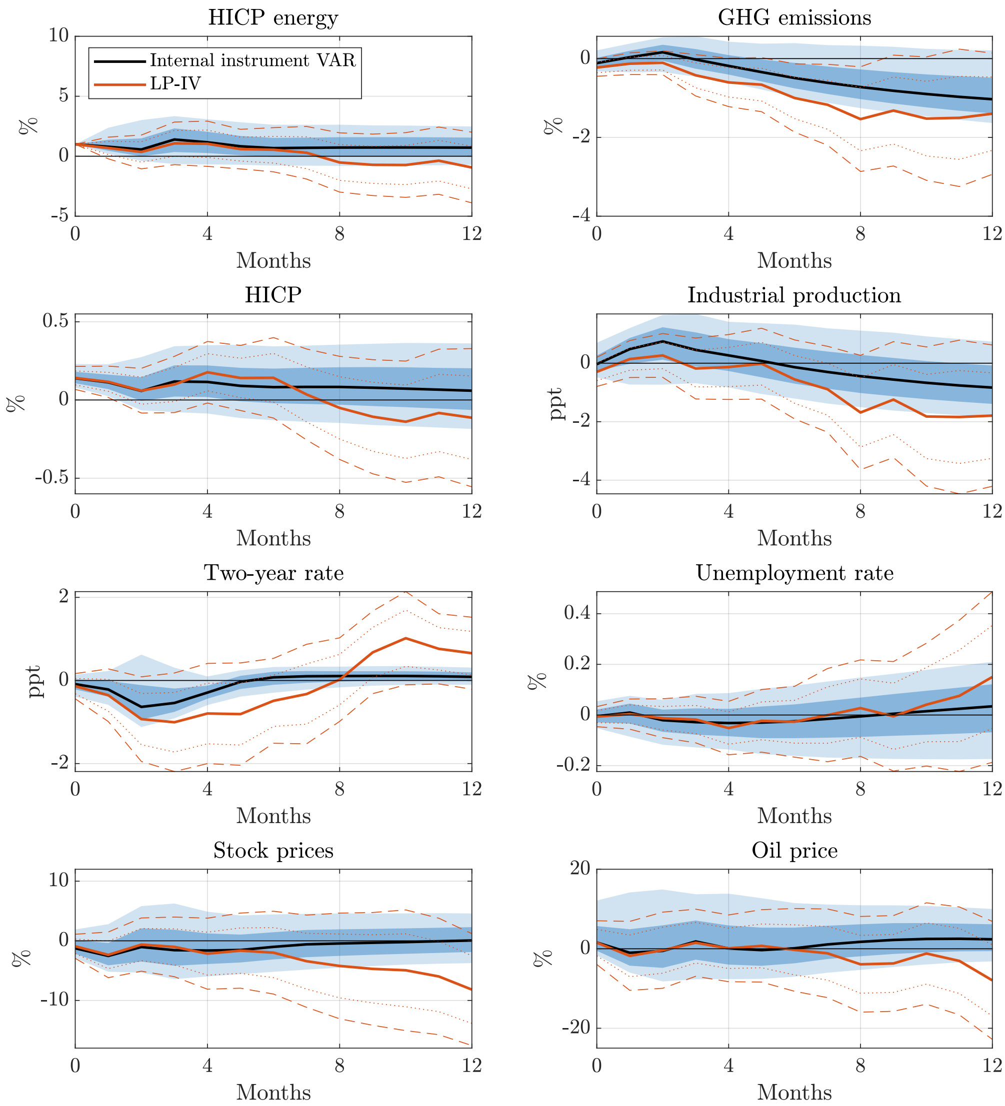

Figure C.20: Internal Instrument VAR Versus LP-IV

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact. The solid dark and red lines are the point estimates for the internal instrument VAR and the LP-IV, respectively, and the shaded areas / dashed lines are 68 and 90 percent confidence bands.

C.6. Estimation method

As discussed in the main text, I use VAR techniques for estimation because the sample is relatively short, and VARs offer a parsimonious characterization of the data. Another important assumption underlying this approach is that the finite-order VAR adequately approximates the data generating process, implying negligible lag truncation bias.

To assess the extent of lag truncation bias, I conduct two exercises. First, I estimate the responses using local projections, relaxing the dynamic VAR structure while maintaining the invertibility assumption. Second, I employ Bayesian methods to estimate a VAR with a longer lag order.

Frequentist approach. Section 3 in the main text discusses how local projections can be used to estimate effects on additional outcome variables. Here, I apply this approach to re-estimate the responses of the baseline model variables. This involves regressing each variable of interest on the identified VAR shock and its own lags, as in (8).

Alternatively, I employ the approach originally proposed in Jordà (2005). This amounts to directly estimating the reduced-form responses $\boldsymbol{\psi}_0^h$:

$$\mathbf{y}_{t+h} = \mathbf{b} + \boldsymbol{\psi}_0^h \mathbf{y}_t + \cdots + \boldsymbol{\psi}_p^h \mathbf{y}_{t-p} + \mathbf{u}_{t+h}, \tag{14}$$

and computing the structural responses as $\boldsymbol{\theta}_h = \boldsymbol{\psi}_0^h \boldsymbol{\theta}_0$. Both estimators are similar in spirit: they rely on the invertibility assumption but relax the dynamic structure of the VAR. In other words, both approaches allow for a more robust estimation of the autocovariance function. The key difference is that the first approach directly estimates structural impulse responses, whereas the second estimates reduced-form responses and translates them into structural responses using the structural impact vector.

Figure C.21 presents the results. The responses closely resemble those from the baseline VAR, with the most notable difference appearing in the emissions response, which is somewhat less stable. In addition, the local projection estimates are less precisely estimated.

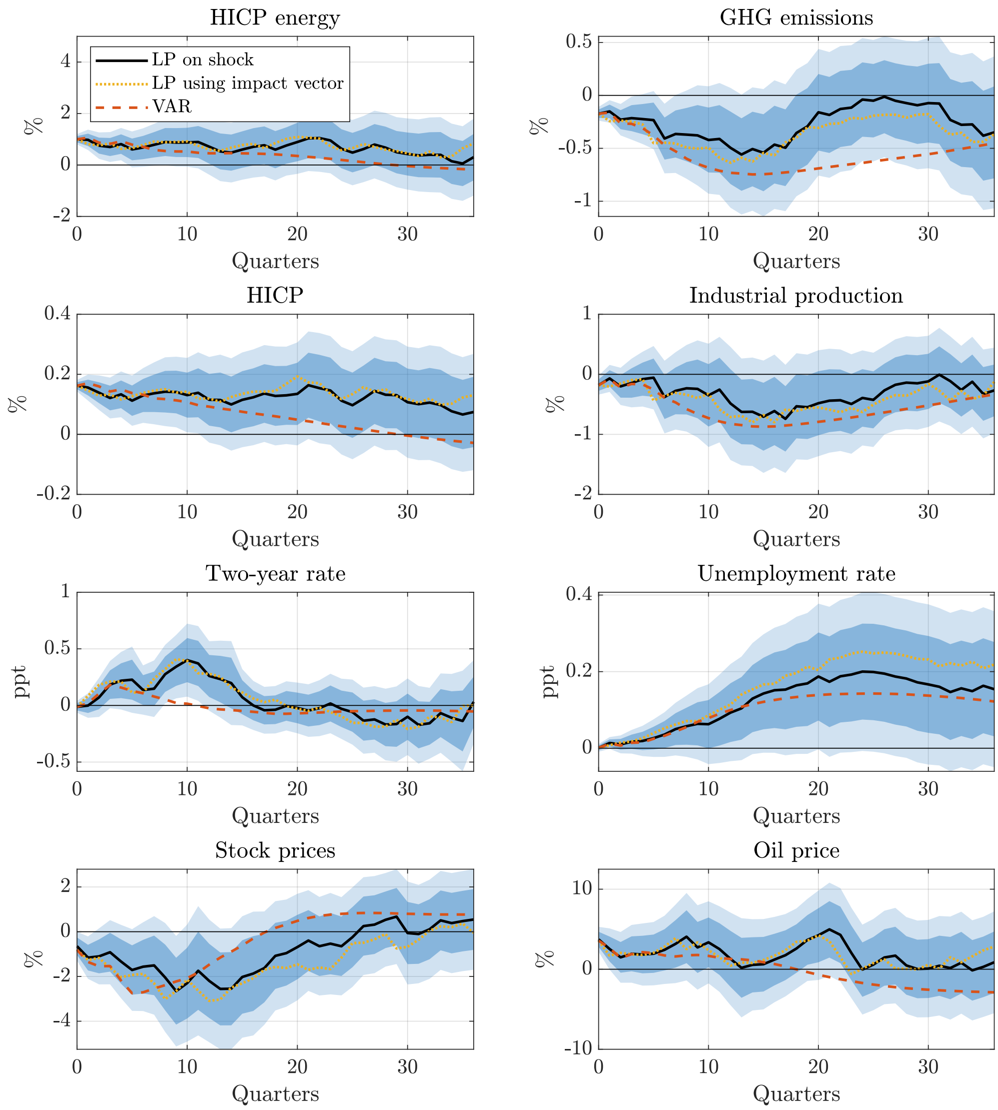

Figure C.21: Alternative Frequentist Estimation

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, estimated using (i) simple local projections on the carbon policy VAR shock and (ii) based on reduced-form LP responses and the structural impact vector from the VAR, compared to the baseline VAR responses. Solid and dashed lines: point estimates. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

Bayesian approach. Alternatively, I use Bayesian techniques for estimation. Specifically, I estimate a Bayesian VAR with shrinkage priors, allowing me to extend the number of lags up to 18. Specifically, I use a Minnesota prior with tightness 0.3 and a decay of 1. Figure C.22 presents the results. The responses based on the Bayesian models align well with the responses from the baseline VAR estimated using frequentist methods.

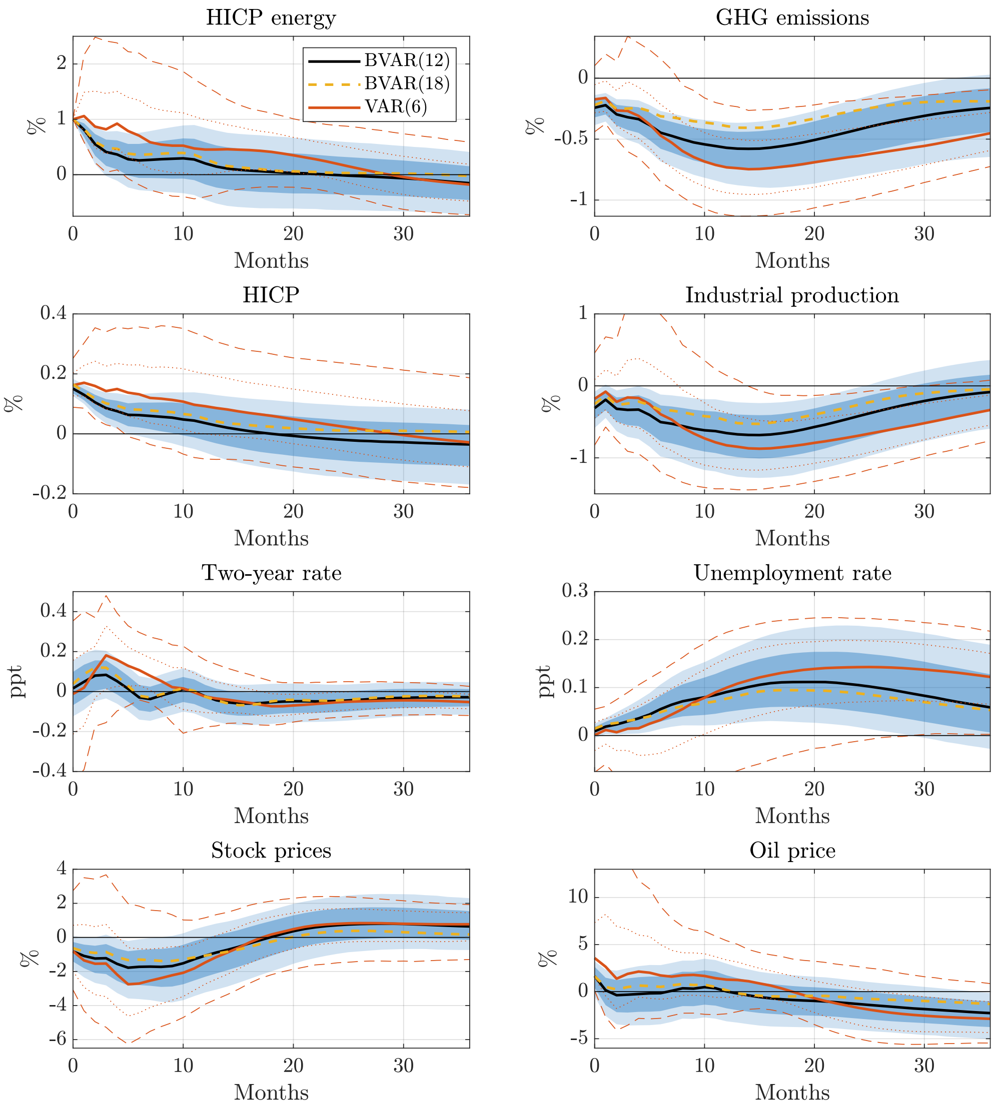

Figure C.22: Alternative Bayesian Estimation

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, estimated using a Bayesian VAR with Minnesota priors (tightness=0.3, decay=1). We consider BVARs with 12 and 18 lags, compared to the baseline frequentist VAR(6) responses. Solid and dashed lines: posterior medians/point estimate. Dark and light shaded areas: 68 and 90 percent HPD intervals. Light dashed lines: 68 and 90 percent bootstrapped confidence bands.

Overall, these findings suggest that the results are robust to the estimation method, suggesting that lag truncation bias is unlikely to be a major concern in this application.

C.7. Inference in local projections

The lag-augmentation approach used for inference in the local projections (8) does not account for the estimation uncertainty in the VAR shock, treating the estimated shock as if it were directly observed. To assess how consequential this potential generated regressor problem is for inference, I alternatively compute the standard errors using bootstrapping techniques. Specifically, I independently sample from the residuals of the local projection model and the carbon policy shock, taking the sampling error of the carbon policy shock into account.

Figure C.23 compares the confidence bands for the responses of real GDP, consumption, investment and wages based on the baseline standard errors from the lag-augmentation approach to the standard errors constructed based on the bootstrap.

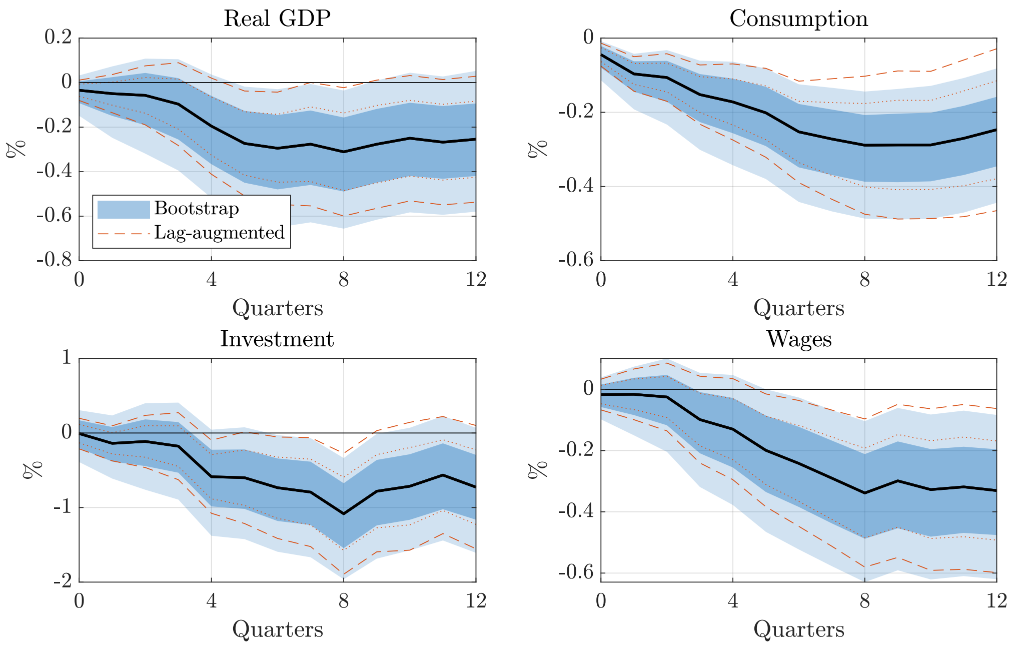

Figure C.23: Inference in Local Projections on Carbon Policy Shock

Notes: Impulse responses of a selection of quarterly variables estimated using local projections on the carbon policy shock. The responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline model. The bootstrapped 68 and 90 percent confidence bands are depicted as the light and dark blue shaded areas, respectively. For comparison, the confidence bands based on the lag-augmentation approach, that do not account for sampling uncertainty of the carbon policy shock, are included as the dashed and dotted orange lines.

Interestingly, the confidence bands from both methods are very similar, with the bootstrap-based bands being only marginally wider. This suggests that accounting for the sampling error in the generated regressor does not significantly alter the conclusions in the present application.

D. Additional Aggregate Results

In this Appendix, I present some additional results pertaining to the analysis in Section 4 of the paper.

D.1. Main results with 95 percent bands

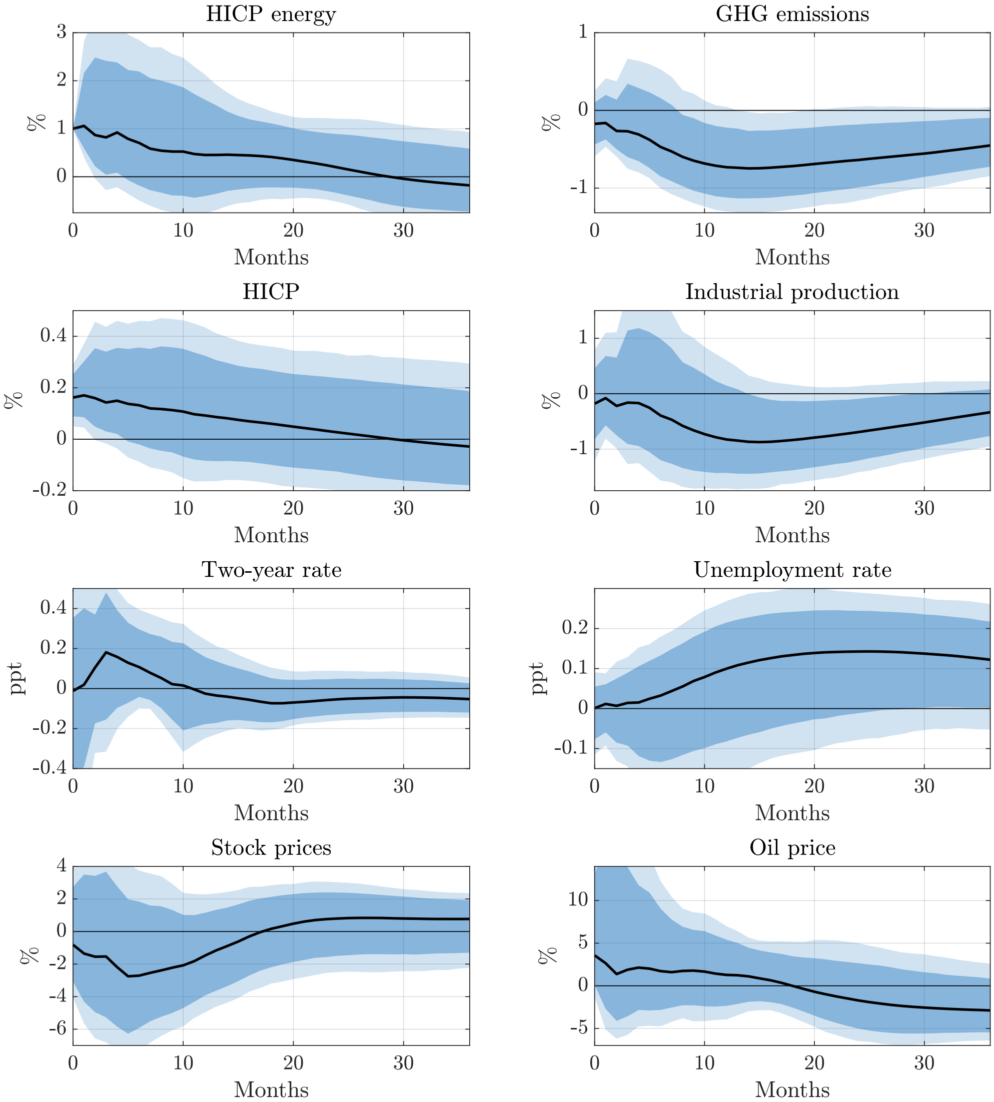

First stage regression: robust F-statistic: 16.85, R^2: 2.38%

Figure D.1: VAR Responses with 95 Percent Bands

Notes: Impulse responses to a carbon policy shock, normalized to increase the HICP energy by 1 percent on impact, estimated based on the VAR model (4) using the refined carbon policy surprise series from specification (d) as an instrument. Lag order: 6. Solid line: point estimate. Dark and light shaded areas: 90 and 95 percent confidence bands based on moving-block bootstrap.

Throughout the paper, I report 68 and 90 percent confidence bands, which is standard practice in the macroeconomics literature, acknowledging the power problems resulting from the relatively short available sample periods. In microeconomic work and in science, 90 and 95 percent confidence bands are more standard. To allow for better comparability with these literatures, I reproduce the main results from the paper with 90 and 95 percent bands in Figures D.1-D.2.

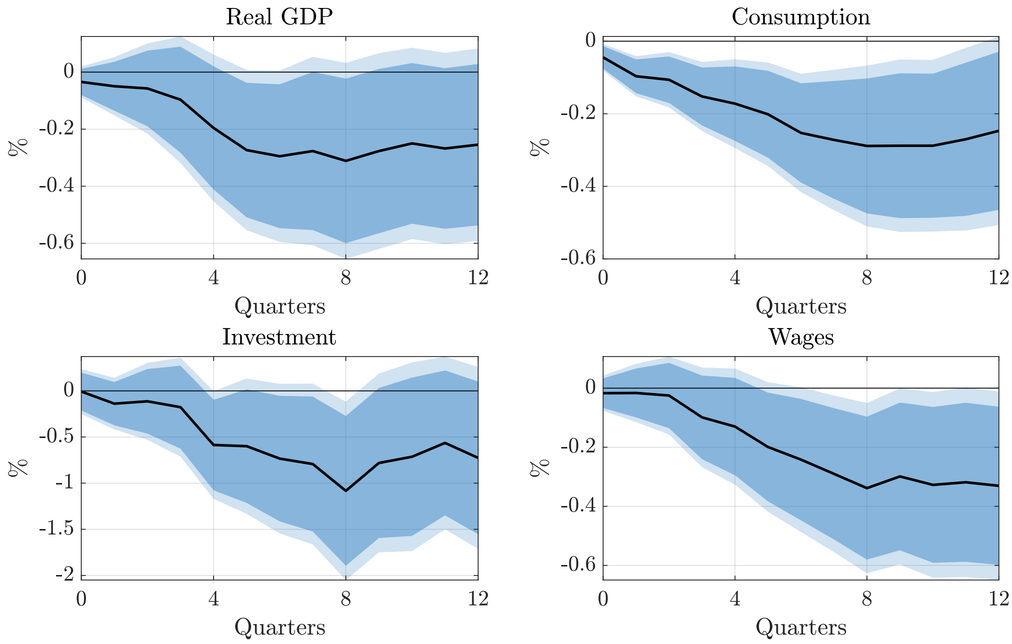

Figure D.2: Local Projection Responses with 95 Percent Bands

Notes: Impulse responses of a selection of quarterly variables, estimated using local projections (8) of the variable of interest on the carbon policy shock from the baseline VAR. Responses are normalized to have the same quarterly peak effect on HICP energy as in the baseline VAR. Controls: 3 lags of outcome variable and linear trend. Solid line: point estimate. Dark and light shaded areas: 90 and 95 percent confidence bands based on lag-augmentation approach.

D.2. Core versus headline HICP

In the main text, I document a significant and persistent increase in headline HICP. An important question that has also relevant implications for the conduct of monetary policy is how the shock transmits to core consumer prices. To this end, I re-estimate the model substituting headline for core HICP. Figure D.3 presents the response for core HICP together with the HICP headline and energy component from the baseline model. The response of core consumer prices is more muted and less precisely estimated. This illustrates that this is really a shock to relative prices. Reassuringly, all other responses from the model with core HICP are very similar to the baseline case.

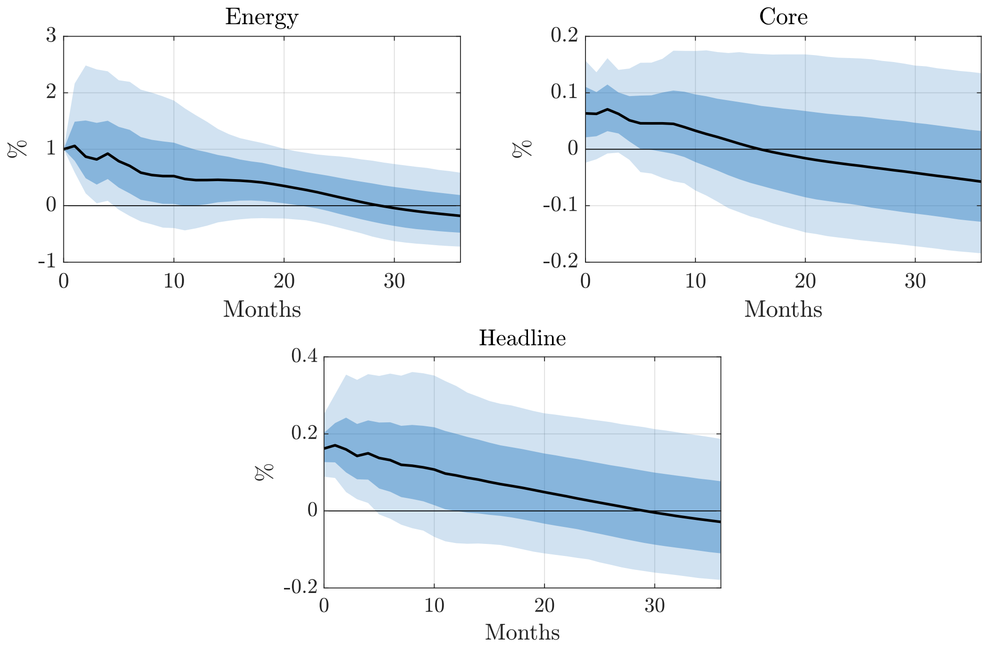

Figure D.3: Headline Versus Core HICP

Notes: Impulse responses of the headline, energy and core HICP to a carbon policy shock. The headline and energy indices are from the baseline model; the core response is from the model featuring core instead of headline HICP. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

D.3. Variance Decomposition

To better understand how carbon policy shocks have contributed to variations in macroeconomic and financial variables, I perform a variance decomposition exercise in addition to the historical decomposition presented in the paper. I do so both under the invertibility assumption maintained in the external instrument VAR as well as under weaker assumptions in the context of a general SVMA model, as proposed by Plagborg-Møller and Wolf (2022). In particular, I perform a standard forecast error variance decomposition in the SVAR and compute forecast variance ratios for the SVMA. The forecast variance ratio for variable i at horizon h is given by

$$FVR_{i,h} = 1 - \frac{\mathrm{Var}(y_{i,t+h}|\{y_\tau\}_{-\infty<\tau\leq t}, \{\varepsilon_{1,\tau}\}_{t<\tau<\infty})}{\mathrm{Var}(y_{i,t+h}|\{y_\tau\}_{-\infty<\tau\leq t})}, \tag{15}$$

and measures the reduction in the econometrician's forecast variance that would arise from being told the entire path of future realizations of the shock of interest. Plagborg-Møller and Wolf (2022) show that this statistic is interval-identified under the assumption that a valid instrument is available. Under the assumption of recoverablity, the ratio is point-identified by the upper bound.

Table D.1: Variance Decomposition

h	HICP energy	Emissions	HICP	IP	Two-year rate	Unemp. rate	Stock prices	Oil price
Panel A: Forecast variance decomposition (SVAR-IV)								
6	0.21	0.22	0.32	0.03	0.04	0.01	0.06	0.04
	[0.03, 0.39]	[0.02, 0.49]	[0.03, 0.52]	[0.01, 0.27]	[0.01, 0.22]	[0.00, 0.30]	[0.01, 0.31]	[0.01, 0.27]
12	0.15	0.35	0.21	0.09	0.03	0.05	0.07	0.04
	[0.03, 0.36]	[0.03, 0.56]	[0.03, 0.45]	[0.01, 0.32]	[0.02, 0.21]	[0.01, 0.32]	[0.01, 0.31]	[0.01, 0.27]
24	0.13	0.42	0.12	0.18	0.05	0.12	0.06	0.05
	[0.03, 0.33]	[0.03, 0.54]	[0.02, 0.38]	[0.02, 0.35]	[0.02, 0.22]	[0.01, 0.36]	[0.01, 0.30]	[0.02, 0.26]
36	0.12	0.40	0.09	0.20	0.06	0.15	0.07	0.07
	[0.03, 0.30]	[0.03, 0.51]	[0.02, 0.35]	[0.03, 0.34]	[0.03, 0.23]	[0.01, 0.36]	[0.02, 0.31]	[0.03, 0.26]
Forecast variance ratio (SVMA-IV)								
6	0.04, 0.25	0.01, 0.05	0.05, 0.30	0.00, 0.03	0.03, 0.22	0.00, 0.03	0.00, 0.01	0.02, 0.12
	[0.01, 0.39]	[0.00, 0.23]	[0.01, 0.44]	[0.00, 0.23]	[0.01, 0.41]	[0.00, 0.23]	[0.00, 0.20]	[0.00, 0.31]
12	0.03, 0.19	0.02, 0.16	0.03, 0.22	0.00, 0.02	0.04, 0.23	0.00, 0.01	0.00, 0.01	0.02, 0.12
	[0.01, 0.38]	[0.00, 0.39]	[0.00, 0.43]	[0.00, 0.28]	[0.01, 0.41]	[0.00, 0.25]	[0.00, 0.28]	[0.01, 0.32]
24	0.02, 0.16	0.05, 0.30	0.03, 0.16	0.01, 0.06	0.03, 0.22	0.00, 0.02	0.00, 0.01	0.02, 0.11
	[0.01, 0.37]	[0.00, 0.47]	[0.00, 0.44]	[0.00, 0.30]	[0.02, 0.37]	[0.00, 0.30]	[0.00, 0.28]	[0.01, 0.31]
36	0.02, 0.14	0.05, 0.31	0.02, 0.13	0.01, 0.09	0.04, 0.23	0.00, 0.03	0.00, 0.01	0.02, 0.11
	[0.01, 0.33]	[0.00, 0.46]	[0.00, 0.42]	[0.00, 0.31]	[0.02, 0.37]	[0.00, 0.33]	[0.00, 0.27]	[0.01, 0.30]

Notes: The table shows the variance decomposition at horizons ranging from 6 months to 4 years. Panel A includes the forecast error variance decomposition from the external instrument VAR, Panel B shows the identified set for the forecast variance ratio. Bootstrapped 90% confidence intervals are reported in brackets.

The results are shown in Table D.1. Carbon policy shocks have contributed meaningfully to historical variations in the variables of interest. Under the invertibility assumption (Panel A), they account for about 20 percent of the variations in energy prices and around 20 percent of the short-run variations in emissions, which goes up to roughly 40 percent at the 3-year horizon. Turning to the macroeconomic variables, we can see that they explain a substantial part of variations in the headline HICP, especially at shorter horizons, and a non-negligible fraction of the variations in industrial production and the unemployment rate at longer horizons. The shocks explain only little of the variations in the two-year rate, stock prices and oil prices. The forecast variance ratios in Panel B, which dispense from the assumption of invertibility, paint a similar picture.

D.4. Financial conditions and uncertainty

To better understand how the shock transmits to the economy, I study the responses of indicators for financing conditions and financial uncertainty, see Figure D.4. The responses turn out to be largely insignificant, suggesting that these variables do not appear to play a dominant role in the transmission of the carbon policy shock.

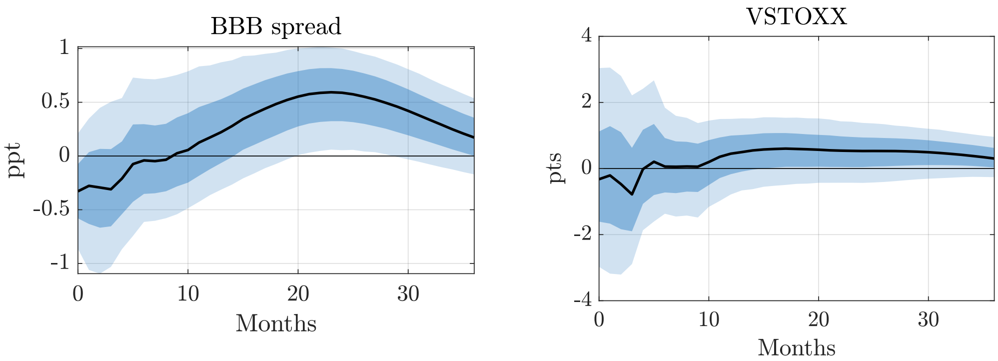

Figure D.4: Financial Conditions and Uncertainty

Notes: Impulse responses of financial conditions, as proxied by the BBB bond spread, and the VSTOXX index as a measure of financial uncertainty.

D.5. Green patenting

I perform two robustness checks on the patenting responses. First, to better account for potentially low-quality patents, I further restrict our sample to patents that have been cited more than once. Second, following Acemoglu et al. (2023) I exclude technologies that do not directly compete with fossil-fuel technologies, including those aimed at reducing pollution from fossil-fuel electricity generation (Y02E20), improving grid efficiency (Y02E40) or storage (Y02E60). The results turn out to be robust, see Figure D.5.

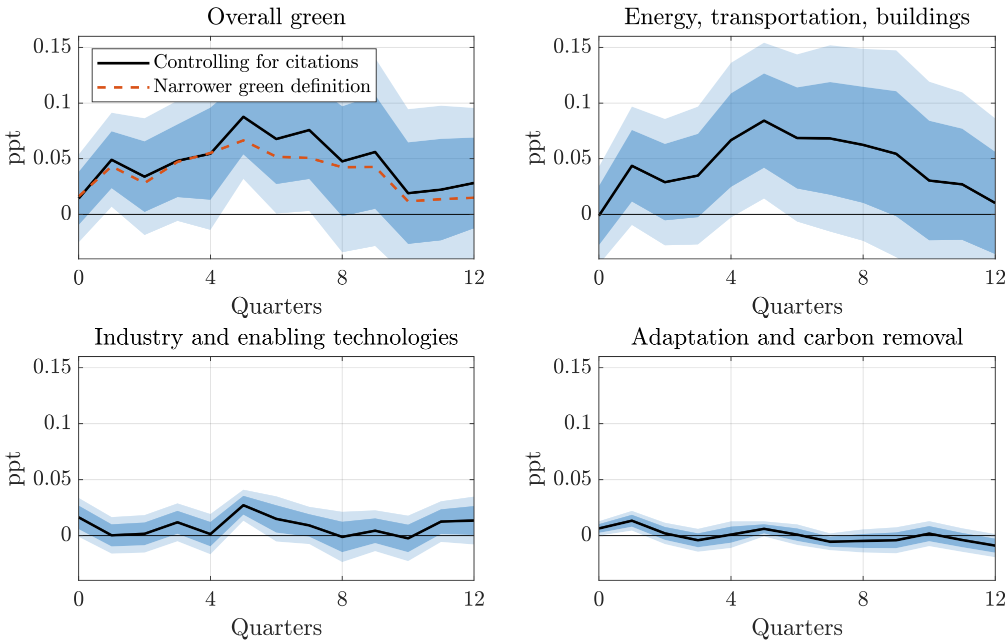

Figure D.5: Controlling for Patent Quality

Notes: Impulse responses of green patenting, as measured by the number of climate change mitigation patents as a share of all biadic patents file, imposing the additional restriction that patents have at least one citation. The figure reports responses for overall green patenting; mitigation technologies in energy generation, transportation, and buildings; mitigation in industry; and adaptation and carbon removal technologies. For overall green patenting, I also show results using a more narrow definition of green patents. Solid and dashed lines: point estimates. Dark and light shaded areas: 68 and 90 percent confidence bands based on lag-augmentation approach.

E. Additional Micro Results

In this Appendix, I present some additional results pertaining to Section 5 on the heterogeneous effects of carbon pricing in the paper.

E.1. Aggregate effects for the UK

Because of data availability, the household-level analysis is carried out for the UK. As a validating exercise, I have verified that the aggregate effects on the UK, as measured by real GDP, consumption and investment, are comparable to the EU level responses, though somewhat less persistent, see Figure E.1.

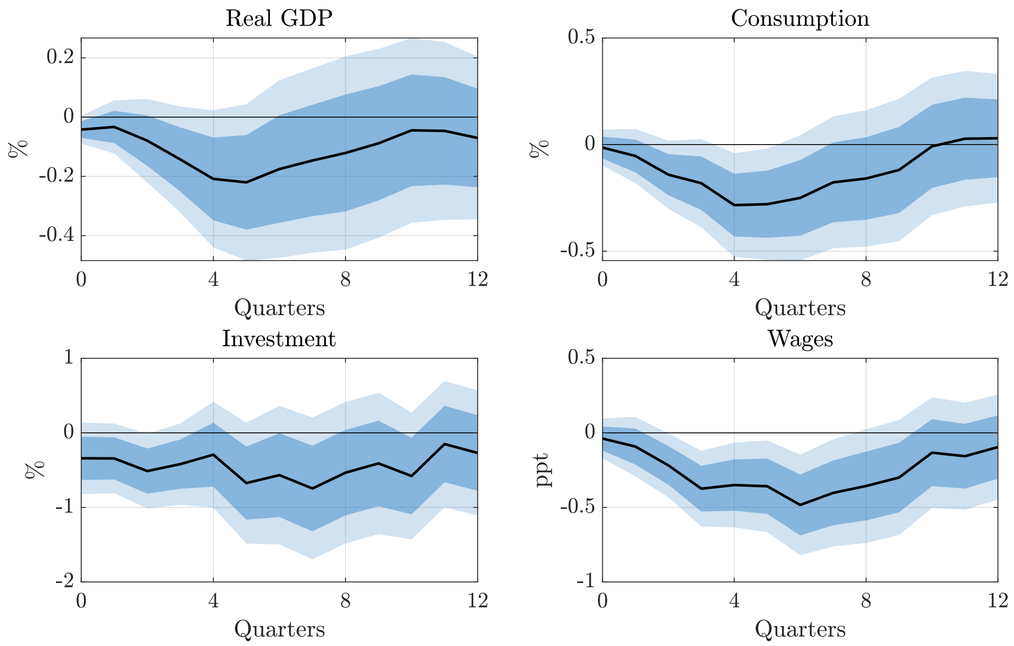

Figure E.1: Effect on GDP, Consumption, Investment and Wages in the UK

Notes: Impulse responses of a selection of quarterly variables estimated using local projections on the carbon policy shock. The responses are normalized to have the same peak effect on HICP energy as in the baseline model.

E.2. Additional descriptive statistics

Figure E.2 compares the empirical distribution of age and total expenditure for the three income groups. We can see that the groups are comparable in terms of their age distribution. As expected, higher income groups tend to have higher expenditure but there is also more within group variation.

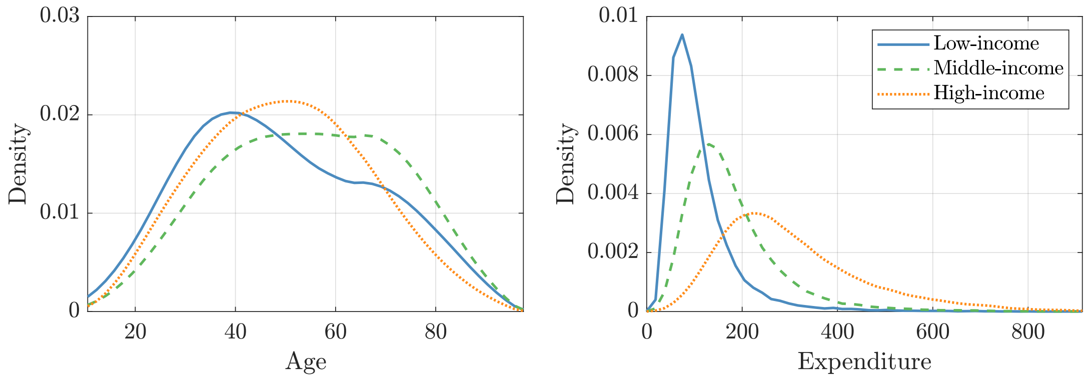

Figure E.2: Empirical Distribution of Age and Total Expenditure in the LCFS

Notes: Empirical probability distribution of age and total expenditure (excl. housing) for all three income groups. The distributions are estimated using an Epanechnikov kernel.

Figure E.3 depicts the evolution of different households characteristics, including age, education and housing tenure, over time. There are some trends in these variables, however, they are rather slow-moving and thus unlikely to confound potential heterogenities in the household responses to carbon policy shocks, which exploit variation at a much higher frequency.

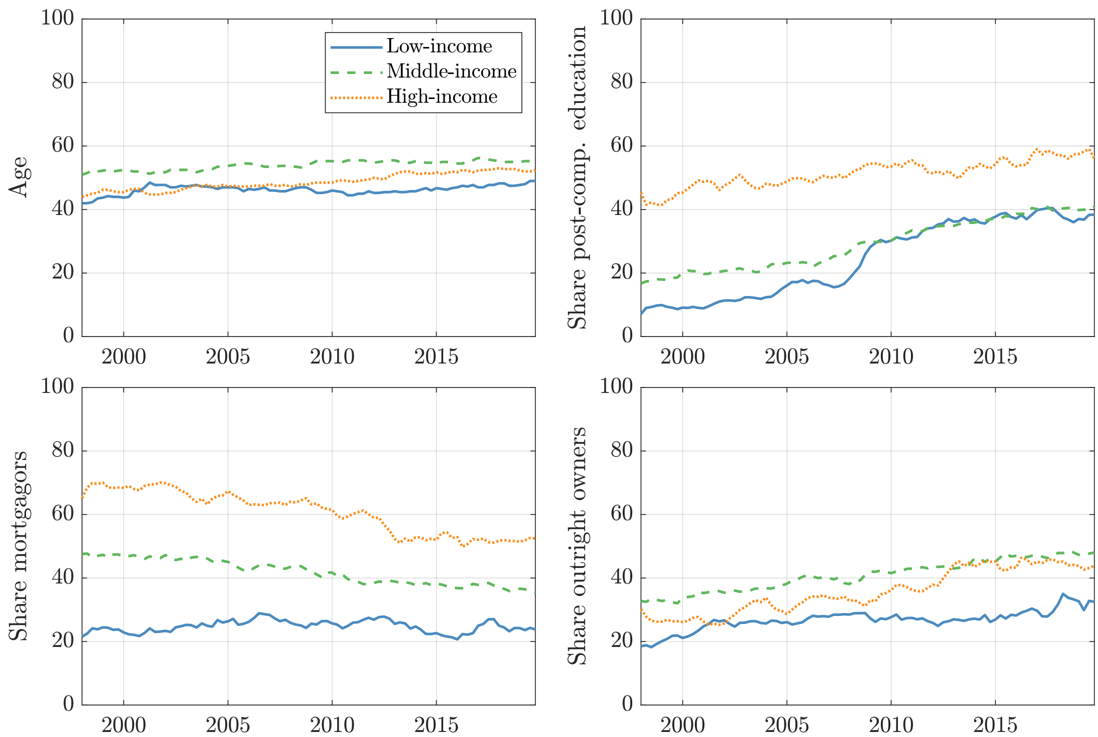

Figure E.3: Evolution of Household Characteristics by Income Group

Notes: Evolution of age, education, and housing tenure status over time by income group.

E.3. Aggregate responses

Before studying at the heterogeneous expenditure responses by income group, I look at the aggregate expenditure responses as a validating exercise. The results are shown in Figures E.4. The response of aggregated expenditure from household micro data is very similar to the consumption response from national statistics—both in terms of shape and magnitude. This supports the notion that the survey data is indeed representative for the macroeconomy. For completeness, I also report the aggregated responses for different expenditure categories.

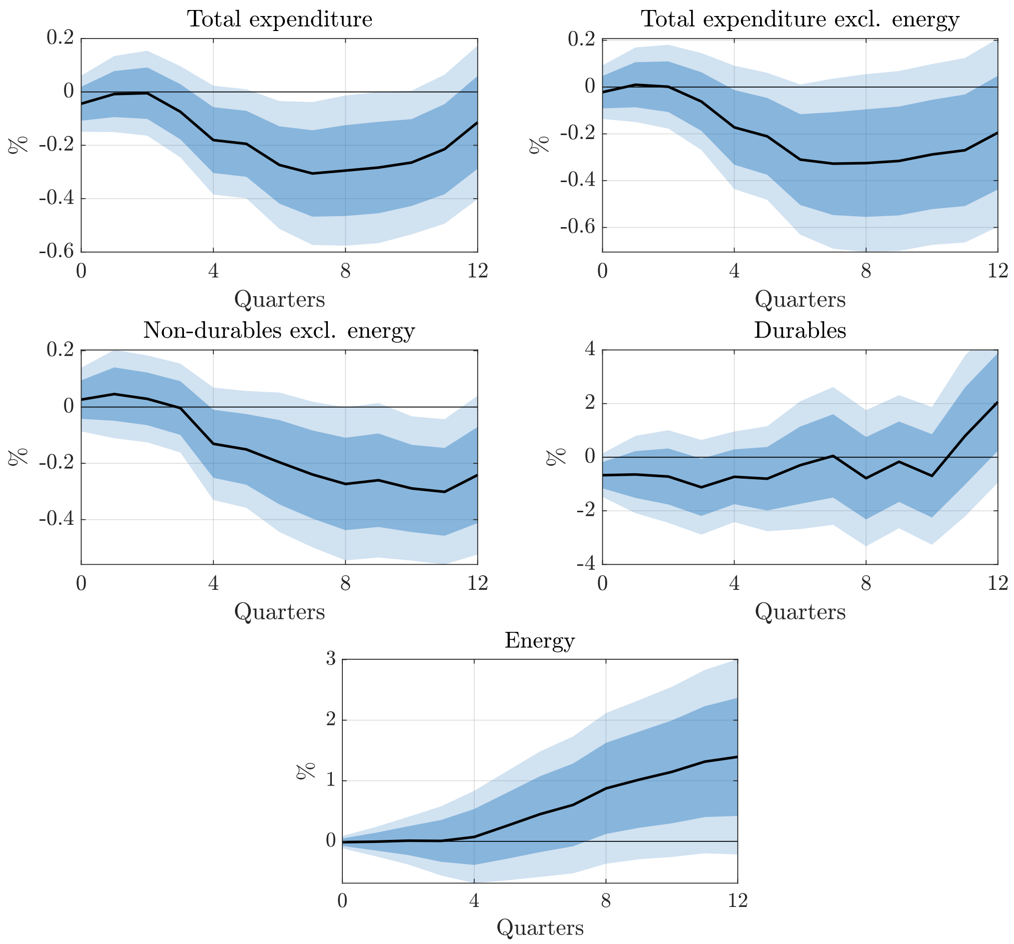

Figure E.4: Aggregate Expenditure Responses

E.4. Smoothing impulse responses

In the LCFS, households interviewed at time t are typically asked to report expenditure over the previous three months (with the exception of non-durable consumption which refers to the previous two weeks). To eliminate some of the noise inherent in survey data, I smooth the expenditure and income measures with a backward-looking (current and previous four quarters) moving average, as in Cloyne, Ferreira, and Surico (2020). However, as shown in Figure E.5, I obtain similar results when using the raw series instead, even though the responses become more jagged and imprecise, or by using smooth local projections as proposed by Barnichon and Brownlees (2019).

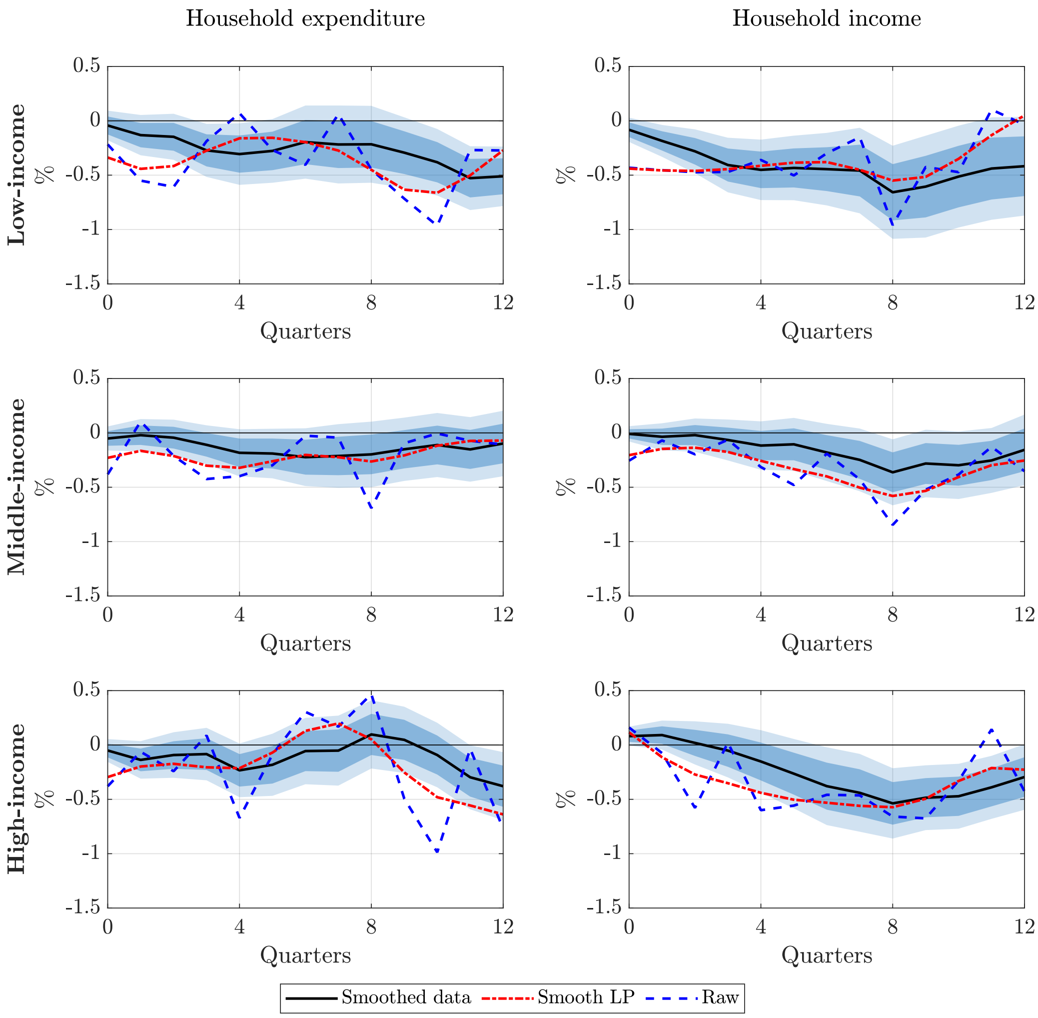

Figure E.5: Sensitivity with Respect to Smoothing of Responses

Notes: Impulse responses of total expenditure excluding housing and current total household income by income group, computed using simple backward-looking moving average (baseline), smooth local projections (red dotted line), and unsmoothed (blue dashed line).

E.5. Robustness concerning grouping

To mitigate concerns about endogenous changes in the grouping variable, I study the responses of current and normal disposable income in Figure E.6. We can see that both variables are rather slow-moving. Current income starts to fall significantly after about a year. In contrast, the response of normal disposable income moves less and is insignificant, supporting its validity as a grouping variable.

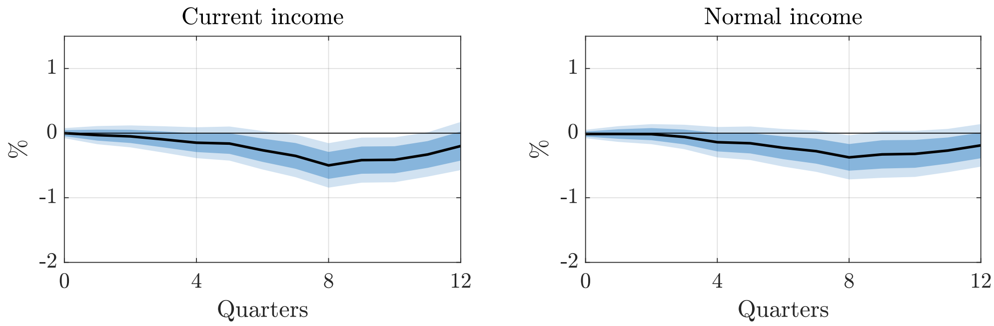

Figure E.6: Responses of Current and Normal Income

As discussed in the main text, the normal income variable can be thought of as a proxy for permanent income. As a robustness check, I compute estimates for permanent income from a Mincerian-type regressions. Specifically, I use age, education, ethnicity, sex, martial status, occupation, the source of the main household income, as well as interactions between age and education, and between age and sex as predictors, as in Alves et al. (2020).

Figure E.7 shows the responses by permanent income group. The results turn out to be robust.

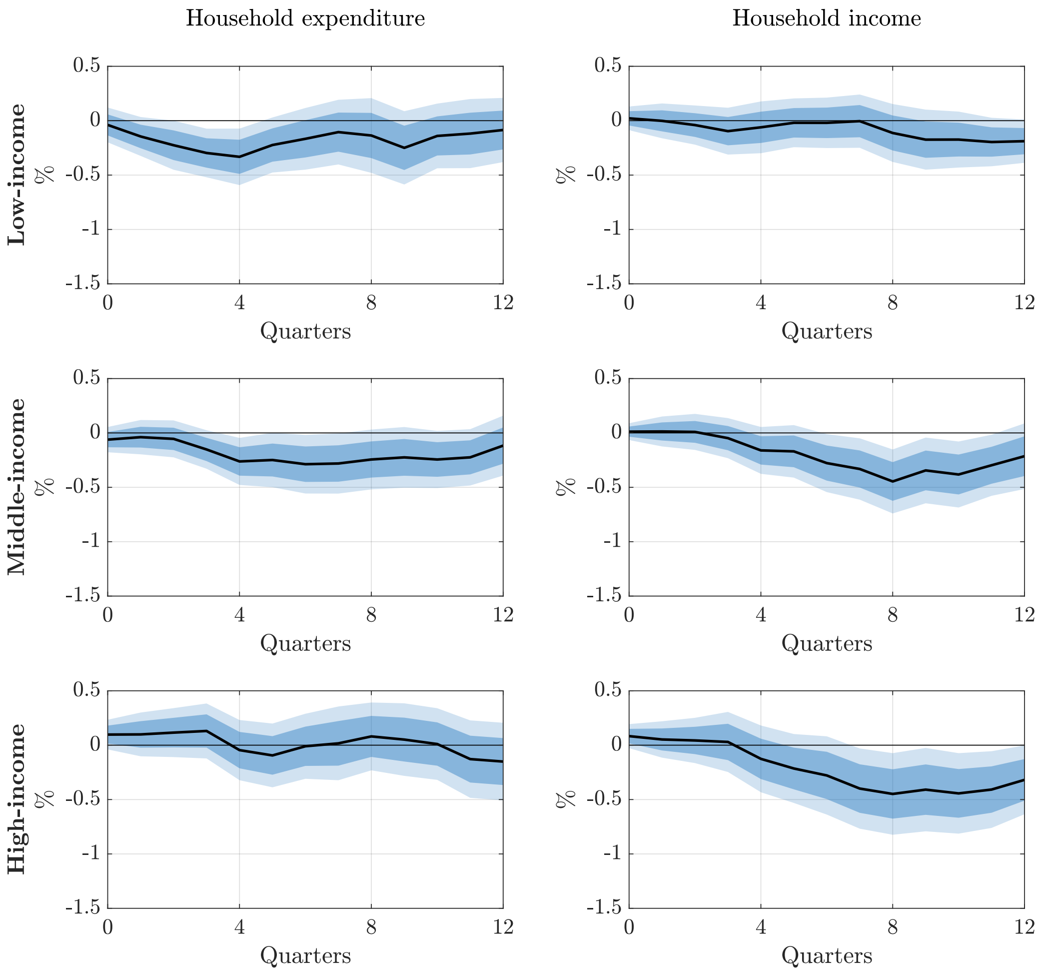

Figure E.7: Expenditure and Income Responses by Permanent Income

Notes: Impulse responses of total expenditure excluding housing and current total household income by permanent income, estimated using a Mincerian-type regression using age, education, ethnicity, sex, martial status, occupation, the source of the main household income, as well as interactions between age and education, and between age and sex (bottom 25 percent, middle 50 percent, top 25 percent).

E.6. Selection

To mitigate concerns about selection, I use a number of different grouping variables, including age, education and housing tenure. From Figures E.8-E.10, we can see that none of these alternative grouping variables can account for the patterns uncovered for income, suggesting that the estimates are not spuriously picking up differences in other household characteristics. Similarly, the uncovered heterogeneity can also not be accounted for by occupation, sex and region. These results are available from the author upon request.

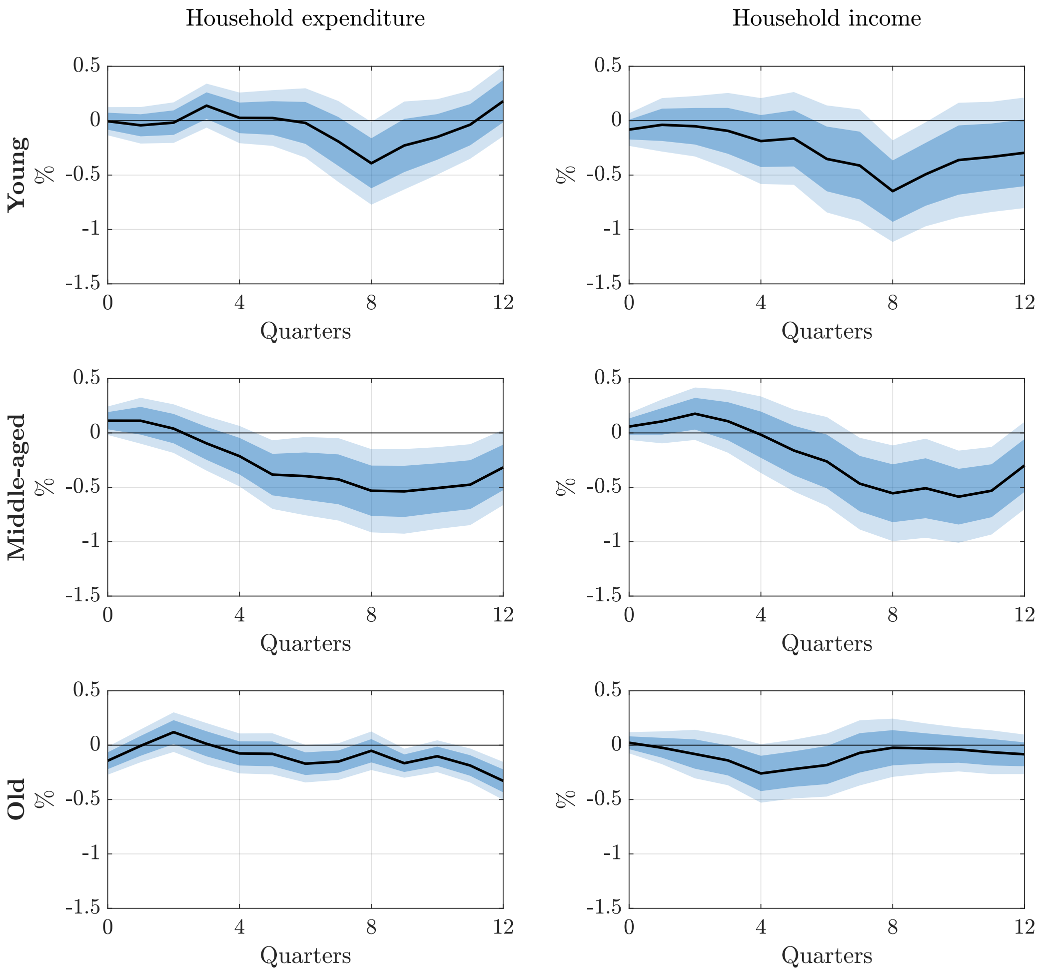

Figure E.8: Household Expenditure and Income Responses by Age Group

Notes: Impulse responses of total expenditure excluding housing and current total household income for young (bottom 33 percent), middle-aged (middle 33 percent) and older households (top 33 percent), based on the age of the household head.

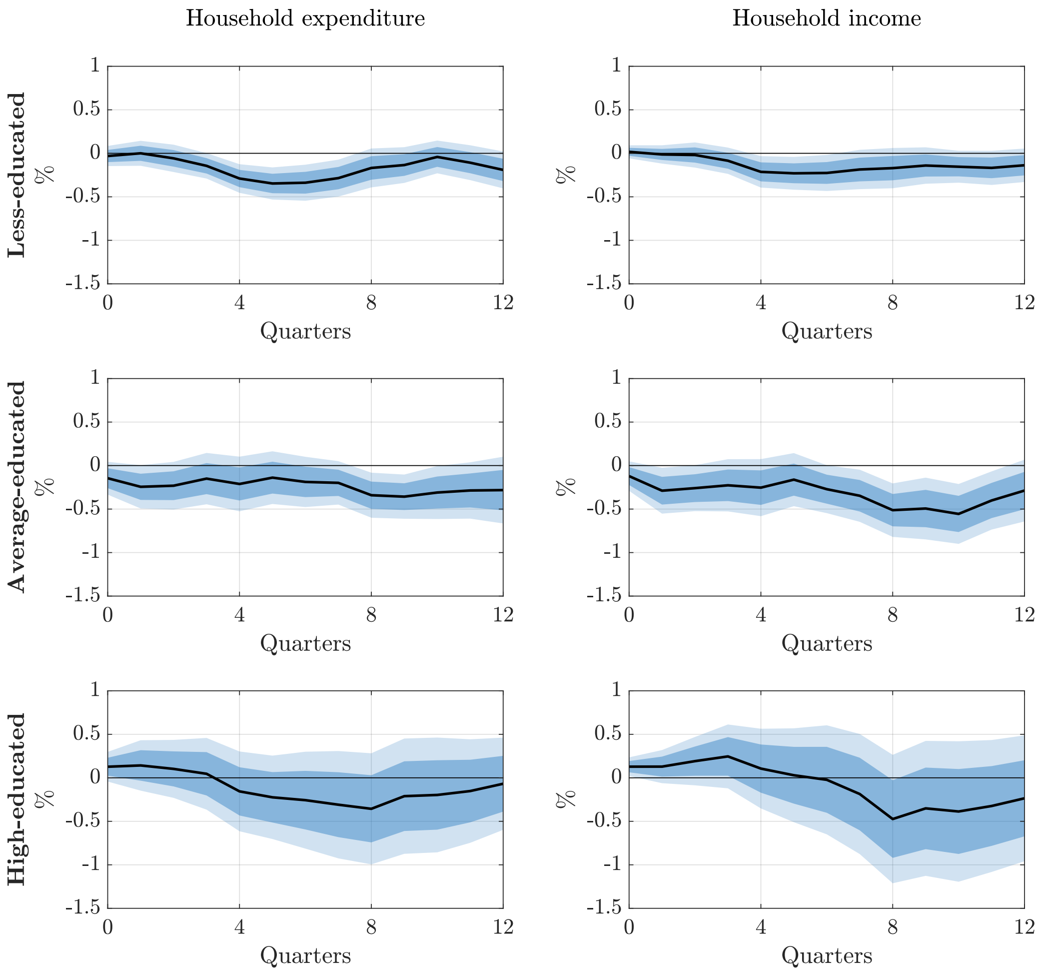

Figure E.9: Household Expenditure and Income Responses by Education Status

Notes: Impulse responses of total expenditure excluding housing and current total household income for less educated, normally educated and well educated households. Education status is proxied by the highest age a household member has completed full-time education and the three groups are below 16 years, between 17 and 18 years (compulsory education), and 19 years or above (post-compulsory).

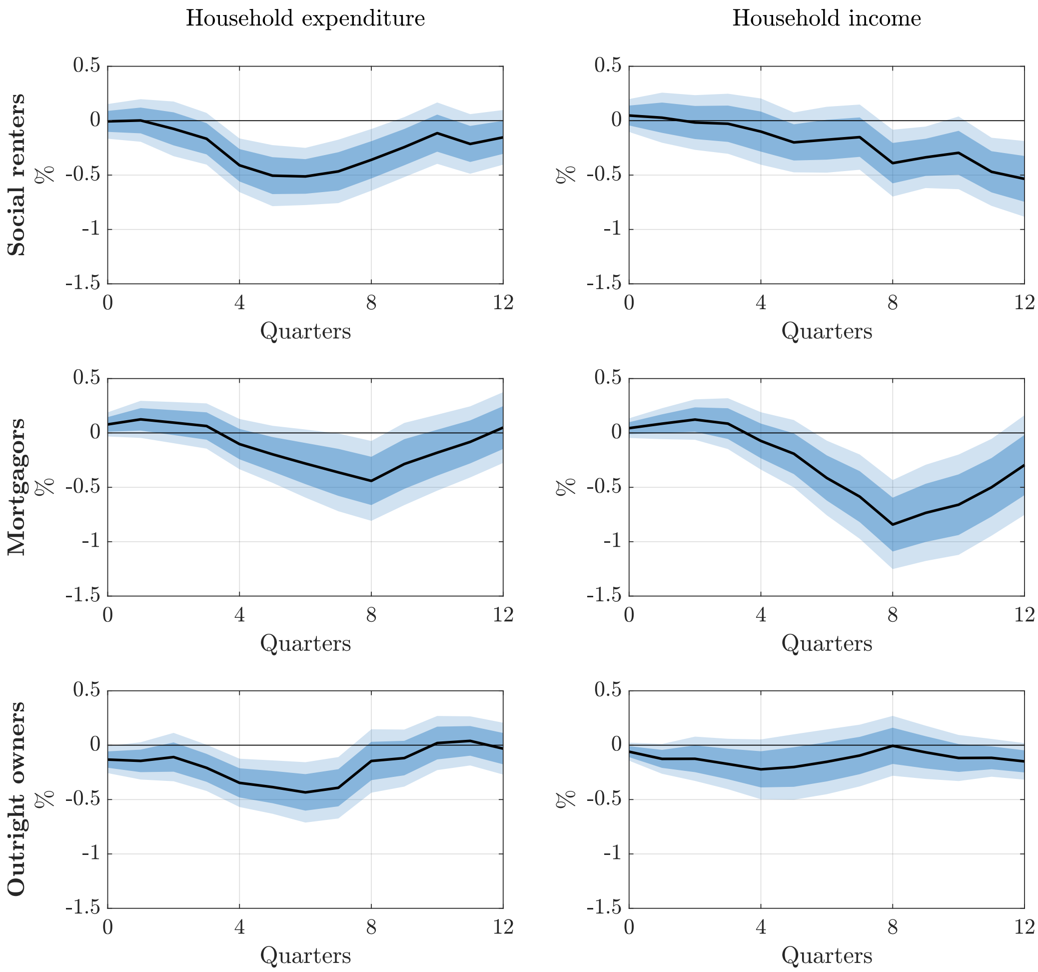

Figure E.10: Household Expenditure and Income Responses by Housing Tenure

Notes: Impulse responses of total expenditure excluding housing and current total household income for social renters, mortgagors and outright owners.

E.7. Expenditure responses

How do different components of household expenditure respond to carbon policy shocks? Figure E.11 shows the expenditure responses, broken down by energy, non-durables excluding energy and durable spending. These numbers are used to arrive at the cumulative figures in Table 3.

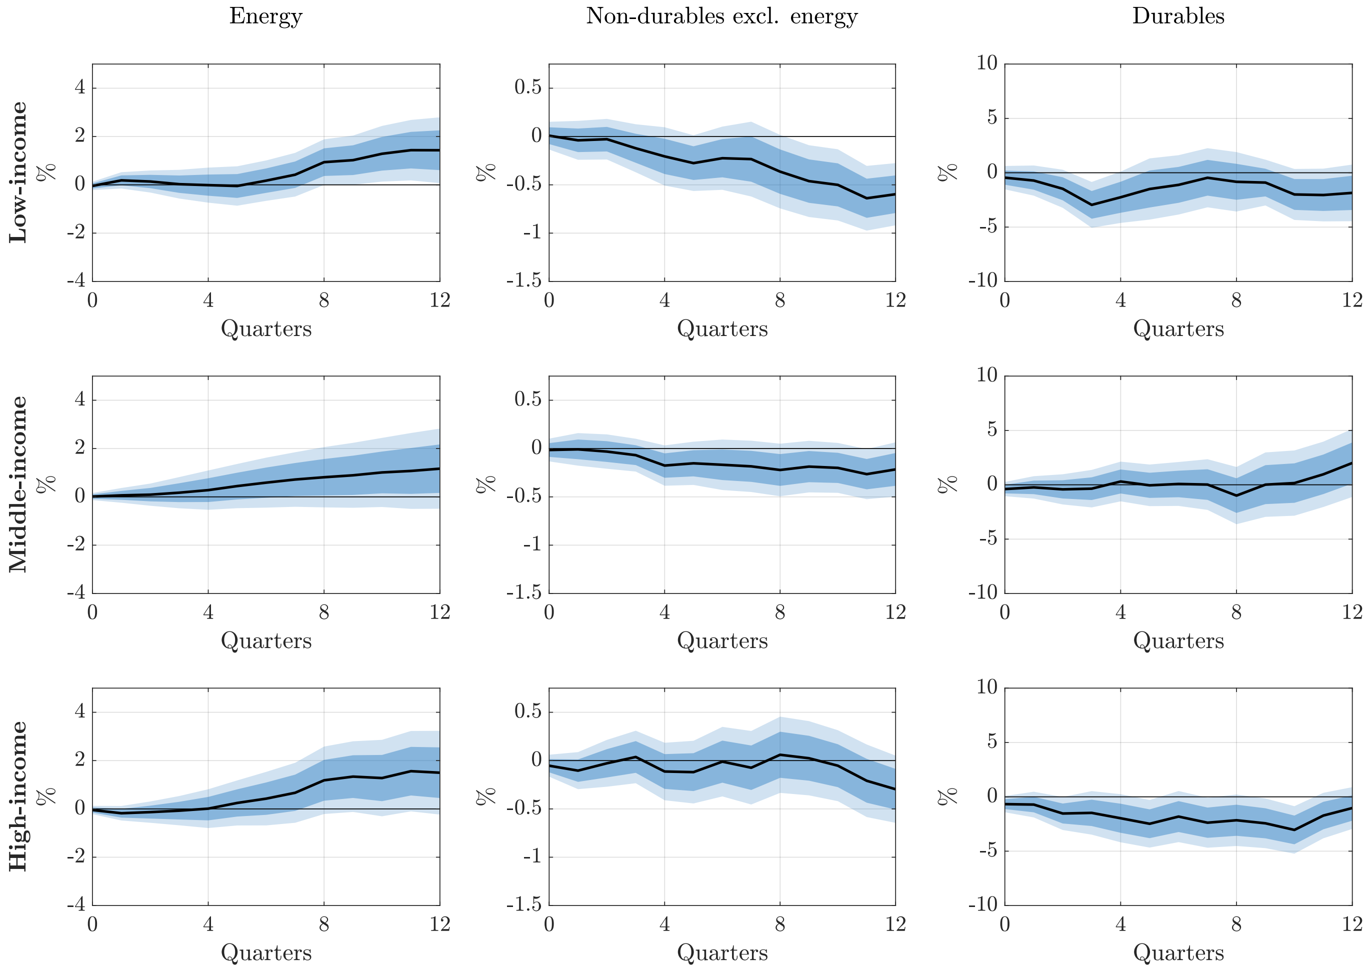

Figure E.11: Energy, Non-durables and Durables Expenditure Responses by Income Group

Notes: Impulse responses of energy, non-durables excluding energy and durables expenditure for low-income (bottom 25 percent), middle-income (middle 50 percent) and high-income households (top 25 percent). The households are grouped by total normal disposable income and the responses are computed based on the median of the respective group.

E.8. What drives the income response?

To understand what is driving the heterogeneity in the income responses, we study how the labor income responses vary by sector of employment using data from the LFS. To this end, I grouped sectors according to their SIC 2003 sections by their energy intensity and their "demand sensitivity", i.e. how much sectoral labor income changes after changes in aggregate income. The data on energy intensities is from the ONS. The demand sensitivity is proxied by the elasticity of sectoral labor income to aggregate labor income, using sectoral data from the LFS and wage data from national accounts. Similar results are obtained when estimating the elasticity with respect to the unemployment rate.

Table E.1 shows the data on sectoral energy intensity and estimated demand sensitivity together with the resulting classification. I define high energy intensive sectors as sectors with an energy intensity above 5 and high demand sensitive sectors as sectors with a demand sensitivity in excess of 0.5. Choosing the threshold involves some judgment. As a robustness check, I have excluded/included the sectors closest to the two thresholds for both groupings. The results turn out to be not sensitive to the precise level of the threshold.

Table E.1: Sectors by Energy Intensity and Demand Sensitivity

Panel A: Energy intensity and estimated demand sensitivity

Sectors		Energy intensity (TJ/£m)	Demand sensitivity ($\varepsilon_u y_i$)
A-B:	Agriculture, forestry and fishing	11.4	0.43
C,E:	Mining and quarrying; energy, gas and water	12.8	0.16
D:	Manufacturing	11.6	0.44
F:	Construction	2.6	0.52
G-H:	Wholesale and retail trade; hotels and restaurants	3.0	0.51
I:	Transport, storage and communication	9.4	0.19
J-K:	Banking, finance and insurance	0.7	0.41
L-N:	Public admin, education and health	1.3	0.35
O-Q:	Other services	1.1	0.72

Panel B: Sector classification

Group	Sectors	SIC sections
High energy intensity	Agriculture, forestry, and fishing; mining and quarrying; manufacturing; electricity, gas and water supply (utilities); transport, storage and communications	A-E, I
Lower energy intensity	Construction; Wholesale and retail trade; Hotels and restaurants; Financial intermediation; Real estate, renting and business; Public administration and defense; Education; Health and social work; Other community, social and personal services	F-H, J-Q
High demand sensitivity	Construction; Wholesale and retail trade; Hotels and restaurants; Other community, social and personal services	F-H, O-Q
Lower demand sensitivity	Agriculture, forestry, and fishing; mining and quarrying; manufacturing; electricity, gas and water supply (utilities); transport, storage and communications; Financial intermediation; Real estate, renting and business; Public administration and defense; Education; Health and social work	A-E, J-N

Notes: The sectors are grouped based on SIC 2003 sections. Note that the grouping is not perfect, as the LFS only has information on groups of sections over the entire sample of interest. The data on the energy intensity by sector from 1999-2019 is from the ONS.

Finally, another source of heterogeneity in the income response is the income composition. To better understand this, I study the responses of labor earnings and financial income. The earnings of low-income households fall more promptly than for higher-income households, consistent with the results on total income. On the other hand, the financial income of low- and middle-income households barely shows a response, reflecting the fact that these households own very little financial assets. In contrast, high-income households experience a temporary fall in their financial income in the short run, which however subsequently reverts (consistent with the stock market response).

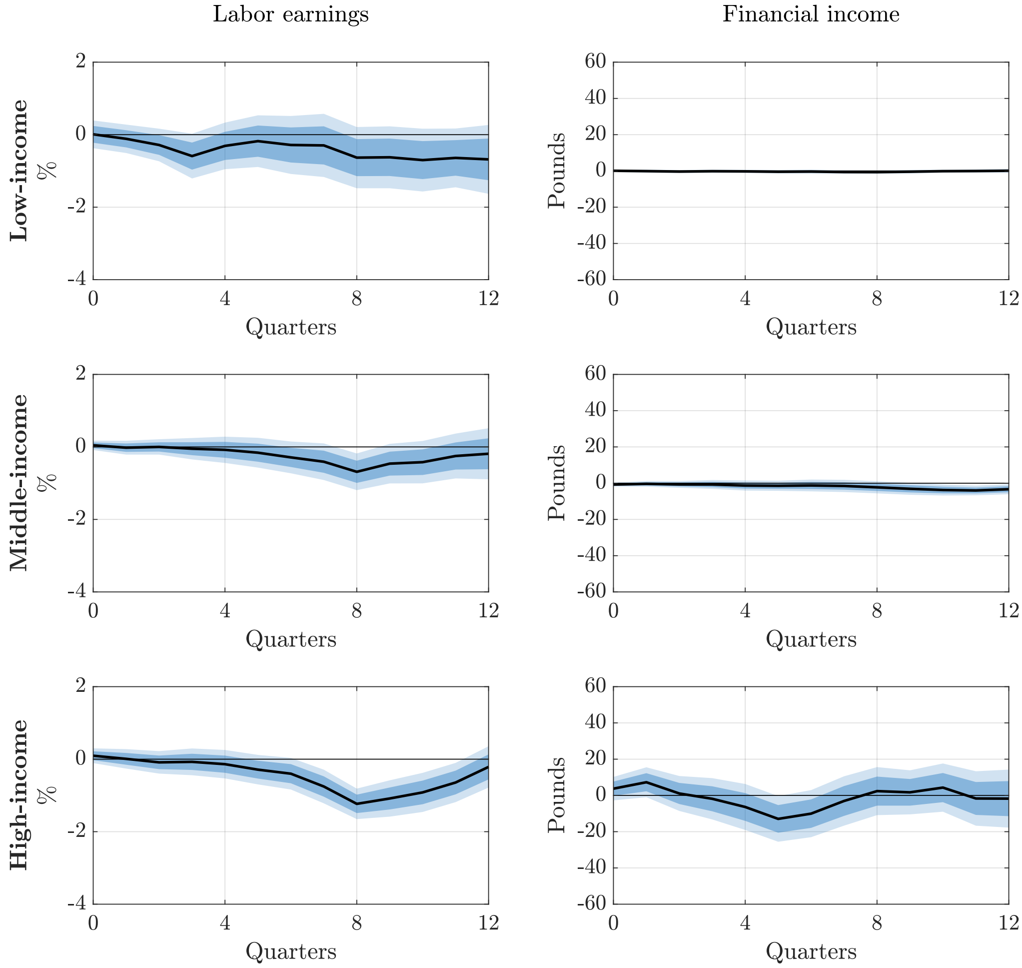

Figure E.12: Responses of Earnings and Financial Income

Notes: Impulse responses of labor earnings (wages from main occupation) and financial income (interest, dividend, rents) by income group (bottom 25 percent, middle 50 percent, top 25 percent).

E.9. External validity

To mitigate concerns regarding external validity, I confirm the main results on the heterogeneity in household expenditure by income group using data for Denmark and Spain. As can be seen from Figure E.13, the expenditure response turns out to be significant and persistent for low-income households, while high-income households are much less affected. These findings confirm the results for the UK, supporting the external validity of the results.

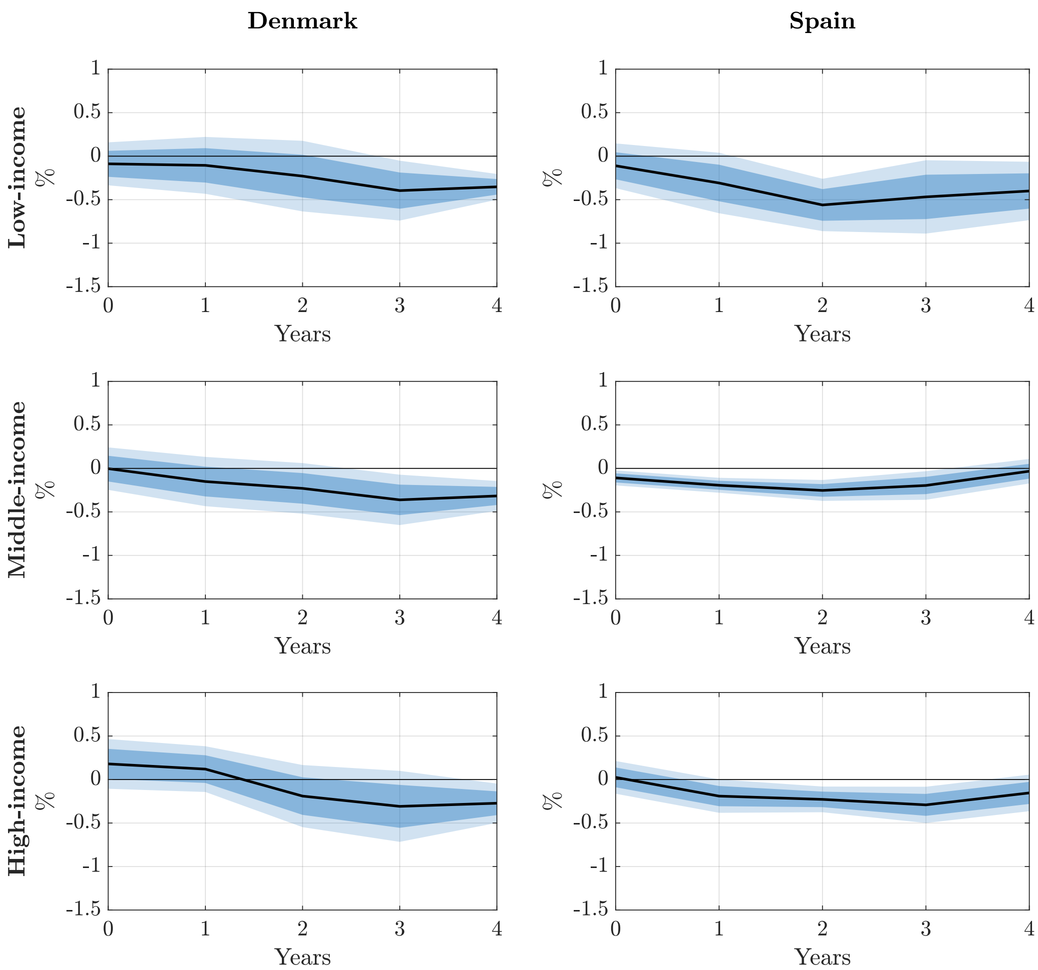

Figure E.13: Expenditure by Income Group for Other European Countries

Notes: Impulse responses of total expenditure for low-income, middle-income and high-income households in Denmark and Spain. The Danish data are from the Danish household budget survey (HBS) available for 1999-2019, accessed via the StatBank Denmark database, and expenditure is grouped by total annual income (under 250K DKK, 250-999K DKK, 1000K DKK or over). The Spanish data are from the Spanish HBS available for 2006-2019, accessed via the INE website, and expenditure is grouped by regular net monthly household income (under 1000 euros, 1000-2499 euros, 2500 euros or over).

E.10. Attitudes towards climate policy

As discussed in the paper, public opposition can be an impediment for climate policy. Thus, it is interesting to see how carbon pricing affects the public attitude towards climate policy. To analyze this question, I use data from the British social attitudes (BSA) survey. The BSA is an annual survey that asks about the attitudes of the British population towards a wide selection of topics, ranging from welfare to genomic science. The BSA is used to inform the development of public policy and is an important barometer of public attitudes. Some of the questions in the BSA are repeated over time and thus, it is possible to analyze how certain attitudes have changed over time.

To proxy the public attitude towards climate policy, I rely on a question from the transportation module of the survey, which asks about the attitude towards environmentally-motivated fuel taxes. In particular, the question asks whether the respondent agrees with the following statement: "For the sake of the environment, car users should pay higher taxes". The BSA also includes information about the income of the respondent, thus it is possible to analyze how the attitudes of different income groups have evolved. Figure E.14 shows how the attitude towards climate policy has changed among low- and higher-income households. The support of climate policy has remained relatively stable at moderate levels for a large part of the sample. In the early to middle 2010s, the support started increasing for higher-income households. In contrast, the support of low-income households has remained stable until the end of the sample.

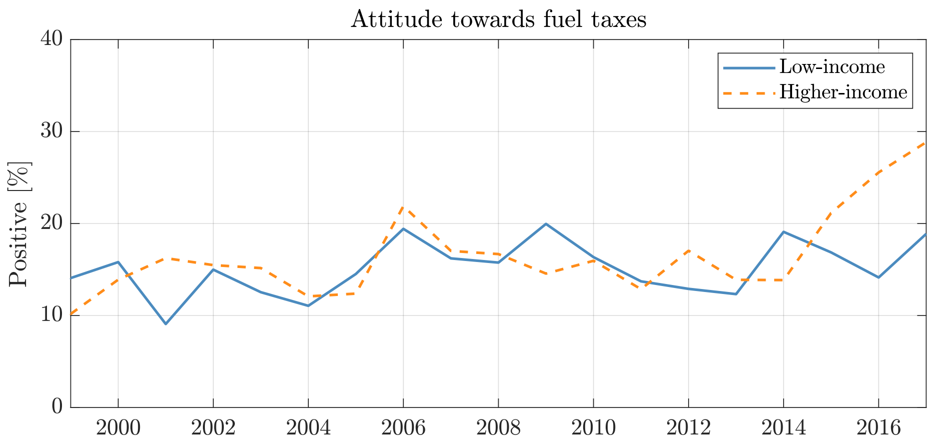

Figure E.14: Public support for climate policy by income group

Notes: Attitudes towards climate policy by income group over time, as proxied by the share of households in the British social attitudes survey that agree to the following statement: "For the sake of the environment, car users should pay higher taxes".

F. A Heterogeneous-agent Climate-economy Model

To study the role of redistributing carbon revenues more formally, I build a climate-economy model. The aim is to obtain a framework that can account for the empirical findings and can be used as a laboratory for policy experiments. The model belongs to the dynamic stochastic general equilibrium (DSGE) class. It augments the climate-economy structure in Golosov et al. (2014) with nominal rigidities and household heterogeneity, as in Bilbiie, Känzig, and Surico (2022), to allow for the demand channels identified in the data.

F.1. Households

The household sector consists of a continuum of infinitely lived households. Households have identical preferences and derive utility from consumption x and disutility from labor h. The consumption good is a composite of an energy and a non-energy good. To retain tractability, I consider a model with limited heterogeneity. There are two types of households: a share λ of households are *hand-to-mouth* (H) and a share $1-\lambda$ are *savers* (S) who choose their consumption intertemporally and save/invest in capital and risk-free bonds. Apart from the difference in MPC, households differ in their energy expenditure share and income incidence. Consistent with the data, I assume that the hand-to-mouth have a higher energy share and that their income is more elastic to changes in aggregate income than savers'.

Households face idiosyncratic risk as they switch exogenously between types. In particular, the exogenous change of type follows a Markov chain: the probability to stay a saver is s and the probability to remain hand-to-mouth is h (with transition probabilities $1-s$ and $1-h$, respectively). I focus on the stationary equilibrium with $\lambda = (1-s)/(2-s-h)$, which is the *unconditional* probability of being hand-to-mouth. I assume that only bonds are liquid and can be used to self-insure.

Savers. There is limited asset market participation. Only savers are able to self-insure themselves using liquid bonds.[6]

[6] This is a tractable way of introducing idiosyncratic risk and liquidity in spirit of full-blown HANK models à la Kaplan, Moll, and Violante (2018), see Bilbiie (2020) and Bilbiie, Känzig, and Surico (2022) for a detailed discussion.

Savers maximize their lifetime utility

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U(x_{S,t}, h_t) \right], \tag{16}$$

choosing how much to consume $x_{S,t}$, save $b_{S,t+1}$ and invest $i_{S,t}$. Their consumption bundle $x_{S,t}$ is a composite of a non-energy good $c_{S,t}$ and energy $e_{S,t}$: $x_{S,t} = \left(a_{S,c}^{\frac{1}{\epsilon_x}} c_{S,t}^{\frac{\epsilon_x - 1}{\epsilon_x}} + a_{S,e}^{\frac{1}{\epsilon_x}} e_{S,t}^{\frac{\epsilon_x - 1}{\epsilon_x}} \right)^{\frac{\epsilon_x}{\epsilon_x - 1}}$, where $a_{S,c}$ and $a_{S,e}$ are distribution parameters satisfying $a_{S,c} + a_{S,e} = 1$, and ϵ_x is the elasticity of substitution between non-energy and energy goods. This gives rise to standard non-energy and energy demand functions: $c_{S,t} = a_{S,c} \left(\frac{1}{p_{S,t}} \right)^{-\epsilon_x} x_{S,t}$ and $e_{S,t} = a_{S,e} \left(\frac{p_{e,t}}{p_{S,t}} \right)^{-\epsilon_x} x_{S,t}$.

The savers budget constraint equates their consumption, savings and investment to their income, accounting for the flows of liquid assets between types (see Bilbiie, Känzig, and Surico, 2022, for details). Their income is given by $y_{S,t} = w_t h_t + \frac{R_{t-1}^b}{\Pi_t} b_{S,t} + (1 - \tau^k) r_t k_{S,t} + \frac{(1-\tau^d) d_t}{1-\lambda} + \omega_{S,t}$, where $p_{S,t}$ is the price of the savers' final consumption bundle, $\frac{R_{t-1}^b}{\Pi_t}$ is the risk-free rate deflated by inflation, r_t is the rental rate of capital, d_t are dividends, and $\omega_{S,t}$ are transfers from the government. Capital accumulation follows $k_{S,t+1} = i_{S,t} + (1 - \delta) k_{S,t}$.

Maximizing (16) subject to the budget constraint and the capital accumulation equation, we obtain the Euler equations for investment and bond holdings:

$$\frac{U_x(x_{S,t}, h_{S,t})}{p_{S,t}} = \beta \, \mathbb{E}_t \left[\frac{R_t^b}{\Pi_{t+1}} \left(s \frac{U_x(x_{S,t+1}, h_{S,t+1})}{p_{S,t+1}} + (1 - s) \frac{U_x(x_{H,t+1}, h_{H,t+1})}{p_{H,t+1}} \right) \right] \tag{17}$$

$$\frac{U_x(x_{S,t}, h_{S,t})}{p_{S,t}} = \beta \, \mathbb{E}_t \left[(1 + (1 - \tau^k) r_{t+1} - \delta) \frac{U_x(x_{S,t+1}, h_{S,t+1})}{p_{S,t+1}} \right] \tag{18}$$

Note that only the Euler equation for bonds includes the marginal utility in the H-state, reflecting the fact that only bonds are liquid and can be used to self-insure against idiosyncratic risk.

Hand-to-mouth. Hand-to-mouth households have no assets and thus consume all of their income in every period:

$$p_{H,t} x_{H,t} = y_{H,t}. \tag{19}$$

The income of the hand-to-mouth is given by $y_{H,t} = w_t h_t^d + \omega_{H,t}$, where $\omega_{H,t}$ are government transfers. The non-energy and energy demand functions and the associated price index are analogous to the expressions for the savers.

Labor unions. Labor supply decisions are relegated to a labor union, which sets wages according to the following schedule:

$$w_t = \varphi h_t^{\theta} \left(\lambda \frac{1}{p_{H,t}} U_x(x_{H,t}, h_t) + (1-\lambda) \frac{1}{p_{S,t}} U_x(x_{S,t}, h_t) \right)^{-1}, \tag{20}$$

where w_t is the real wage charged by the union, $p_{H,t}$ and $p_{S,t}$ are the relative prices of the hand-to-mouth and the savers' consumption baskets, respectively, and $U_x(\cdot)$ is the marginal utility of consumption. The labor market structure equalizes labor income across households; thus all income heterogeneity in the model will come from heterogeneity in financial income.[7]

F.2. Firms

The firm block of the model consists of two sectors: energy and non-energy producers. Energy firms produce energy using labor as an input. Non-energy firms produce the non-energy consumption good using capital, energy, and labor as inputs. Consistent with the data, I assume that energy firms can adjust their prices flexibly while non-energy firms face nominal price rigidities (Dhyne et al., 2006).

Energy sector. The energy firm produces energy according to the following technology

$$e_t = a_{e,t} h_{e,t}, \tag{21}$$

as in Golosov et al. (2014). I assume that there is only a single source of energy (e.g. coal) that is available in approximately infinite supply. Without loss of generality, energy is measured in terms of carbon content (carbon amount emitted). Energy firms are subject to a carbon sales tax τ_t.[8] The optimal energy supply is characterized by $w_t = (1-\tau_t) p_{e,t} \frac{e_t}{h_{e,t}}$.

[7]This is a reduced-form way of capturing the income responses observed in the data. In the model, this labor market structure helps to mitigate varying labor supply responses offsetting income heterogeneity.

[8]For simplicity, I consider here a carbon tax, however, I could equivalently consider regulating the quantity (see e.g. the discussion in Heutel, 2012).

Non-energy sector. The non-energy sector consists of standard New Keynesian firms that produce different varieties of non-energy goods and set prices subject to nominal rigidities. The final non-energy good is assembled by a CES aggregator.

The non-energy variety j is produced according to the following technology, using capital $k_t(j)$, energy $e_{y,t}(j)$, and labor $h_{y,t}(j)$ as inputs

$$y_t(j) = e^{-\gamma s_t} \left[(1-\nu)^{\frac{1}{\epsilon_y}} \left(a_t k_t(j)^{\alpha} h_{y,t}(j)^{1-\alpha} \right)^{\frac{\epsilon_y - 1}{\epsilon_y}} + \nu^{\frac{1}{\epsilon_y}} \left(e_{y,t}(j) \right)^{\frac{\epsilon_y - 1}{\epsilon_y}} \right]^{\frac{\epsilon_y}{\epsilon_y - 1}}, \quad (22)$$

where a_t is a technology shifter. The function $e^{-\gamma s_t}$ captures climate damages, where s_t is the atmospheric carbon concentration. This generates a feedback loop between climate and the economy. Higher economic activity increases carbon emissions via higher energy use, which in turn increases the carbon concentration. A higher carbon concentration will have economic damages in turn (e.g. via weather events etc.), which reduce output. The functional form is taken from Golosov et al. (2014), where γ governs the size of climate damages.

The cost-minimization problem gives rise to the factor demands for capital $r_t = \alpha v_{1,t} mc_t \frac{y_t}{k_t}$, labor $w_t = (1-\alpha) v_{1,t} mc_t \frac{y_t}{h_{y,t}}$ and energy $p_{e,t} = v_{2,t} mc_t \frac{y_t}{e_{y,t}}$, where mc_t are real marginal costs and $v_{1,t} = \left(e^{-\gamma s_t} a_t k_t^{\alpha} h_{y,t}^{1-\alpha} / y_t \right)^{\frac{\epsilon_y - 1}{\epsilon_y}}$ and $v_{2,t} = \left(e^{-\gamma s_t} e_{y,t} / y_t \right)^{\frac{\epsilon_y - 1}{\epsilon_y}}$ are auxiliary terms. Note that factor demands are common across firms.

The price setting problem gives rise to a standard Phillips curve, which in log-linear form reads $\hat{\pi}_t = \kappa \hat{mc}_t + \beta E_t \hat{\pi}_{t+1}$, where hatted variables denote log-deviations from steady state. Finally, profits are given by $d_t = \int_0^1 [\frac{P_t(j)}{P_t} y_t(j) - mc_t y_t(j)] dj$.

F.3. Climate block

As in Golosov et al. (2014), the current level of atmospheric carbon concentration is a function of current and past emissions, $s_t = (1-\varphi) s_{t-1} + \varphi_0 e_t$, where φ_0 captures the share of emissions that do not immediately exit the atmosphere, and $1-\varphi$ measures how emission decay over time.

F.4. Fiscal and monetary policy

The government runs a balanced budget in every period, i.e. all transfers are financed by tax revenues. We consider the following transfer policy

$$\lambda\omega_{H,t} = \mu\tau_t p_{e,t} e_t \quad \text{and} \quad (1-\lambda)\omega_{S,t} = (1-\mu)\tau_t p_{e,t} e_t. \tag{23}$$

The distribution of carbon tax revenues are governed by the parameter μ. As the baseline, I assume that all carbon revenues accrue to the savers $\mu = 0$. Later, we will study alternative transfer policies. Carbon taxes τ_t are set according to the following rule: $\tau_t = (1-\rho_\tau)\tau + \rho_\tau \tau_{t-1} + \epsilon_{\tau,t}$. Finally, the monetary authority follows a standard Taylor rule, targeting headline inflation (in log-linear form): $\hat{r}^b_t = \rho_r \hat{r}^b_{t-1} + (1-\rho_r)(\phi_\pi \hat{\pi}_{T,t} + \phi_y \hat{y}_t) + \epsilon_{mp,t}$, where $\hat{\pi}_{T,t}$ is headline inflation.

F.5. Aggregation and market clearing

Because capital is only held by S, we have that $(1-\lambda)k_{S,t} = k_t$ and $(1-\lambda)i_{S,t} = i_t$. Because bonds are in zero net supply, we have $z_{S,t} = z_{H,t} = b_{S,t} = b_{H,t} = 0$.

Aggregate total, non-energy, and energy consumption are given by $x_t = \lambda x_{H,t} + (1-\lambda)x_{S,t}$, $c_t = \lambda c_{H,t} + (1-\lambda)c_{S,t}$, and $e_{c,t} = \lambda e_{H,t} + (1-\lambda)e_{S,t}$, respectively. Labor market clearing requires $\lambda h_{H,t} + (1-\lambda)h_{S,t} = h_{y,t} + h_{e,t}$. The energy market clears if $e_t = e_{c,t} + e_{y,t}$.

Finally, goods market clearing requires that $c_t + i_t = y_{d,t}$.

F.6. Calibration

I parameterize the model as follows. The time period is a quarter. The discount factor β takes the standard value 0.99, which implies an annualized steady-state interest rate of 4 percent. The intertemporal elasticity of substitution $1/\sigma$ and the labor supply elasticity $1/\theta$ are set to 1. These are standard values in the literature.

The labor weight in the utility function, φ_i is set such that steady-state hours worked h_i are normalized to one. I calibrate the share of hand-to-mouth λ to 25 percent, corresponding to the low-income threshold used in the LCFS. Such a share is also in line with the estimates of hand-to-mouth households in Kaplan, Violante, and Weidner (2014). Idiosyncratic risk is calibrated to $1-s = 0.04$, as in Bilbiie (2020). The distribution parameters $a_{H,e}$ and $a_{S,e}$ are set to match the energy expenditure shares of 9.5 percent for the hand-to-mouth and 6.5 percent for the savers as observed in the LCFS. Note that the elasticity of substitution ϵ_x is the same as the own price elasticity in this model. I calibrate ϵ_x to 0.2, as my empirical evidence points to rather low sustainability at the horizons considered.

The elasticity is consistent with Labandeira, Labeaga, and López-Otero (2017) who perform a meta analysis on the price elasticity of energy demand and find an average short-run elasticity of around 0.21.

Turning to the production side, I set the depreciation rate δ to 0.025, implying an annual depreciation on capital of 10 percent. The capital adjustment cost parameter is set to $\varphi_k = 4$, which implies an elasticity of investment to Tobin's marginal q of 10. I set α to 0.3, implying a steady-state capital share of around 70 percent (see e.g. Smets and Wouters, 2003). Using data on non-household energy consumption and energy prices in the EU, I estimate a energy share of around 7 percent. To approximate that share, I thus set $\nu = 0.07$. The elasticity of substitution between energy and capital/labor is set to 0.21, drawing again on the evidence in Labandeira, Labeaga, and López-Otero (2017). The elasticity of substitution between non-energy varieties is assumed to be 6, which is a standard value and implies a steady-state markup of 20 percent, consistent with the evidence in Christopoulou and Vermeulen (2012). The Calvo parameter θ_p is set to 0.825, which implies an average price duration of 5-6 quarters, in line with the empirical estimates in Alvarez et al. (2006). These parameter choices imply a relatively flat Phillips curve with a slope of 0.04.

For the climate block, I rely on the values in Golosov et al. (2014). I abstract from uncertainty about the damage parameter and use the deterministic, long-run value from Golosov et al. (2014). Note, however, that carbon emissions in my model are in arbitrary units. Thus, following Heutel (2012) I scale the damage parameter to make the increase in output damages from doubling the steady-state carbon stock consistent with the projected increase in damages from doubling CO2 levels in 2005. Turning to the carbon cycle, note that the excess carbon has a half-life of about 300 years (Archer, 2005). This implies a value of $1 - \varphi = 0.9994$.[9] Furthermore, according to the 2007 IPCC reports, about half of the CO2 pulse to the atmosphere is removed after a time scale of 30 years. This implies that $\varphi_0 = \frac{0.5}{(1-\varphi)^{120}} = 0.5359$.

Turning to fiscal and monetary policy, I compute the steady-state carbon tax as the implied tax rate implied by the average EUA price which is around 3.9 percent (the average real EUA price as a share of gross electricity prices in emission units). The persistence of the tax shock is set to 0.85, which is broadly consistent with the shock persistence estimated in the external instruments VAR. Finally, the Taylor rule coefficient on inflation is set to 1.5, and interest smoothing is assumed to be 0.8. These values are standard in the literature.

[9] From the carbon cycle, we have $E_t s_{t+h} = (1 - \varphi)^h s_t = 0.5 s_t$. Thus, I impose $(1 - \varphi)^{1200} = 0.5$ to get φ.

All other taxes are assumed to be zero in the baseline case, later I will use them to equalize the income incidence. Furthermore, I assume that all carbon tax revenues accrue to the savers, $\mu = 0$, motivated by the fact that there is no redistribution scheme in the current EU ETS in place. The calibration is summarized in Table F.1.

Table F.1: Calibration

Parameter	Description	Value	Target/Source/Comments
β	Discount factor	0.99	Standard value
$1/\sigma$	Intertemporal elasticity of substitution	1	Standard value
$1/\theta$	Labor supply elasticity	1	Standard value
λ	Share of hand-to-mouth	0.2	Share of low-income households, LCFS
$1-s$	Probability of becoming H	0.04	Bilbiie (2020)
$a_{H,e}$	Distribution parameter H	0.078	Energy share of 9.5%, LCFS
$a_{S,e}$	Distribution parameter S	0.056	Energy share of 6.5%, LCFS
ϵ_x	Elasticity of substitution energy/non-energy households	0.2	LCFS, Labandeira, Labeaga, and López-Otero (2017)
ϵ_y	Elasticity of substitution energy/non-energy firms	0.21	Labandeira, Labeaga, and López-Otero (2017)
δ	Depreciation rate	0.025	Standard value
φ_k	Capital adjustment costs	4	Standard value
α	Capital returns-to-scale	0.3	Standard value
ν	Energy returns-to-scale	0.07	Steady-state energy share of $\approx$ 7%; Eurostat
ϵ_p	Price elasticity	6	Steady-state markup of 20%
θ_p	Calvo parameter	0.825	Average price duration of 5-6 quarters
γ	Climate damage parameter	$5.3 * 10^{-5}$	Golosov et al. (2014)
φ_0	Emissions staying in atmosphere	0.5359	Golosov et al. (2014)
$1-\varphi$	Emissions decay parameter	0.9994	Golosov et al. (2014)
ϕ_π	Taylor rule coefficient inflation	1.5	Standard value
ρ_r	Interest smoothing	0.8	Standard value
τ	Steady-state carbon tax	0.039	Implied tax rate from average EUA price
ρ_τ	Persistence carbon tax shock	0.85	Persistence in the data

F.7. Results

Model evaluation. The impulse responses to a carbon policy shock, leading to an increase in the energy price by 1 percent, are shown in Figure F.1. In what follows, we focus on the peak responses, as the model is not designed to match the hump-shaped responses in the data. The model is successful in generating consumption and income responses, overall and by household group, that are in the same order of magnitude as the estimated responses in Section 5. As in the data, consumption and income are more responsive to carbon policy shocks for the low-income, hand-to-mouth households. In contrast, the responses of high-income savers are much less pronounced.

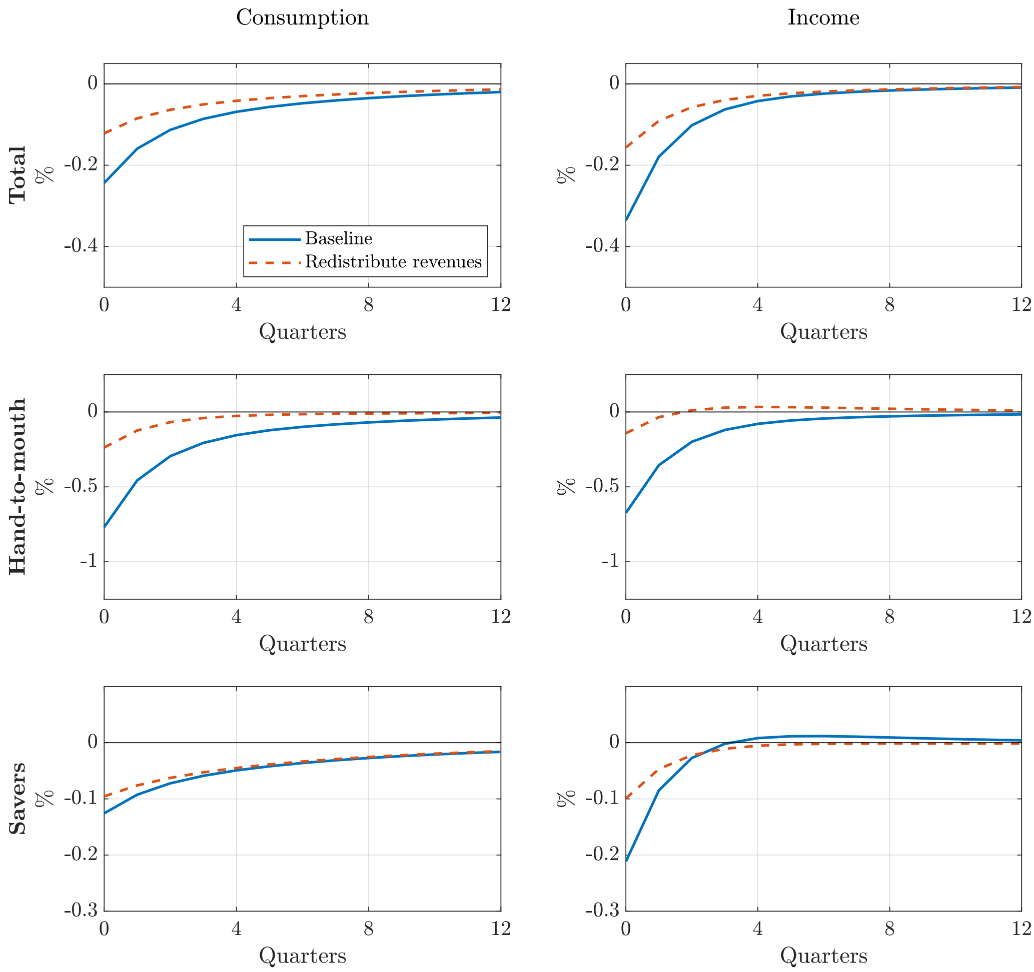

Figure F.1: Model Responses for Consumption and Income

Notes: Impulse responses of consumption and income, in the aggregate as well as for hand-to-mouth and savers, to a carbon policy shock normalized to increase the energy price by 1 percent. The blue line is the baseline response when carbon revenues solely accrue to the savers; the red dashed line is the response when carbon revenues are redistributed equally among hand-to-mouth and savers.

Redistributing carbon revenues. We are now in a position to study how different carbon revenue redistribution schemes affect the transmission of carbon policy shocks. Figure F.1 compares the baseline case when all carbon revenues accrue to the savers (blue line) to the case where the revenues are distributed equally across households $\mu = \lambda$ (red dashed line).

Redistributing carbon revenues has important consequences: the aggregate effect on consumption and income is much smaller than in the baseline case of no redistribution. In contrast, redistributing revenues has a smaller impact on the response of emissions, see Figure F.2. The intuition is that the redistribution

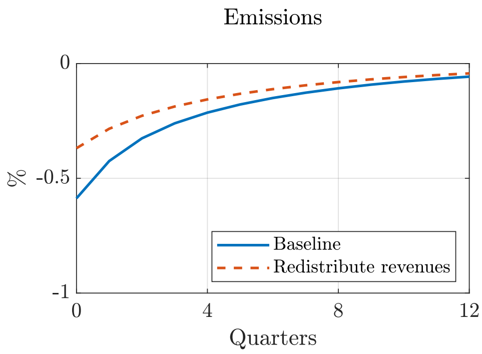

Figure F.2: Emissions Response

scheme stabilizes the income of the hand-to-mouth which translates into a significantly smaller consumption response as they have a high MPC. Savers, on the other hand, face a somewhat more prolonged fall in their income but the effect on their consumption is more muted as they are able to smooth the effects of the shock. Thus, redistributing carbon revenues also leads to a reduction in consumption inequality. Emissions on the other hand change by less as low-income households' energy demand is inelastic and they make up only a small share of aggregate emissions to start with.

The above findings speak directly to the recent debate on carbon pricing and inequality in Europe. The model confirms the intuition that redistributing carbon revenues could mitigate the effect on aggregate consumption and alleviate the distributional impact without compromising emission reductions to a significant extent. An interesting case in point in this context is the carbon tax in British Columbia. Contrary to the EU ETS, the tax was introduced alongside substantial reductions in income taxes and direct subsidies to the most affected households. The existing empirical evidence finds that the tax also reduced emissions significantly but the effects on economic activity turn out to be smaller (see Metcalf, 2019; Bernard and Kichian, 2021)—consistent with the predictions of my model.

Role of heterogeneity. Household heterogeneity plays an important role for the magnitudes of the responses. In particular, heterogeneity in MPCs linked to heterogeneity in energy shares and income incidence can amplify the responses further. This is illustrated in Figure F.3, which compares the responses of the heterogeneous agent to the corresponding representative agent version of the model.

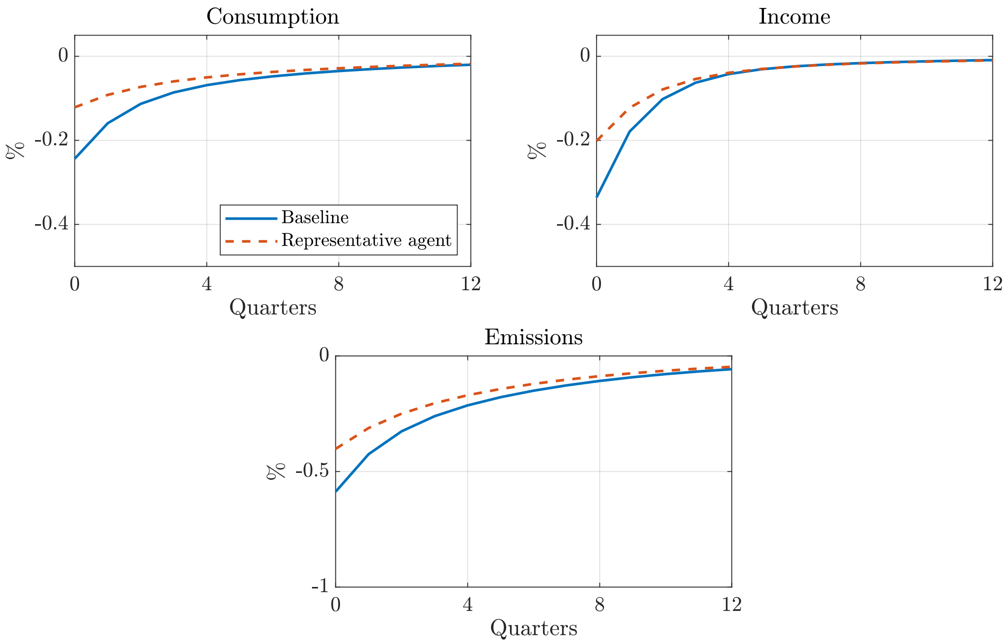

Figure F.3: Heterogeneity Versus Representative Agent

To get a better understanding of how much the heterogeneity matters for the direct and indirect channels I identify, I perform a decomposition. In particular, I compare four different scenarios: (i) a model where there is no heterogeneity in income incidence and energy share (this is achieved by perfectly redistributing income over the cycle and calibrating the energy share for H and S to the same level), (ii) a model with equal incidence but heterogeneity in energy shares, (iii) a model with unequal incidence and no energy share heterogeneity, and (iv) our baseline case with both heterogeneities. From Figure F.4, we can see that the heterogeneity in income incidence turns out to be crucial, accounting for the bulk of the amplification of the aggregate consumption response. This can be seen from the fact that the model with unequal incidence is already very close to the baseline with heterogeneous energy shares and income incidence.

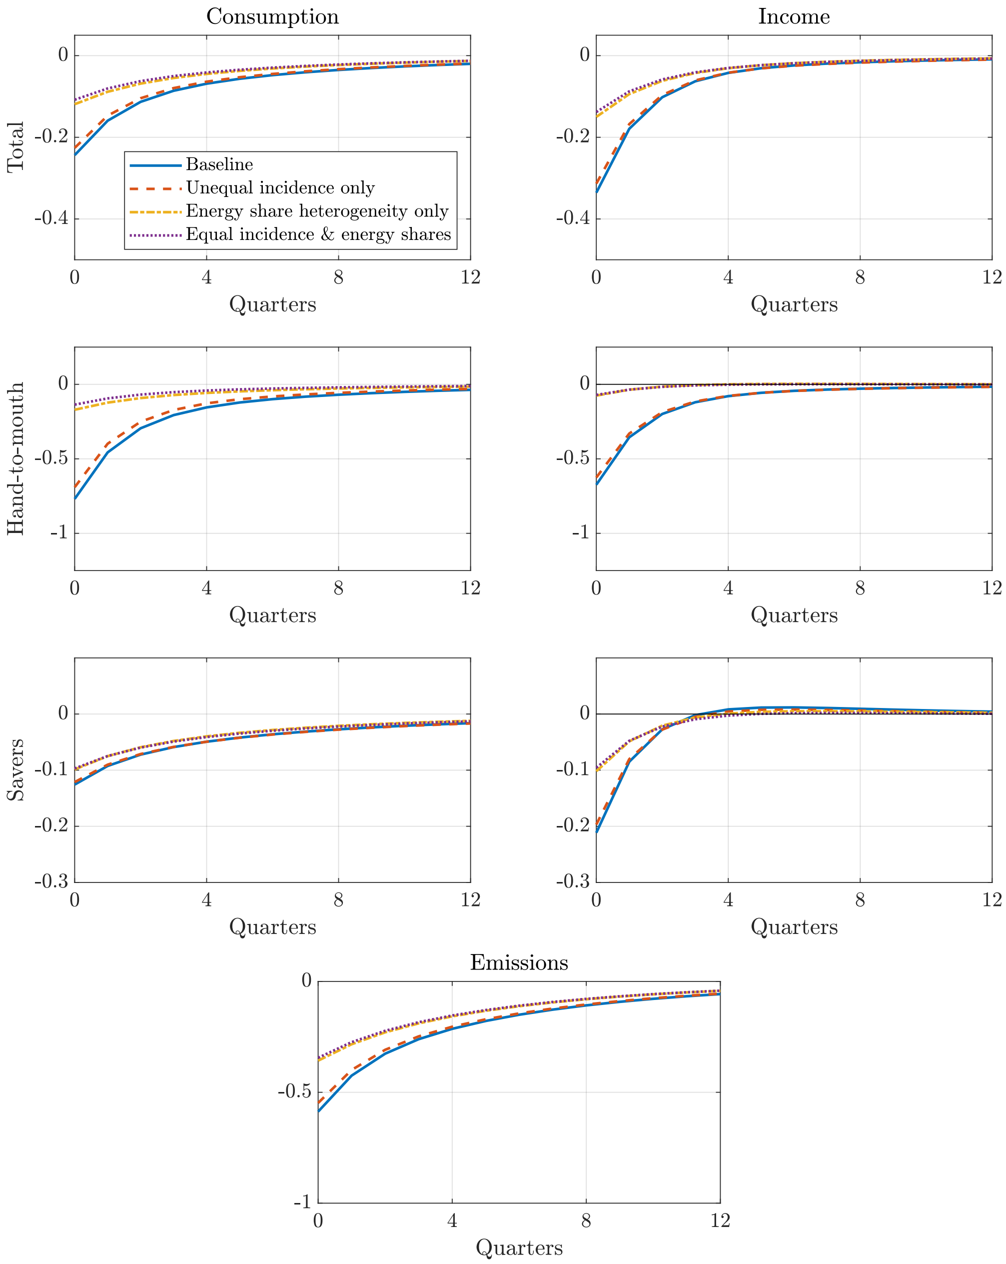

Figure F.4: Role of Unequal Income Incidence and Energy Share Heterogeneity

References Appendix

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hémous. 2023. "Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution." National Bureau of Economic Research.

Alloza, Mario, Pablo Burriel, and Javier J. Pérez. 2019. "Fiscal policies in the euro area: Revisiting the size of spillovers." *Journal of Macroeconomics*, 61: 103132.

Almgren, Mattias, José-Elías Gallegos, John Kramer, and Ricardo Lima. 2022. "Monetary policy and liquidity constraints: Evidence from the euro area." *American Economic Journal: Macroeconomics*, 14(4): 309–340.

Alvarez, Luis J., Emmanuel Dhyne, Marco Hoeberichts, Claudia Kwapil, Hervé Le Bihan, Patrick Lünnemann, Fernando Martins, Roberto Sabbatini, Harald Stahl, and Philip Vermeulen. 2006. "Sticky prices in the euro area: a summary of new micro-evidence." *Journal of the European Economic Association*, 4(2-3): 575–584.

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. 2020. "A further look at the propagation of monetary policy shocks in HANK." *Journal of Money, Credit and Banking*, 52(S2): 521–559.

Archer, David. 2005. "Fate of fossil fuel CO2 in geologic time." *Journal of geophysical research: Oceans*, 110(C9).

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring economic policy uncertainty." *The Quarterly Journal of Economics*, 131(4): 1593–1636.

Barnichon, Regis and Christian Brownlees. 2019. "Impulse response estimation by smooth local projections." *Review of Economics and Statistics*, 101(3): 522–530.

Baumeister, Christiane and James D. Hamilton. 2019. "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks." *American Economic Review*, 109(5): 1873–1910.

Bernard, Jean-Thomas and Maral Kichian. 2021. "The Impact of a Revenue-Neutral Carbon Tax on GDP Dynamics: The Case of British Columbia." *The Energy Journal*, 42(3).

Bilbiie, Florin O. 2020. "Monetary Policy and Heterogeneity: An Analytical Framework."

Bilbiie, Florin O, Diego R Känzig, and Paolo Surico. 2022. "Capital and income inequality: An aggregate-demand complementarity." *Journal of Monetary Economics*, 126: 154–169.

Bloom, Nicholas. 2009. "The impact of uncertainty shocks." *Econometrica*, 77(3): 623–685.

Botero, Juan C., Simeon Djankov, Rafael La Porta, Florencio Lopez-de Silanes, and Andrei Shleifer. 2004. "The regulation of labor." *The Quarterly Journal of Economics*, 119(4): 1339–1382.

Caldara, Dario, Michele Cavallo, and Matteo Iacoviello. 2019. "Oil price elasticities and oil price fluctuations." *Journal of Monetary Economics*, 103: 1–20.

Christopoulou, Rebekka and Philip Vermeulen. 2012. "Markups in the Euro area and the US over the period 1981–2004: a comparison of 50 sectors." *Empirical Economics*, 42(1): 53–77.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary policy when households have debt: new evidence on the transmission mechanism." *The Review of Economic Studies*, 87(1): 102–129.

Dhyne, Emmanuel, Luis J. Alvarez, Hervé Le Bihan, Giovanni Veronese, Daniel Dias, Johannes Hoffmann, Nicole Jonker, Patrick Lunnemann, Fabio Rumler, and Jouko Vilmunen. 2006. "Price changes in the euro area and the United States: Some facts from individual consumer price data." *Journal of Economic Perspectives*, 20(2): 171–192.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. "Optimal taxes on fossil fuel in general equilibrium." *Econometrica*, 82(1): 41–88.

Haščič, Ivan and Mauro Migotto. 2015. "Measuring environmental innovation using patent data."

Hémous, David, Morten Olsen, Carlo Zanella, and Antoine Dechezleprêtre. 2025. "Induced Automation Innovation: Evidence from Firm-level Patent Data."

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

Jarociński, Marek and Peter Karadi. 2020. "Deconstructing monetary policy surprises—the role of information shocks." *American Economic Journal: Macroeconomics*, 12(2): 1–43.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American Economic Review*, 95(1): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2015. "Betting the house." *Journal of International Economics*, 96: S2–S18.

Känzig, Diego R. and Maximilian Konradt. 2024. "Climate Policy and the Economy: Evidence from Europe's Carbon Pricing Initiatives." *IMF Economic Review*, 72(3): 1081–1124.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary policy according to HANK." *American Economic Review*, 108(3): 697–743.

Kaplan, Greg, Giovanni L. Violante, and Justin Weidner. 2014. "The Wealthy Hand-to-Mouth." *Brookings Papers on Economic Activity*, 2014(1): 77–138.

Kilian, Lutz. 2009. "Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market." *American Economic Review*, 99(3): 1053–69.

Känzig, Diego R. 2021. "The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements." *American Economic Review*, 111(4): 1092–1125.

Labandeira, Xavier, José M Labeaga, and Xiral López-Otero. 2017. "A meta-analysis on the price elasticity of energy demand." *Energy policy*, 102: 549–568.

Li, Dake, Mikkel Plagborg-Møller, and Christian K. Wolf. 2024. "Local projections vs. vars: Lessons from thousands of dgps." *Journal of Econometrics*, 244(2): 105722.

Mansanet-Bataller, Maria and Angel Pardo. 2009. "Impacts of regulatory announcements on CO2 prices." *The Journal of Energy Markets*, 2(2): 1–33.

Metcalf, Gilbert E. 2019. "On the economics of a carbon tax for the United States." *Brookings Papers on Economic Activity*, 2019(1): 405–484.

Montiel Olea, José Luis and Mikkel Plagborg-Møller. 2021. "Local projection inference is simpler and more robust than you think." *Econometrica*, 89(4): 1789–1823.

Nakamura, Emi and Jón Steinsson. 2018. "High-frequency identification of monetary non-neutrality: The information effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Plagborg-Møller, Mikkel and Christian K. Wolf. 2021. "Local projections and VARs estimate the same impulse responses." *Econometrica*, 89(2): 955–980.

Plagborg-Møller, Mikkel and Christian K Wolf. 2022. "Instrumental variable identification of dynamic variance decompositions." *Journal of Political Economy*, 130(8): 2164–2202.

Ramey, Valerie A. 2011. "Identifying government spending shocks: It's all in the timing." *The Quarterly Journal of Economics*, 126(1): 1–50.

Ramey, Valerie A. and Sarah Zubairy. 2018. "Government spending multipliers in good times and in bad: evidence from US historical data." *Journal of Political Economy*, 126(2): 850–901.

Rigobon, Roberto. 2003. "Identification through heteroskedasticity." *Review of Economics and*

Statistics, 85(4): 777–792.

Rigobon, Roberto and Brian Sack. 2004. “The impact of monetary policy on asset prices.” *Journal of Monetary Economics*, 51(8): 1553–1575.

Smets, Frank and Raf Wouters. 2003. “An estimated dynamic stochastic general equilibrium model of the euro area.” *Journal of the European Economic Association*, 1(5): 1123–1175.