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

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 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 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
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 document that a tighter carbon pricing regime leads to higher energy prices, lower emissions and more green innovation. This comes at the cost of a fall in economic activity, which is borne unequally across society: poorer households lower their consumption significantly while richer households are less affected. The poor are more exposed because of their higher energy share and, importantly, also experience a larger fall in income. Targeted fiscal policy can help alleviate these costs while maintaining emission reductions.

Diego R. Känzig
Department of Economics
Northwestern University
Kellogg Global Hub
2211 Campus Drive
Evanston, IL 60208
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. Yet, little is known about the effects of such policies in practice. Is carbon pricing effective at reducing emissions? What is the impact on output, employment and inflation, and who bears the economic costs of these policies?

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) is the largest carbon market in the world, covering around 40 percent of the EU's greenhouse gas (GHG) emissions. The market was established in phases and the regulations have been updated frequently. Using an event study approach, I collect 126 regulatory update events concerning the supply of emission allowances. By measuring the change in the carbon futures price in a tight window around the regulatory news, I isolate a series of carbon policy surprises. Reverse causality can be plausibly ruled out because economic conditions are known and priced by the market prior to the regulatory news, and they are unlikely to change within the tight window I consider. Using the surprise series as an instrument, 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 a strong, immediate increase in energy prices and a persistent fall in overall GHG emissions. Thus, carbon pricing is successful in achieving its goal of reducing emissions. However, this does not come without a cost. Consumer prices rise significantly and economic activity falls, as reflected in lower output and higher unemployment. Monetary policy leans against the inflationary pressures, likely exacerbating the effects on activity. Stock prices fall, but the shock does not appear to strongly transmit through financial markets. The main transmission channel appears to work through higher carbon prices, which passing through energy prices lead to a fall in income, and thus consumption and investment. Interestingly, the fall in activity turns out to be somewhat less persistent than the fall in emissions – improving the emissions intensity in the longer term. Consistent with that, I document a significant uptick in low-carbon patenting as carbon pricing creates an incentive for green innovation.

Carbon policy shocks also contribute meaningfully to historical variations in prices, emissions and macroeconomic aggregates. However, they did not account for the fall in emissions associated with the global financial crisis – supporting the validity of my identification strategy.

My results illustrate a trade-off between reducing emissions and the economic costs of climate policies. Importantly, these costs are not equally distributed across society. Using detailed household-level data, I document pervasive heterogeneity in the expenditure response to carbon policy shocks. While the expenditure of higher-income households only falls marginally, low-income households reduce their expenditure significantly and persistently. These households are more severely affected in two ways. First, they spend a larger share of their disposable income on energy and thus the higher energy bill leaves significantly fewer resources for other expenditures. Second, they experience a stronger fall in income, as they tend to work in sectors that are more impacted by the policy. Interestingly, these are not the sectors with the highest energy intensity but sectors that are more sensitive to changes in demand, producing more discretionary goods and services. Crucially, the magnitudes of the expenditure responses are larger than what can be accounted for by the direct effect through energy prices alone. This points to an important role of indirect, general equilibrium effects via income and employment. Based on my estimates, indirect effects can account for about two-thirds of the total effect on consumption.

My findings on the distributional impact of carbon pricing suggest that targeted fiscal policies could be an effective way to reduce the economic costs. To the extent that energy demand is inelastic, which turns out to be particularly the case for poorer households, this should not compromise emission reductions. This intuition is confirmed in a climate-economy model with nominal rigidities and heterogeneity in households' energy expenditure shares, income incidence and marginal propensities to consume (MPCs). The model can account for the observed empirical responses to carbon policy reasonably well. Using the model, I show that redistributing carbon revenues can mitigate the fall in aggregate consumption and reduce the regressive distributional consequences of carbon pricing, without compromising emission reductions to a significant extent. Finally, I provide some suggestive evidence that carbon pricing leads to a fall in the support for climate-related policies that is particularly pronounced among low-income households. Therefore, mitigating the distributional impact may also help to increase the public support for climate policy.

A comprehensive series of sensitivity checks indicate that the results are robust along a number of dimensions including the selection of event dates, the construction of the instrument, the estimation technique, the model specification, and the sample period. Importantly, the results are also robust to accounting for confounding news over the event window using an heteroskedasticity-based estimator.

Related literature and contribution. This paper contributes to a growing literature studying the effects of climate policy and carbon pricing in particular. While there is mounting evidence on the effectiveness of such policies for emission reductions (Martin, De Preux, and Wagner, 2014; Andersson, 2019, among others), less is known about the economic effects. A number of studies have analyzed the macroeconomic effects of the British Columbia carbon tax, finding no significant impacts on GDP (Metcalf, 2019; Bernard and Kichian, 2021). Metcalf and Stock (2020*a*,*b*) study the macroeconomic impacts of carbon taxes in European countries. They find no robust evidence of a negative effect of the tax on employment or GDP growth.[1] In a similar vein, Konradt and Weder di Mauro (2021) find that carbon taxes in Europe and Canada do not appear to be inflationary. In contrast, theoretical studies based on computable general equilibrium models tend to find contractionary output effects (see e.g. McKibbin et al., 2017; Goulder and Hafstead, 2018). I contribute to this literature by providing new estimates based on the EU ETS, the largest carbon market in the world.

A large literature has studied the macroeconomic effects of discretionary tax changes more generally. To address the endogeneity of tax changes, the literature has used SVAR techniques (Blanchard and Perotti, 2002) and narrative methods (Romer and Romer, 2010). The narrative approach points to large macroeconomic effects: a tax increase leads to a significant and persistent decline of output and its components (see also Mertens and Ravn, 2013; Cloyne, 2013). However, it is unclear how much we can learn from these estimates with respect to carbon pricing, which is enacted to correct an externality and not because of past decisions or ideology. While the motivation behind carbon pricing is arguably long-term and thus more likely unrelated to the current state of the economy – similar to the tax changes considered in Romer and Romer (2010) – it is still perceivable that regulatory decisions also take economic conditions into account.

To address this potential endogeneity in carbon pricing, I propose a novel identification strategy exploiting high-frequency price variation. From a methodological viewpoint, my approach is closely related to the literature on high-frequency identification, which was developed in the monetary policy setting

[1] Metcalf and Stock (2020*a*,*b*) study the effects of national carbon taxes, which are present in many European countries and cover sectors that are not included in the EU ETS. A key difference is that European carbon taxes generally do not cover the power sector, which is part of the ETS.

(Kuttner, 2001; Gürkaynak, Sack, and Swanson, 2005; Gertler and Karadi, 2015; Nakamura and Steinsson, 2018, among others) and more recently employed in the global oil market context (Känzig, 2021). In this literature, policy surprises are identified using high-frequency asset price movements around policy events, such as FOMC or OPEC meetings. The idea is to isolate the impact of policy news by measuring the change in asset prices in a tight window around the events.

I contribute to this literature by extending the high-frequency identification approach to climate policy, exploiting institutional features of the European carbon market. A number of studies have used event study techniques to analyze the effects of regulatory news on carbon, energy and stock prices (Mansanet-Bataller and Pardo, 2009; Fan et al., 2017; Bushnell, Chong, and Mansur, 2013, among others). To the best of my knowledge, however, this paper is the first to exploit these regulatory updates to analyze the economic effects of carbon pricing. The approach is very general and could also be employed to evaluate the performance of other cap and trade systems.

Equipped with this novel identification strategy, I provide new evidence not only on the aggregate but also on the distributional consequences of carbon pricing. Among policymakers, there is growing consensus that the transition towards a low-carbon economy should involve fairness and equity considerations (European Comission, 2021). Against this backdrop, it is crucial to understand how carbon pricing affects economic inequality. I find that carbon pricing in the EU has been more regressive than commonly thought, burdening lower-income households substantially more than richer ones. This stands in contrast to existing studies, which tend to find a more modest regressive impact (Beznoska, Cludius, and Steiner, 2012; Ohlendorf et al., 2021). My findings illustrate the importance of accounting for indirect, general-equilibrium effects via income and employment; solely focusing on the direct effects via higher energy prices can understate the actual distributional impact.

Finally, I show that carbon-policy induced changes in energy prices transmit through a powerful demand channel that can outweigh the traditional cost channel. This has important implications for the transmission of energy price shocks more broadly and speaks to a growing literature studying the role of Keynesian supply shocks (see e.g. Guerrieri et al., 2022). The demand channel can be reinforced by the monetary policy reaction (Bernanke, Gertler, and Watson, 1997) and the unequal incidence on constrained households (Bilbiie, 2008; Auclert, 2019; Patterson, 2021). Thus, the distributional consequences do not only matter for inequality but also for the transmission of the policy to the macroeconomy. To formalize this in the context of carbon tax policy, I develop a climate-economy model (in spirit of Heutel, 2012; Golosov et al., 2014; Annicchiarico and Di Dio, 2015) with nominal rigidities and household heterogeneity. In this sense, I also contribute to an influential literature studying the role of heterogeneity in the transmission of fiscal policies (see e.g. Johnson, Parker, and Souleles, 2006; Kaplan and Violante, 2014; Cloyne and Surico, 2017, among many others).

Roadmap. The paper proceeds as follows. In the next section, I provide institutional background on the European carbon market and discuss the high-frequency identification strategy. Section 3 covers the econometric approach. Section 4 presents the results on the aggregate effects of carbon pricing, on emissions and the macroeconomy. Section 5 looks into the heterogeneous effects of carbon pricing, using detailed household-level data. I analyze the distributional impact, the relative importance of different transmission channels, and the role of redistributing carbon revenues. Section 6 looks beyond the short term and analyzes the impact on attitudes towards climate policies and the effects on green innovation. Section 7 concludes.

2. Institutional Background and Identification

2.1. The European carbon market

The European emissions trading system is the cornerstone of the EU's policy to combat climate change. It is the largest carbon market in the world and also has one of the longest implementation histories. Established in 2005, it covers more than 11,000 heavy energy-using installations and airlines, accounting for around 40 percent of the EU's greenhouse gas emissions.

The market operates under the cap and trade principle. Different from a carbon tax, a cap is set on the total amount of certain greenhouse gases that can be emitted by installations in the system. The cap is reduced over time so that total emissions fall. Within the cap, emission allowances are auctioned off or allocated for free among the companies in the system, and can subsequently be traded. Alternatively, companies can also use limited amounts of international credits from emission-saving projects around the world. Regulated companies must monitor and report their emissions. Each year, the companies must surrender sufficient allowances to cover their emissions. This is enforced with heavy fines. If a company reduces its emissions, it can keep the spare allowances for future needs or sell them to another company short of allowances (European Comission, 2020*a*).

A brief history of the EU ETS. The development of the EU ETS was divided into different phases. Figure 1 shows the evolution of the carbon price over the phases of the system. The first phase lasted three years, from 2005 to 2007. This period was a pilot phase to prepare for phase two, where the system had to run efficiently to help the EU meet its Kyoto targets. In this initial phase, almost all allowances were freely allocated at the national level. In absence of reliable emissions data, phase one caps were set on the basis of estimates. In 2006, the carbon price fell significantly as it became apparent that the total amount of allowances issued exceeded total emissions, and eventually converged to zero as phase one allowances could not be transferred to phase two.

Figure 1: The EU Carbon Price

Notes: The EU carbon price, as measured by the price of the first EUA futures contract over the different phases of the EU ETS.

The second phase ran from 2008 until 2012, coinciding with the first commitment period of the Kyoto Protocol where the countries in the EU ETS had concrete emission targets to meet. Because verified annual emissions data from the pilot phase was now available, the cap on allowances was reduced in this phase, based on actual emissions. The proportion of free allocation fell slightly, several countries started to hold auctions, and businesses were allowed to buy limited amounts of international credits. The commission also started to extend the system to cover more gases and sectors; in 2012 the aviation sector was included, even though this only applied for flights within the European Economic Area. Despite these changes, EU carbon prices remained at moderate levels. This was mainly because the 2008 economic crisis led to large fall in emissions. As this was not reflected in the way the caps were set, this led to a large surplus of allowances weighing down on prices.

The subsequent third phase began in 2013 and ran until the end of 2020. Learning from the previous phases, the system was changed in a number of key respects. In particular, the new system relies on a single, EU-wide cap on emissions in place of the previous national caps, auctioning became the default way of allocating allowances with harmonized allocation rules for the allowances still allocated for free, and the system covers more sectors and gases, in particular nitrous oxide and perfluorocarbons in addition to carbon dioxide. In 2014, the Commission postponed the auctioning of 900 million emission allowances to address the surplus of allowances that has built up since the Great Recession ('back-loading'). Later, the Commission introduced a market stability reserve, which became operative in January 2019. This reserve has the aim to reduce the current surplus of allowances and improve the system's resilience to major shocks by adjusting the supply of allowances. To this end, back-loaded and unallocated allowances were transferred to the reserve rather than auctioned in the last years of phase three.

The current, fourth phase spans the period from 2021 to 2030. The legislative framework for this trading period was revised in early 2018. To achieve the EU's 2030 emission reduction targets, the pace of annual reductions in total allowances is increased to 2.2 percent from the previous 1.74 percent and the market stability reserve is reinforced to improve the systems resilience to future shocks. More recently, the Commission has proposed to further revise and expand the scope of the EU ETS, with the aim to achieve a climate-neutral EU by 2050 (see European Comission, 2020*a*).

Regulatory events. Given its pioneering role, the establishment of the European carbon market has followed a learning-by-doing process. As illustrated above, since the start in 2005, the system has been expanded considerably and its institutions and rules have been continuously updated to address issues encountered in the market, improve market efficiency, and reduce information asymmetry and market distortions.

Building on the event study literature, I collect a comprehensive list of regulatory events in the EU ETS. These regulatory update events can take the form of a decision of the European Commission, a vote of the European Parliament or a judgment of a European court. Of primary interest in this paper are regulatory news regarding the *supply* of emission allowances. Thus, I focus on news concerning the overall cap in the EU ETS, the free allocation of allowances, the auctioning of allowances as well as the use of international credits. Going through the official journal of the European Union as well as the European Commission Climate Action news archive, I could identify 126 such events during the period from 2005 to 2019. The events as well as the sources are detailed in Table A.1 in the Appendix. There are only a few events that concern the setting of the overall cap in the system. In the first two phases, the key events concern decisions on the national allocation plans (NAP) of the individual member states, e.g. the commission approving or rejecting allocation plans or court rulings in legal conflicts about the free allocation of allowances. With the move to auctioning as the default way of allocating allowances, decisions on the timing and quantities of emission allowances to be auctioned became the most important regulatory news in phase three. Finally, starting from phase two, there were also a number of important events related to the use and entitlement of international credits.

The selection of events is a crucial factor in event studies. As the baseline, I use all of the identified events, however, in Appendix C.1, I study the sensitivity of the results with respect to different event types in detail.

Carbon futures markets. There exist several organized markets where EU emission allowances (EUAs) can be traded. An EUA is defined as the right to emit one ton of carbon dioxide equivalent gas and is traded in spot markets such as Bluenext in Paris, EEX in Leipzig or Nord Pool in Oslo. Furthermore, there exist also futures markets on EUAs, such as the EEX in Leipzig and ICE in London. In 2018, the cumulative trading volume in the relevant futures and spot markets was about 10 billion EUA (DEHSt, 2019). The most liquid markets to trade emission allowances are the futures markets. In this paper, I focus on data from the ICE, which has been found to dominate the price discovery process in the European carbon market (Stefan and Wellenreuther, 2020).

2.2. High-frequency identification

Since policies to fight climate change are long-term in nature, they are likely less subject to endogeneity concerns than other fiscal polices (Romer and Romer, 2010). However, to properly address the concern that regulatory decisions in the carbon market may take economic conditions into account, I adopt a high-frequency identification approach.

The institutional framework of the European carbon market provides an ideal setting in this respect. First, as discussed above, there are frequent regulatory updates in the market that can have significant effects on the price of emission allowances. Second, there exist liquid futures markets for trading allowances. This motivates the idea to construct a series of carbon policy surprises by looking at how carbon prices change around regulatory events in the carbon market. By measuring the price change within a sufficiently tight window around the event, reverse causality of the state of the economy can be plausibly ruled out because it is known and priced prior to the decision, and unlikely to change within the tight window.

Specifically, I compute the carbon policy surprise series as the change in the EUA futures price on the day of the regulatory event compared to the last trading day before the event. To account for the fact that carbon prices were close to zero at the end of the first phase, I measure the surprises as the EUR change in carbon prices relative to the prevailing wholesale electricity price on the day before the event:

$$CPSurprise_{t,d} = \frac{F^{carbon}_{t,d} - F^{carbon}_{t,d-1}}{P^{elec}_{t,d-1}}, \tag{1}$$

where d and t indicate the day and the month of the event, respectively, $F_{t,d}$ is the settlement price of the EUA futures contract, and $P^{elec}_{t,d-1}$ is the wholesale electricity price. This allows me to isolate some variation in the carbon price that is driven by the regulatory news, assuming that risk premia do not change over the narrow event window.[2] An alternative approach is to compute the surprise series as the percentage change in the carbon price around the event. Reassuringly, this produces very similar results, especially when excluding the second half of 2007 when carbon prices were approaching zero as the trial phase was coming to an end, see Appendix C for more details.

The daily surprises, $CPSurprise_{t,d}$, are then aggregated to a monthly series, $CPSurprise_t$, by summing over the daily surprises in a given month. In months without any regulatory events, the series takes zero value.

Figure 2 shows the resulting carbon policy surprise series. We can see that regulatory news can have a significant impact on carbon prices, with some news moving carbon prices by close to 1.5 percent, relative to wholesale electricity prices.[3] In the first phase, there were a number of significant events concerning the free allocation of allowances. For instance, in June 2005 the initial national allocation plans were finally completed, which lead to a significant increase in prices. The beginning of the second phase was characterized by only few regulatory news. This changed dramatically from the second half of phase two through the first years of phase three, which were marked by many significant carbon policy surprises. For example, carbon prices jumped up in March 2011 after the Commission proposed to start the auctioning of allowances earlier than originally

[2]While futures prices are in general subject to risk premia, there is evidence that these premia vary primarily at lower frequencies (Piazzesi and Swanson, 2008; Hamilton, 2009; Nakamura and Steinsson, 2018).

[3]Carbon prices, per se, turn out to be more volatile, with some announcements moving prices in excess of 40 percent, see Appendix C.1.

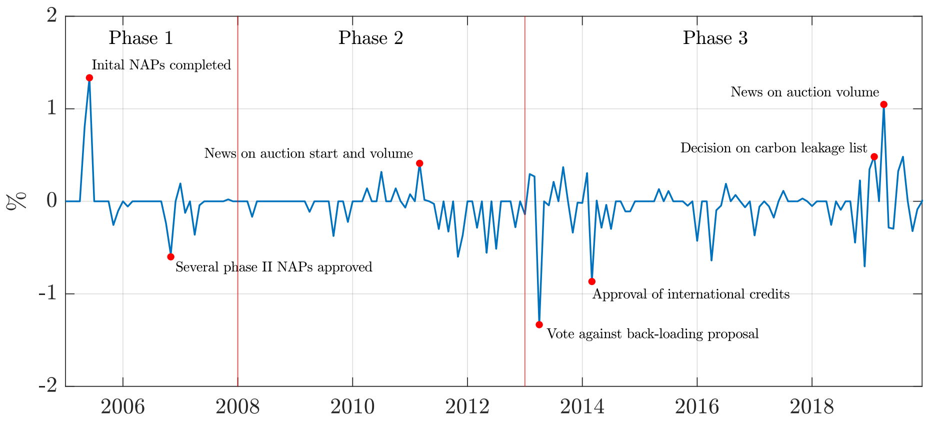

Figure 2: The Carbon Policy Surprise Series

Notes: The carbon policy surprise series, constructed as the change of the EUA futures price around regulatory policy events concerning the supply of EU emission allowances relative to the prevailing wholesale electricity price.

planned. On April 16, 2013 the European Parliament voted against the Commission's back-loading proposal, which led to a massive fall in carbon prices. And in March 2014, the Commission approved two batches of international credit entitlement tables, causing a significant fall in prices, just to name a few. There were also a number of significant surprises towards the end of phase three. In February 2019, carbon prices jumped up following news on the adoption of a stricter carbon leakage list, and in April 2019, carbon prices increased further, after an update on auction volumes in EFTA countries contributing to bullish sentiment in the market.

Construction choices and diagnostics. A crucial choice in high-frequency identification concerns the size of the event window. There is a trade-off between capturing the entire response to the announcement and the threat of other news confounding the response, so-called background noise (Nakamura and Steinsson, 2018). To give markets enough time to respond to the 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 events are mostly unavailable. However, to mitigate concerns about background noise when using a daily window, I also present results from a heteroskedasticity-based approach that allows for background noise in the surprise series (see Appendix C.2).

Another choice concerns the maturity of the futures contract. I focus here on the front contract (the contract with the closest expiry date) for two reasons. First, it is the most liquid contract and thus gives the best price signal. Second, near-dated contracts also tend to be less sensitive to risk premia (Baumeister and

Kilian, 2017; Nakamura and Steinsson, 2018). Thus, focusing on the front contract helps to further mitigate concerns about time-varying risk premia.

To be able to interpret the resulting series as carbon policy surprises, it is crucial that the events do not release other information such as news about the demand of emission allowances or economic activity in the EU more generally. To address these concerns, I select only regulatory update events that were specifically about changes to the supply of emission allowances in the European carbon market and do not include broader events such as outcomes of Conference of the Parties (COP) meetings or other international conferences. In a series of sensitivity checks, I also show that the results are not driven by a particular subset of events. In particular, the results are robust to excluding events from the first trial phase or excluding event days in periods of economic distress, such as the Great Recession or the European debt crisis (see Appendix C.1).

Finally, I perform a number of additional diagnostic checks on the surprise series as proposed in Ramey (2016), in particular with regards to autocorrelation, forecastability and correlation with other structural shocks. I find no evidence that the series is serially correlated. The p-value for the Q-statistic that all autocorrelations are zero is 0.97. I also find no evidence that macroeconomic or financial variables have any power in forecasting the surprise series. For all variables considered, the p-values for the Granger causality test are far above conventional significance levels, with the joint test having a p-value of 0.93. Finally, 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. Overall, this evidence supports 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 has many desirable properties. Nonetheless, it is only an imperfect measure of the shock of interest because it may not capture all relevant instances of regulatory news in the carbon market and could be measured with error (see also Stock and Watson, 2018).

Therefore, I do not use it as a direct shock measure but as an *instrument*. Provided that the surprise series is correlated with the carbon policy shock but uncorrelated with all other shocks, we can use it to estimate the dynamic causal effects of a carbon policy shock. Because of the short sample at hand, I rely on VAR techniques for estimation. For identification, I rely on the external instrument approach (Stock, 2008; Stock and Watson, 2012; Mertens and Ravn, 2013). While this approach tends to be very efficient, it provides biased estimates if the VAR is not invertible. Thus, I also present results from an internal instrument approach (Ramey, 2011; Plagborg-Møller and Wolf, 2019), which includes the instrument in the VAR and is robust to problems of non-invertibility.

An alternative strategy would be to estimate the dynamic causal effects using local projections (see Jordà, Schularick, and Taylor, 2015; Ramey and Zubairy, 2018). However, this approach is quite demanding given the short sample, as it involves a distinct IV regression for each impulse horizon. Importantly, Plagborg-Møller and Wolf (2019) show that the internal instrument VAR and the LP-IV rely on the same invertibility-robust identifying restrictions and identify, in population, the same relative impulse responses. In Appendix B.2, I compare the LP-IV to the internal instrument VAR responses in the sample at hand. Reassuringly, the responses turn out to be similar, even though the LP responses are more jagged and less precisely estimated.

3.1. Framework

Consider the standard VAR model

$$\mathbf{y}_t = \mathbf{b} + \mathbf{B}_1\mathbf{y}_{t-1} + \cdots + \mathbf{B}_p\mathbf{y}_{t-p} + \mathbf{u}_t, \tag{2}$$

where p is the lag order, $\mathbf{y}_t$ is a $n \times 1$ vector of endogenous variables, $\mathbf{u}_t$ is a $n \times 1$ vector of reduced-form innovations with covariance matrix $\mathrm{Var}(\mathbf{u}_t) = \mathbf{\Sigma}$, $\mathbf{b}$ is a $n \times 1$ vector of constants, and $\mathbf{B}_1, \ldots, \mathbf{B}_p$ are $n \times n$ coefficient matrices.

Under the assumption that the VAR is invertible, we can write the innovations $\mathbf{u}_t$ as linear combinations of the structural shocks $\boldsymbol{\varepsilon}_t$:

$$\mathbf{u}_t = \mathbf{S}\boldsymbol{\varepsilon}_t. \tag{3}$$

By definition, the structural shocks are mutually uncorrelated, i.e. $\mathrm{Var}(\boldsymbol{\varepsilon}_t) = \mathbf{\Omega}$ is diagonal. From the invertibility assumption (3), we get the standard covariance restrictions $\mathbf{\Sigma} = \mathbf{S}\mathbf{\Omega}\mathbf{S}'$.

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

$$\mathbb{E}[z_t \varepsilon_{1,t}] = \alpha \neq 0 \tag{4}$$

$$\mathbb{E}[z_t \boldsymbol{\varepsilon}_{2:n,t}] = \mathbf{0}, \tag{5}$$

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 (4) is the relevance requirement and assumption (5) is the exogeneity condition. These assumptions, in combination with the invertibility requirement (3), identify $\mathbf{s}_1$ up to sign and scale:

$$\mathbf{s}_1 \propto \frac{\mathbb{E}[z_t \mathbf{u}_t]}{\mathbb{E}[z_t \mathrm{u}_{1,t}]}, \tag{6}$$

provided that $E[z_t \mathrm{u}_{1,t}] \neq 0$.[4] To facilitate interpretation, we 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, as proposed by Jentsch and Lunsford (2019).

Internal instrument approach. To address potential problems of non-invertibility, I also employ an internal instrument approach. For identification, we have to assume in addition to (4)-(5) that the instrument is orthogonal to leads and lags of the structural shocks:

$$\mathbb{E}[z_t \boldsymbol{\varepsilon}_{t+j}] = \mathbf{0}, \quad \text{for } j \neq 0. \tag{7}$$

In return, we can dispense of the invertibility assumption underlying equation (3). Under these assumptions, we can estimate the dynamic causal effects by augmenting the VAR with the instrument ordered first, $\bar{\mathbf{y}}_t = (z_t, \mathbf{y}_t')'$, and computing the impulse responses to the first orthogonalized innovation, $\bar{\mathbf{s}}_1 = [\mathrm{chol}(\bar{\boldsymbol{\Sigma}})]_{.,1} / [\mathrm{chol}(\bar{\boldsymbol{\Sigma}})]_{1,1}$. As Plagborg-Møller and Wolf (2019) show, this approach consistently estimates the relative impulse responses even if the instrument is contaminated with measurement error or if the shock is non-invertible.

[4] To be more precise, the VAR does not have to be fully invertible for identification with external instruments. As Miranda-Agrippino and Ricco (2018) show, it suffices if the shock of interest is invertible in combination with a limited lead-lag exogeneity condition.

3.2. Empirical specification

Studying the macroeconomic impact of carbon policy requires modeling the European economy and the carbon market jointly. The baseline specification consists of eight variables. For the climate block, I use the energy component of the HICP as well as total GHG emissions.[5] To proxy the state of the economy, I include the headline HICP, industrial production, and the unemployment rate. Given that 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. However, using the shadow rate or other longer-term rates produces similar results. Finally, I include a stock market index and the Brent crude oil price, deflated by the HICP, as financial indicators. More information on the data and its sources can be found in Appendix A.2.

The sample period starts in January 1999, when the euro was introduced, and runs until December 2019, stopping before the outbreak of the Covid pandemic. Recall that the carbon policy surprise series is only available from 2005 when the carbon market was established. To deal with this discrepancy, the missing values in the surprise series are censored to zero (see Noh, 2019, for a formal justification of this approach). The motivation for using a longer sample is to increase the precision of the estimates. However, restricting the sample to 2005-2019 produces very similar results.[6]

Following Sims, Stock, and Watson (1990), I estimate the VARs in levels. Apart from the unemployment and the two-year rate, all variables enter in log-levels. As controls I use six lags of all variables and in terms of deterministics only a constant term is included. However, the results turn out the be robust with respect to all of these choices (see Appendix C.3).

4. The Aggregate Effects of Carbon Pricing

4.1. The impact on emissions and the macroeconomy

In this section, we study how carbon policy shocks affect the macroeconomy through the lens of the baseline model. Recall, the main identifying assump-

[5]Unfortunately, GHG emissions are only available at the annual frequency. Therefore, I construct a monthly measure of emissions using the Chow-Lin temporal disaggregation method with indicators from Quilis's (2020) code suite. As the relevant monthly indicators, I include the HICP energy and industrial production. The results are robust to extending the list of indicators used.

[6]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).

tion behind the external instrument approach is that the instrument is correlated with the structural shock of interest but uncorrelated with all other structural shocks. However, to be able to conduct standard inference, the instrument has to be sufficiently strong. To analyze whether this is the case, I perform the weak instruments test by Montiel Olea and Pflueger (2013).

The carbon policy surprise series turns out to be a strong instrument. The heteroskedasticity-robust F-statistic in the first stage is 17.43. As this is clearly above conventional critical values, we conclude that the instrument appears to be sufficiently strong to conduct standard inference.

Having established that the carbon policy surprise series is a strong instrument, we can now turn to the discussion of the macroeconomic and environmental impacts of carbon policy shocks. Figure 3 shows the impulse responses to the identified carbon policy shock, normalized to increase the HICP energy component by one percent on impact. The solid black lines are the point estimates and the shaded areas are 68 and 90 percent confidence bands based on 10,000 bootstrap replications.

A restrictive carbon policy shock leads to a strong, immediate increase in energy prices and a significant and persistent fall in GHG emissions. Thus, carbon pricing appears to be successful at reducing emissions and mitigating climate change by increasing the cost of emitting. Turning to the macroeconomic variables, we can see that the fall in emissions does not come without cost. Industrial production falls and the unemployment rate rises significantly. The labor market response turns out to be particularly pronounced. Consumer prices, as measured by the HICP, increase. The pass-through is strong for headline, however, core consumer prices tend to increase as well but the response is more short-lived (see Figure B.4 in the Appendix). Monetary policy appears to lean against the inflationary pressures, which likely exacerbates the effects on activity. Stock prices do not respond significantly on impact but then tend to fall, anticipating the fall in activity. However, the response is imprecisely estimated. Oil prices on the other hand increase significantly, reflecting the fact that European oil producers and refineries are also covered by the emissions trading scheme.[7]

In terms of magnitudes, the shock leads to an increase in energy prices of about 1.6 percent at peak. GHG emissions and industrial production decline by around 0.6 percent, the unemployment rate rises by about 0.2 percentage points and consumer prices increase by slightly more than 0.2 percent. The two-year

[7]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.

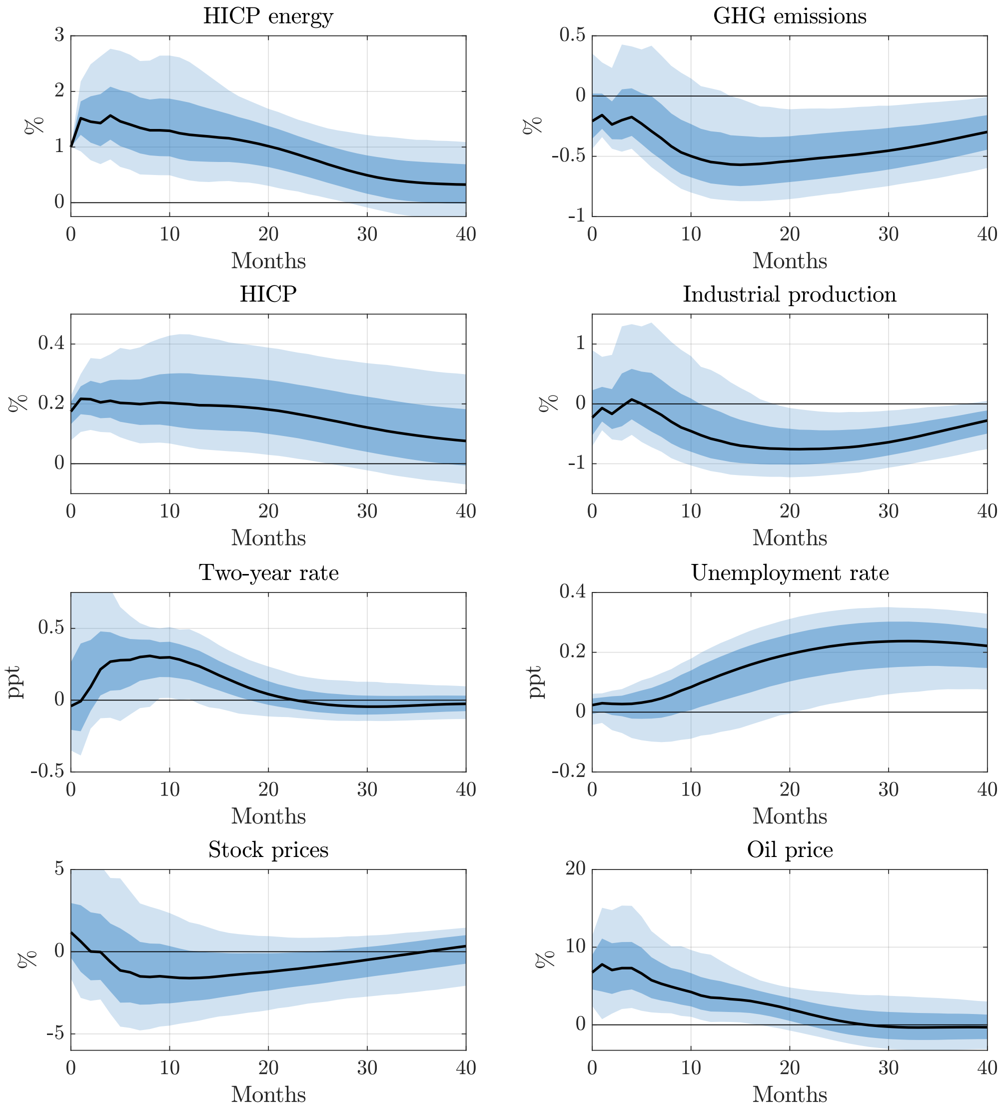

First stage regression: F-statistic: 17.43, R^2: 2.85%

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. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

rate increases by about 25 basis points, stock prices fall by over 1.5 percent and oil prices increase by around 8 percent – all measured at the peak of the responses. Thus, the responses are not only statistically but also economically significant. While the price impacts materialize rather quickly, economic activity and GHG emissions only fall with a substantial lag. The implied price elasticity of emissions lies in the ballpark of the estimates in Metcalf (2019). It is also interesting to observe that the fall in output appears to be less persistent than the fall in emissions – implying an improvement in the emissions intensity at longer horizons. We will revisit this finding in Section 6, where I study the effects of carbon pricing on green innovation.

The results from the internal instrument model turn out to be very similar, see Appendix B.2. The signs are all consistent and the responses are of similar shape and magnitude. Only the estimated response of the two-year rate is somewhat less stable. The pre-test for invertibility by Plagborg-Møller and Wolf (2019) can also not reject the null of invertibility at the 10 percent level. Overall, these findings suggest that the results are robust to relaxing the assumption of invertibility.

To summarize, the above findings clearly illustrate the policy trade-off between reducing emissions and thus the future costs of climate change and the current economic costs associated with climate change mitigation policies. It is useful to contrast these results to Metcalf and Stock (2020*a*), who study the economic and environmental impact of European carbon taxes. They find that these taxes were successful at reducing emissions but had no robust negative effect on output and employment.

A crucial difference is that European carbon taxes do not include the power sector, which is covered by the EU ETS, and plays a crucial role for the macroeconomic effects that I estimate. In fact, in terms of magnitudes my results are consistent with previous evidence on energy price shocks, such as oil shocks (see e.g. Kilian, 2009; Baumeister and Hamilton, 2019; Känzig, 2021). Furthermore, in many European countries, carbon taxes were implemented as part of a broader tax reform which often included other changes to the tax code to cushion the impact of carbon taxes. As we will discuss in Section 5.5, the distribution of carbon revenues plays an important role in the transmission of carbon policy shocks. Finally, given that the EU is a monetary union, we would not expect a monetary response to national carbon tax policies. By contrast, monetary policy seems to lean against the inflationary pressures from the EU ETS, which also helps explain the larger economic impacts (Bernanke, Gertler, and Watson, 1997).

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 indicate that the results are robust along a number of dimensions including the selection of event dates, the construction of the instrument, the estimation technique, the model specification, and the sample period.

4.2. Historical importance

In the previous section, 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 4 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.

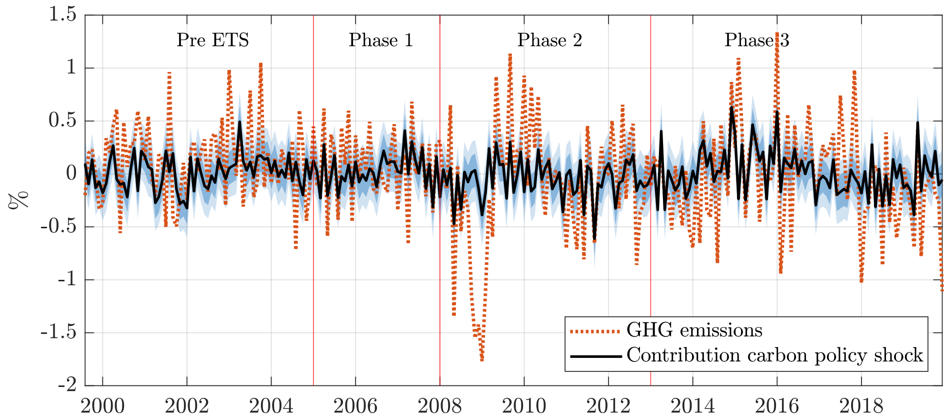

Figure 4: Historical Decomposition of GHG Emissions Growth

Notes: The figure shows the cumulative historical contribution of carbon policy shocks over the estimation sample for GHG emissions growth against the actual evolution of emissions growth. The solid line is the point estimate and the dark and light shaded areas are 68 and 90 percent confidence bands, respectively.

On average, carbon policy shocks account for close to a quarter of the variations in emissions at horizons up to two years. Furthermore, carbon policy shocks also explain a non-negligible share of the variations in energy prices and other macroeconomic and financial variables (see the variance decomposition in Appendix B.2).

4.3. Wider effects and propagation channels

The above results suggest that energy prices play an important role in the transmission of carbon policy shocks. Power producers seem to pass through the emission costs to energy prices to a significant extent, as can be seen from the strong energy price response. This is in line with previous empirical evidence (see e.g. Fabra and Reguant, 2014).

To get a better understanding of how carbon policy shocks transmit to the economy, I analyze the effects on a wider range of macroeconomic and financial variables. To compute the impulse responses, I extract the carbon policy shock from the monthly VAR as $CPShock_t = \mathbf{s}_1'\mathbf{\Sigma}^{-1}\mathbf{u}_t$ (for a derivation, see Stock and Watson, 2018) and estimate the effects using simple local projections:

$$y_{i,t+h} = \beta^i_{h,0} + \psi^i_h CPShock_t + \beta^i_{h,1} y_{i,t-1} + \ldots + \beta^i_{h,p} y_{i,t-p} + \xi_{i,t,h}, \tag{8}$$

where ψ^i_h is the effect on variable i at horizon h. Importantly, we can also use this approach to estimate the effects on variables that are only available at the quarterly or even annual frequency. In this case, we aggregate the shock $CPShock_t$ by summing over the respective months before running the local projections. Using the shock series directly in the local projections instead of the high-frequency surprises increases the statistical power of these regressions, as the shock series is consistently observed and spans the entire sample. Note, however, that this comes at the cost of assuming invertibility. Throughout the paper, I normalize the responses to have the same peak effect on HICP energy as in the baseline model to facilitate comparison of the results. The confidence bands are computed using the lag-augmentation approach (Montiel Olea and Plagborg-Møller, 2020).[8]

Figure 5 shows the impulse responses of real GDP, consumption, investment, and wages. Consistent with the monthly evidence, we find that the shock leads to a significant fall in real GDP. Looking at the different components, we can see that the fall in activity appears to be driven by lower consumption and investment. The consumption response turns out to be particularly pronounced.

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 considered to be inelastic, consumers and firms have less money to spend and invest after paying their energy bills (see e.g. 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.

Interestingly, the magnitudes of the effects are much larger than what can be accounted for by the direct effect through higher energy prices alone. If energy demand is completely inelastic, the direct price effect is bounded by the energy share in expenditure, which is around 10 percent in Europe. Given the shock magnitude, we would thus expect a direct impact on consumption of around 15 percent. However, the estimated consumption response is substantially larger than that, suggesting indirect effects play an important role in the transmission of carbon policy shocks. In fact, the significant fall in wages coupled with the em-

[8]As controls in the local projections, I use 7 lags for monthly variables, 3 lags for quarterly variables and 1 lag for annual variables.

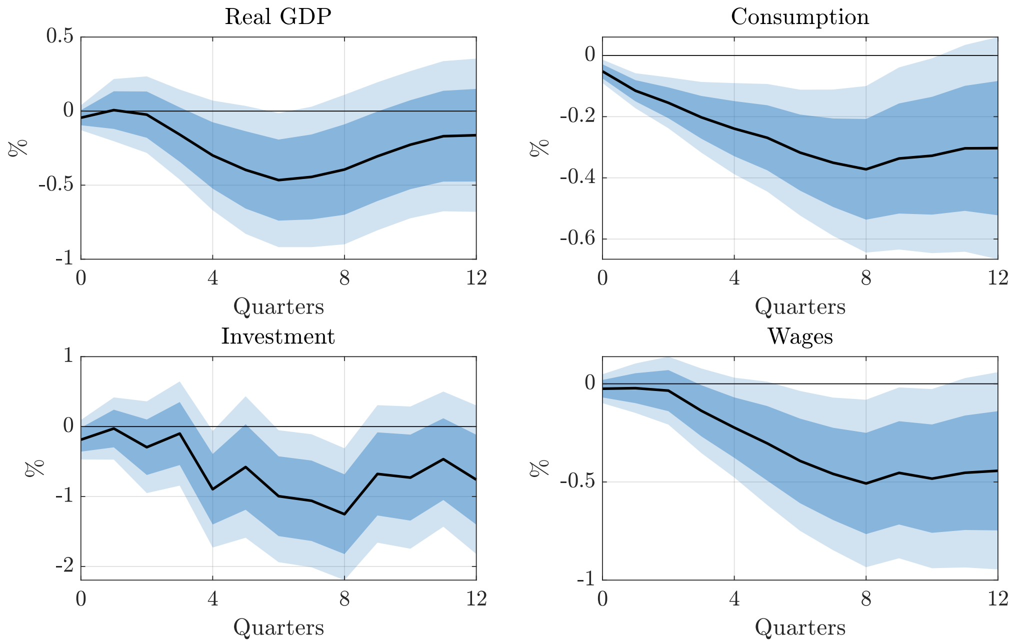

Figure 5: Effect on GDP, Consumption, Investment and Wages

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.

ployment effects documented in Section 4.1 strongly supports this notion. The mechanism works as follows. After a carbon policy-induced energy price increase, the direct decrease in households' and firms' consumption and investment expenditure leads to lower output and exerts downward pressure on employment and wages. At the same time, interest rates increase as monetary policy leans against the inflationary pressures coming from higher energy prices, further exacerbating these effects. The additional fall in aggregate demand induced by lower employment and wages lies at the core of the indirect effect.

There is little evidence that carbon policy shocks strongly transmit through financial channels or elevated uncertainty. As we have seen, the stock market displays a muted response and measures of financial conditions such as credit spreads do not respond significantly as well. Similarly, there is no significant response of uncertainty measures (see Figure B.5 in the Appendix). Thus, these alternative channels are unlikely to play a dominant role in the transmission of carbon policy.

Finally, there are also transmission channels that work through the supply side of the economy. Baqaee and Farhi (2019) focus on the input-output structure of firms. They show that the centrality of the power sector can amplify the effects of energy shocks in the presence of non-linearities. However, given that the sample of interest was characterized by relatively small shocks, we would not expect non-linearities to play a major role. In the next section, I shed more light on the role of different transmission channels using detailed household micro data.

5. The Heterogeneous Effects of Carbon Pricing

Recently, there has been a big debate in Europe on energy poverty and the distributional effects of climate policy amid the European Green Deal (European Comission, 2021). The situation has since been exacerbated by the Russian invasion of Ukraine, which led to a substantial increase in energy bills.

Against this backdrop, it is crucial to better understand the distributional impact of the EU ETS. If certain groups are left behind, this could ultimately undermine the success of climate policy. To this end, I study the heterogeneous effects of carbon pricing on households. This will help to get a better picture on how carbon pricing affects economic inequality. Furthermore, looking into potential heterogeneities in the consumption responses helps to better understand the transmission channels at work. There is reason to believe that there are important heterogeneities at play. First, the direct effect through energy prices crucially depends on the energy expenditure share, which is highly heterogeneous across households. Second, households can also be affected differently in indirect ways, as they may face different impacts on their incomes. As poorer households tend to have a higher energy share and their income tends to be more cyclical, we expect the impact to be regressive.

5.1. Household survey data

To be able to analyze the heterogeneous effects of carbon policy shocks on households, we need detailed micro data on consumption expenditure and income at a regular frequency for a sample spanning the last 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 that has such data as part of the Living Costs and Food Survey (LCFS).[9]

The LCFS is the major survey on household spending in the UK and provides high-quality, detailed information on expenditure, income, and household characteristics. The survey is fielded in annual waves with interviews being con-

[9]The UK was part of the EU ETS until the end of 2020. Over the sample of interest, the aggregate effects in the UK are comparable to the ones documented at the EU level, see Figure B.6 in the Appendix. To further mitigate concerns about external validity, I show that the results for other European countries are comparable, using similar survey data for Denmark and Spain, see Figure B.18.

ducted throughout the year and across the whole of the UK. I compile a repeated cross-section based on the last 20 waves, spanning the period from 1999 to 2019. Each wave contains around 6,000 households, generating over 120,000 observations in total. To compute measures of income and expenditure, I first express the variables in per capita terms by dividing household variables by the number of household members. In a next step, I deflate the variables by the (harmonized) consumer price index to express them in real terms. For more information, see Appendix A.3.

Ideally, we would like to observe how individual consumption expenditure and income evolve over time. Unfortunately, the LCFS being a repeated cross-section has no such panel dimension. To construct a pseudo-panel, it is common to use a grouping estimator in the spirit of Browning, Deaton, and Irish (1985).

A natural dimension for grouping households is their income. However, as the income may endogenously respond to the shock of interest, we cannot use the current household income as the grouping variable. Luckily, the LCFS does not only collect information about current household income but also about *normal* household income. This can be thought of as a proxy for permanent income.[10] Based on normal disposable household income, I group households into three pseudo-cohorts: low-income, middle-income, and high-income households. Following Cloyne and Surico (2017), I assign each household to a quarter based on the date of the interview, and create the group status as the bottom 25 percent of the normal disposable income distribution for low-income, the middle 50 percent for middle-income, and the top 25 percent for high-income in every quarter of a given year. The individual variables are then aggregated using survey weights to ensure representativeness of the British population.

Table 1 presents some descriptive statistics, overall and by income group. We focus here on expenditure excluding housing, however, the results including housing turn out to be similar. We can see that quarterly household expenditure is increasing in income. While low-income households spend a large part of their budget on non-durables, richer households spend more on durables. Importantly, poorer households spend a significantly higher share of their expenditure on energy: the energy share stands at almost 10 percent for low-income, just above 7 percent for middle-income, and around 5 percent for high-income households. Thus, to the extent that energy demand is inelastic, poorer households are more exposed to increases in energy prices.

[10] I have verified that normal income does not respond significantly to the carbon policy shock. In contrast, current income falls significantly and persistently, as shown in Figure B.12 in the Appendix. Alternatively, I group households by an estimate of permanent income obtained from a Mincerian-type regression. The results again turn out to be robust, see Appendix B.3.

Table 1: 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: The table shows 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 as a share of 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.

The different income groups turn out to be comparable in terms of their age. Higher-income households tend to be better educated, and are more likely to be homeowners, either by mortgage or outright.

5.2. Heterogeneity by household income

We are now in a position to study how households' expenditure and income respond to carbon policy shocks and, more importantly, how the response varies by income group. Figure 6 shows the responses of total household expenditure and current income for the three income groups we consider.[11] The solid black lines are again the point estimates and the dark/light shaded areas are 68 and 90 percent confidence bands.

[11] To eliminate some of the noise inherent in survey data, I smooth the expenditure and income measures with a backward-looking (current and previous three quarters) moving average, as in Cloyne, Ferreira, and Surico (2020). The results are robust to using the raw series instead (even though the responses become more jagged and imprecise) or using smooth local projections as proposed by Barnichon and Brownlees (2019), see Figure B.10 in the Appendix. To flexibly control for seasonal and trending behavior, I include a set of quarterly dummies and a linear trend.

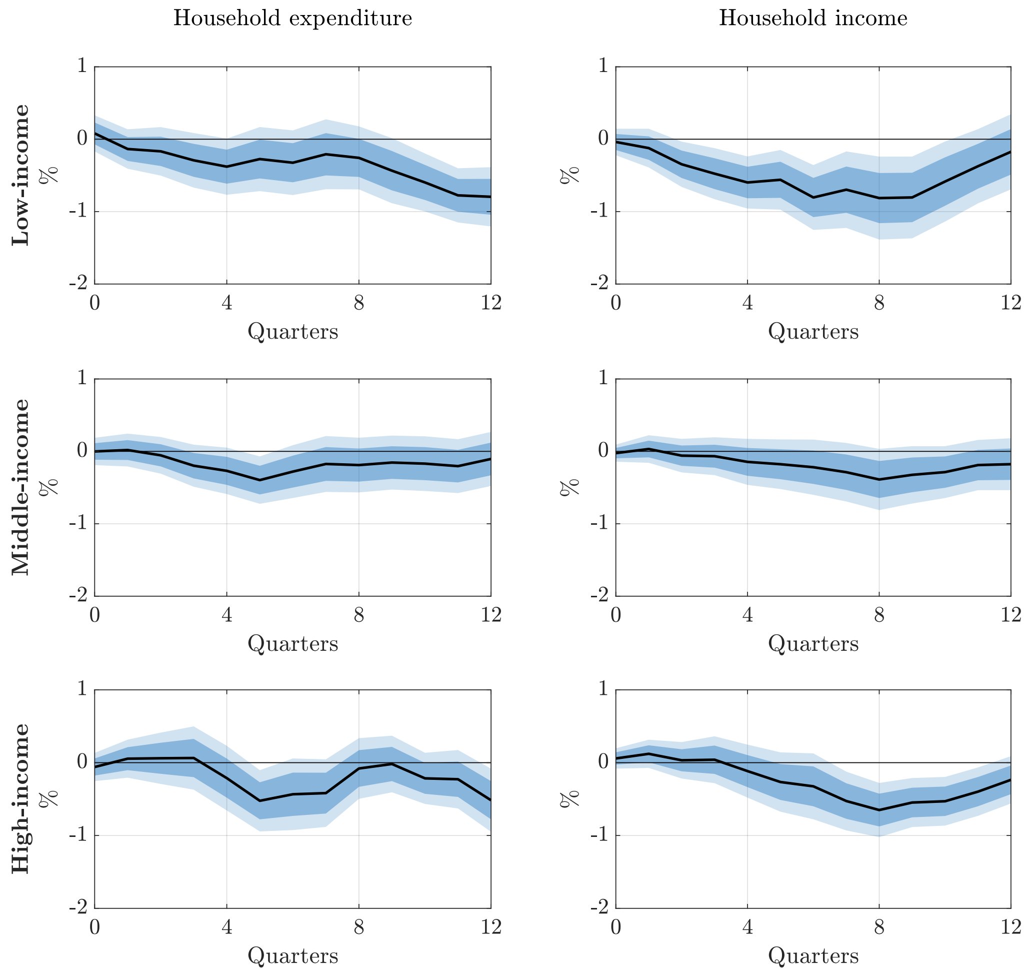

Figure 6: Household Expenditure and Income Responses by Income Group

Notes: Impulse responses of total expenditure (excluding housing) and current total disposable household income 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.

We can see that there is pervasive heterogeneity in the expenditure response across income groups. Low-income households reduce their expenditure significantly and persistently. In contrast, the expenditure response of higher-income households is rather short-lived and only barely statistically significant. This result is even more stark when we separate between different types of expenditure. Figure 7 shows the responses of energy, non-durable goods and services excluding energy, and durable goods expenditure. We can see that poor households substantially lower their non-durable expenditure while higher-income households display an insignificant response. For durable expenditures, the pattern is less clear cut. While low-income household cut durable expenditure, high-income households also display a significant response. Therefore, the heterogeneity in

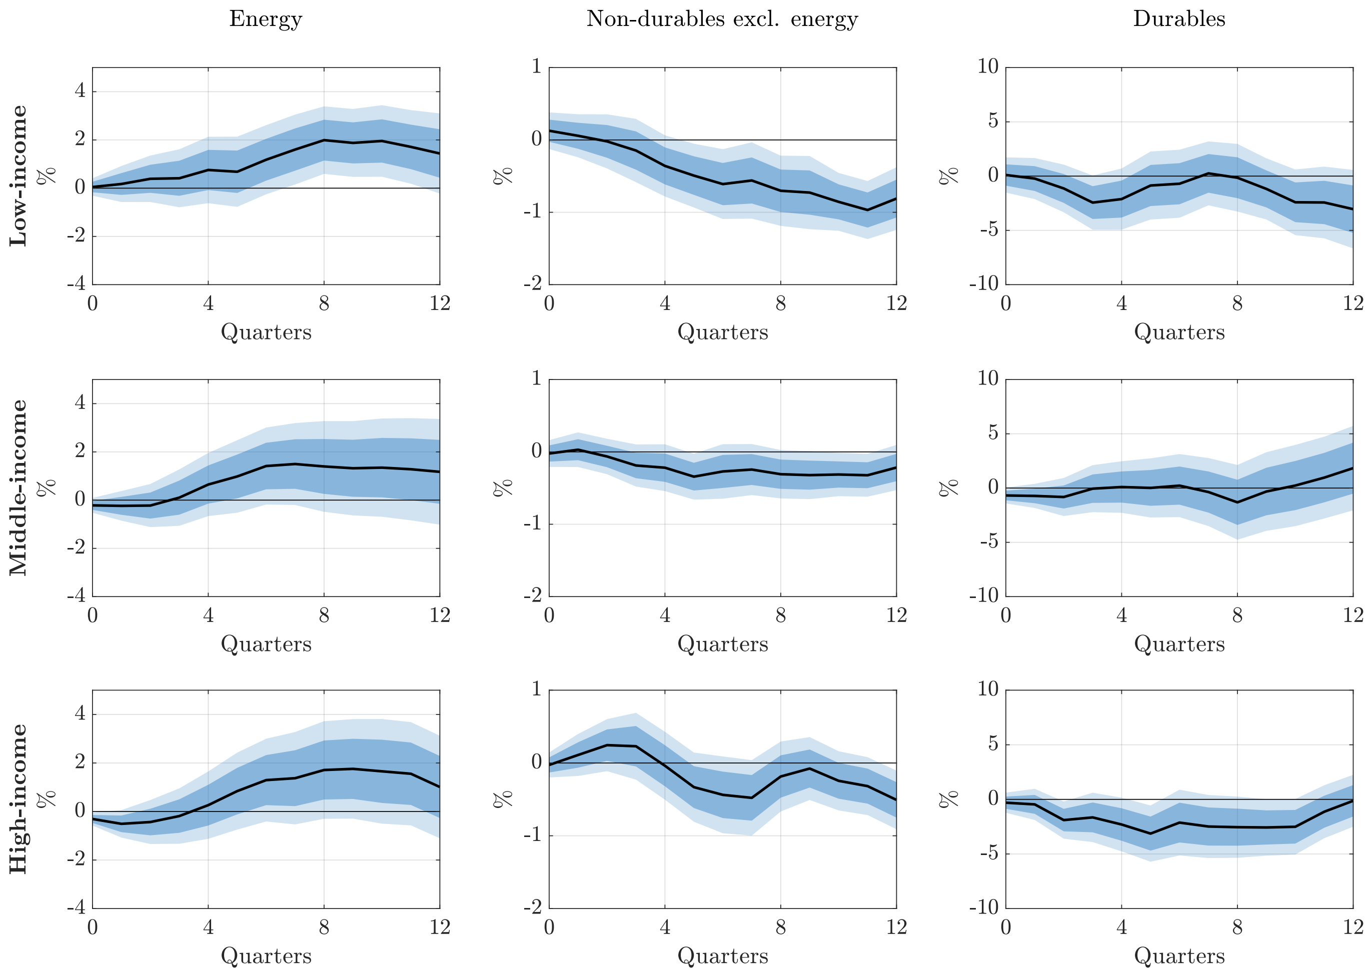

Figure 7: Energy, Non-durables and Durables Expenditure Responses by Income Group

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.

the total expenditure responses appears to be driven by the non-durable expenditure component. Note that these group-specific differences are not only economically but also statistically significant (see the responses of the group differences in Figure B.11 in the Appendix).

Low-income households are more affected in two ways. First, they face a larger and more significant increase in their energy bill. This is consistent with the fact that these households have a higher energy share to start with and their energy demand is particularly inelastic, for instance because of financial constraints. Second, looking at the income responses, we can also see that they face a more significant and substantial fall in their income. As we will see in Section 5.4, this appears to be driven by the fact that they tend to work in sectors that are more affected by the carbon policy shock. Taken together, the shock leads to a substantial reduction in discretionary income, which forces poorer households, who are also more likely to be financially constrained, to cut their expenditure by more.

At this stage, it is worth discussing a potential concern about grouping households concerning selection. The assignments into the income groups are not random and some other characteristics may, potentially, be responsible for the heterogeneous responses I document. To mitigate these concerns, I group the households by a selection of other grouping variables, including age, education and housing tenure. The results are shown in Figures B.14-B.16 in the Appendix. While there is not much heterogeneity by age, less educated households tend to respond more than better educated ones and social renters tend to respond more than homeowners. However, none of the alternative grouping variables can account for the patterns uncovered for income, suggesting that we are not spuriously picking up differences in other household characteristics.

5.3. Direct versus indirect effects

We have seen that there is substantial heterogeneity in the households' expenditure response to carbon policy shocks: while richer households change their expenditure only marginally, low-income households lower their expenditure significantly and persistently. Furthermore, indirect, general equilibrium effects via income and employment seem to play an important role in the transmission of the policy. To shed more light on the role of direct and the indirect effects, it is instructive to convert the responses into an equivalent pound change in income and expenditure over the three-year impulse horizon. This can be interpreted as the overall short-run monetary adjustment following the change in carbon policy.

Table 2 shows the results, overall and by income group. We can see that energy expenditure increases for all income groups, but only low-income households display a strong and significant increase. These households also cut their expenditure significantly, while the adjustment for higher-income households is less pronounced and not statistically significant. Importantly, the increase in energy bills cannot account for the large fall in non-energy expenditure. Note, however, that the shock also leads to a substantial fall in households' incomes, which is again particularly pronounced for low-income households. Coupled with the fact that these households are more likely to be financially constrained (see e.g Jappelli and Pistaferri, 2014), this helps explain the significant expenditure response.

Table 2: Cumulative Monetary Changes over Impulse Horizon

	Overall	By income group		
		Low-income	Middle-income	High-income
Expenditure				
Energy	23.88	28.36	22.53	22.11
	[-16.93, 64.69]	[8.21, 48.51]	[-18.02, 63.07]	[-0.96, 45.17]
Non-durables excl. energy	-103.75	-134.76	-92.33	-95.60
	[-212.38, 4.87]	[-241.21, -28.32]	[-192.67, 8.02]	[-279.87, 88.67]
Durables	-6.95	-2.92	-0.44	-23.99
	[-56.09, 42.20]	[-20.75, 14.92]	[-10.37, 9.50]	[-71.44, 23.45]
Income				
	-203.70	-214.90	-138.65	-322.60
	[-387.13, -20.27]	[-376.38, -53.41]	[-301.82, 24.52]	[-635.44, -9.77]

Notes: The table reports the overall pound change in expenditure and income over the three-year period following a carbon policy shock (in 2015 pounds). Bootstrapped 90 percent confidence intervals are reported in brackets. The overall pound change is computed as the present discounted value of the impulse response, multiplied by the corresponding average quarterly expenditure/income.

By contrast, high-income households also display a significant fall in their income, however, their expenditure responses turn out to be insignificant, consistent with the notion that these households are less financially constrained and thus better able to cushion the adverse effects on their income. Overall, these results point to an important role of indirect effects via income and employment. My estimates suggest that the direct effect through energy prices accounts for less than a third of the aggregate consumption response, as proxied by the increase in energy bills relative to the overall fall in expenditure (23.88/86.82).

The expenditure heterogeneity uncovered in this section is striking, especially against the backdrop that low-income households have much lower levels of expenditure to start with (see in Table 1). Put differently, low-income households account for over 30 percent of the aggregate effect of carbon pricing on consumption, despite the fact that they make up for a much smaller share of consumption in normal times (around 15 percent). Accounting for indirect, general equilibrium effects turns out to be crucial to correctly assess the distributional impact of carbon pricing. Focusing on the direct effect via the energy share alone can understate the actual distributional effects considerably.

The distributional consequences also likely play an important role for the magnitude of the aggregate expenditure response. My findings are consistent with a literature that emphasizes the role of MPC heterogeneity in combination with unequal income incidence for the transmission of aggregate demand shocks (Bilbiie, 2008; Auclert, 2019; Patterson, 2021, among others). These studies show in the context of aggregate-demand policies that the aggregate impact can be amplified when the policy disproportionately affects the incomes of individuals whose consumption is more sensitive. My results suggest that such a mechanism is also at play in the transmission of carbon pricing, following the initial fall in non-energy expenditure. Thus, even though low-income households only make up for a relatively small portion of the population, they play an important role for the transmission of the policy to the macroeconomy.

Alternative channels. Thus far, I focused my analysis on the direct effect via energy prices and the indirect, general-equilibrium effect via income. While there may also other channels at work, I briefly discuss here why these alternative channels are unlikely to play an important role in the transmission of carbon policy. First, carbon pricing may also have an effect on the prices of other goods via substitution effects, which may in turn affect households' budgets. However, as shown in Section 4.1, the response of core consumer prices is much more muted and only barely significant; therefore this channel does likely not play a major role. Second, there may be a number of channels that work through the response of durable expenditure, for instance because of uncertainty or precautionary motives, or via a reduction in durables that are complementary in use with energy (see also Edelstein and Kilian, 2009). However, the overall response of durable expenditure is quantitatively too small to play a dominant role in the transmission of carbon policy. Furthermore, in Section 4.3 I did not find any significant change in aggregate uncertainty after the shock. Finally, households may also adjust their saving behavior as interest rates increase in response to the shock. However, this channel is particularly relevant for higher-income households, which contribute relatively less to the aggregate consumption response.

5.4. What drives the income response?

We have seen that there is significant heterogeneity in the households' income responses. This section aims to shed light on what is driving the income incidence by household group. There are at least two potential sources of heterogeneity. First, households may differ in their labor income, for instance because they work in sectors that are differentially affected by the policy. Second, households may differ in their income composition, as some households also have substantial sources of financial income. I will focus here on the former, which is more relevant to understand the heterogeneity at the lower-end of the income distribution. In Appendix B.3, I also study the role of the household income composition.

To investigate into potential heterogeneities in labor income, I study how the responses vary by the sector of employment using data from the UK Labour Force Survey (LFS).[12] I consider two dimensions to group sectors. First, I group sectors by their energy intensity to gauge the role of the conventional cost channel. Second, I group sectors by how sensitive they are to changes in aggregate demand.[13]

Table 3: 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: The table depicts 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.

Table 3 presents descriptive statistics on the sectoral distribution of households, both overall and by income group. We can see that only few low-income households work in sectors with a high energy intensity such as utilities or man-

[12]Unfortunately, the LCFS does not include any information on the sector of employment. Therefore, I use data from the LFS which provides detailed information on employment sector and income. For more information on the LFS, see Appendix A.3.

[13]I measure the demand-sensitivity by estimating the elasticity of sectoral labor income to changes in aggregate income. Sectors that produce more 'discretionary' goods and services turn out to be more demand-sensitive. See Appendix B.3 for more information.

ufacturing. Thus, the sectors' energy intensity is unlikely to explain the heterogeneous income responses that we observe. A more relevant dimension of heterogeneity appears to be the sectors' demand sensitivity: low-income households work disproportionally in sectors that tend to be more sensitive to aggregate fluctuations, such as retail or hospitality, while a large majority of higher income households work in less demand-sensitive sectors.

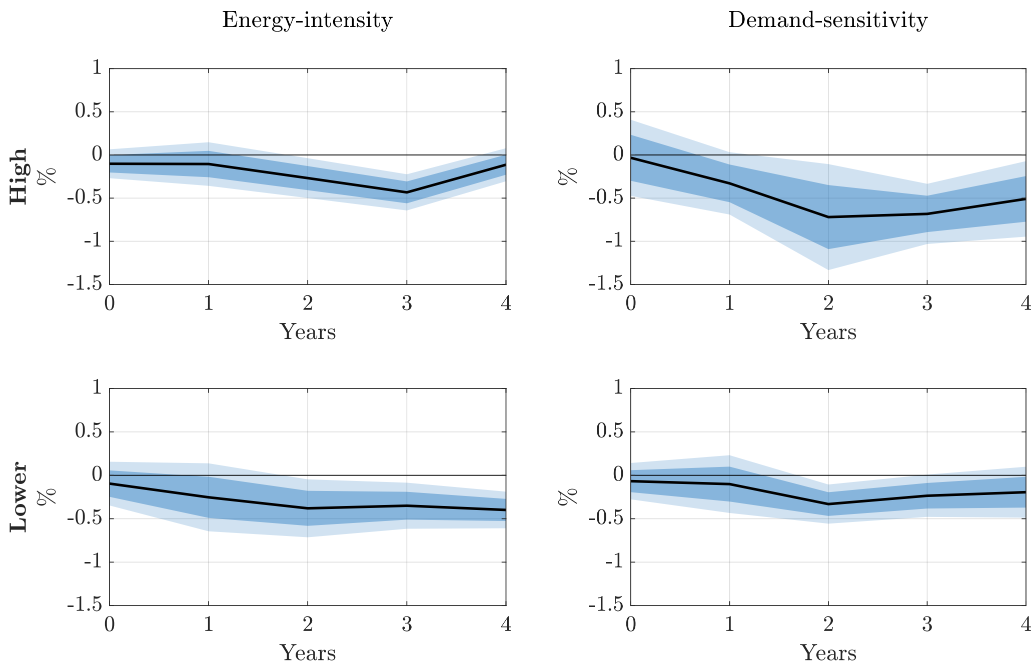

Figure 8: Income Response by Sector of Employment

Notes: Impulse responses of income (pay from main and second job net of deductions and benefits) in different sectors, grouped by their energy-intensity and demand-sensitivity. The response is computed based on the median income in the respective group of sectors. The sector groups are described in detail in Table 3.

Figure 8 shows how the median income across different sectors changes after a carbon policy shock. While sectors with a high energy intensity display a somewhat stronger fall in incomes than sectors with a lower intensity, the differences are not that large quantitatively. By contrast, there is significant heterogeneity by the sectors' demand-sensitivity: households working in demand-sensitive sectors experience the largest and most significant fall in their income while households in less-demand sensitive sectors display more muted income responses. This helps explain the observed heterogeneity in the income responses. In response to a carbon policy shock, these sectors face a stronger decrease in demand, also because households cut expenditure more in these sectors, and thus react by laying off employees and cutting compensation. As low-income households are overrepresented in these sectors, they are disproportionally affected.

These results further support the notion that carbon policy shocks strongly transmit to the economy not only through the traditional cost channel but also through the demand side of the economy, in line with previous evidence by Kilian and Park (2009) on the transmission of energy price shocks. A novel insight is that in the presence of household heterogeneity, the demand channel may be even stronger. This result speaks directly to a growing literature on the role of Keynesian supply shocks (see e.g. Guerrieri et al., 2022; Cesa-Bianchi and Ferrero, 2021).

5.5. The role of redistributing carbon revenues

We have seen that the economic costs of carbon pricing are borne unequally across society. Low-income households are the most affected, having to reduce their expenditures the most, and are contributing disproportionally to the aggregate response. A key question in this context is how the distribution of carbon revenues matters for the transmission of the policy. Since auctioning became the default way of allocating allowances, the system produces a growing share of auction revenues. However, there is no direct redistribution scheme in place that could offset the distributional effects on households that I document.[14] The large majority of revenues are earmarked and used for climate and energy related purposes.

While using the carbon revenues for climate purposes may help to further propel emission reductions, my results indicate that redistributing part of the revenues to the most affected groups in society could mitigate the distributional effects and reduce the economic costs of climate policy. To the extent that energy demand is inelastic, which turns out to be particularly the case for low-income households, this should not compromise the reductions in emissions.

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 by Golosov et al. (2014) with nominal rigidities and household hetero-

[14]For the period from 2012-2020, the revenues generated by the member states of the EU ETS exceeded 57 billion euros (European Comission, 2020*b*). The current ETS does not feature a direct redistribution scheme, however, there are certain other, indirect solidarity measures in place, e.g. via the Cohesion Fund or the Just Transition Fund. Only in the recent 'Fit for 55' plan, the European Commission takes a step in this direction by proposing a Social Climate Fund for the new ETS in transportation and buildings.

geneity, as in Bilbiie, Känzig, and Surico (2021), to allow for the demand channels identified in the data. I will only sketch the relevant parts of the model here, a full description can be found in Appendix D.

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. 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. I assume that only bonds are liquid and can be used to self-insure. There is limited asset market participation. Only savers are able to self-insure themselves using liquid bonds.[15] They choose their consumption intertemporally, according to the following Euler equation:

$$\frac{U_x(x_{S,t}, h_{S,t})}{p_{S,t}} = \beta\, \mathbb{E}_t \left[\frac{R_t^b}{\Pi_{t+1}} \left(s \frac{U_x(x_{S,t+1}, h_{S,t+1})}{p_{S,t+1}} + (1-s) \frac{U_x(x_{H,t+1}, h_{H,t+1})}{p_{H,t+1}} \right) \right], \quad (9)$$

where $x_{i,t}$ is total consumption of household i, $h_{i,t}$ is labor supply, U_x is the marginal utility, $\frac{R_{t-1}^b}{\Pi_t}$ is the real risk-free rate, $p_{i,t}$ is the price index of the household's consumption basket, and $1-s$ is the transition probability of becoming hand-to-mouth. The demand for non-energy and energy goods is given by the following schedules: $c_{S,t} = a_{S,c} \left(\frac{1}{p_{S,t}}\right)^{-\epsilon_x} x_{S,t}$ and $e_{S,t} = a_{S,e} \left(\frac{p_{e,t}}{p_{S,t}}\right)^{-\epsilon_x} x_{S,t}$, where ϵ_x is the elasticity of substitution between non-energy and energy goods. Savers also invest in capital k_t and supply labor $h_{S,t}$. The corresponding optimizing equations are standard and relegated to the Appendix. Savers receive labor income, financial income from dividends and capital returns, and transfers $\omega_{S,t}$ from the government.

Hand-to-mouth households have no assets and thus consume all of their income in every period: $p_{H,t} x_{H,t} = y_{H,t}$. Their income $y_{H,t}$ consists of labor income plus government transfers, $\omega_{H,t}$. The non-energy and energy demand functions, and labor supply equation are analogous to the expressions for the savers.

[15] 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 (2021) for a detailed discussion.

The firm block of the model consists of two sectors: energy and non-energy producers. Energy firms produce energy using labor as an input, and can adjust their prices flexibly. Their production technology is given by

$$e_t = a_{e,t} h_{e,t}, \tag{10}$$

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. Energy firms are subject to a carbon tax τ_t. This conforms well with my empirical analysis, where I study the impacts of plausibly exogenous changes in carbon prices. The optimal energy supply is characterized by $w_t = (1-\tau_t) p_{e,t} \frac{e_t}{h_{e,t}}$.

The non-energy sector consists of standard New Keynesian firms that produce non-energy goods using capital, energy, and labor as inputs and set prices subject to nominal rigidities. Their production technology is given by

$$y_t = e^{-\gamma s_t} \left[(1-\nu)^{\frac{1}{\epsilon_y}} \left(a_t k_t^{\alpha} h_{y,t}^{1-\alpha} \right)^{\frac{\epsilon_y - 1}{\epsilon_y}} + \nu^{\frac{1}{\epsilon_y}} \left(e_{y,t} \right)^{\frac{\epsilon_y - 1}{\epsilon_y}} \right]^{\frac{\epsilon_y}{\epsilon_y - 1}}, \tag{11}$$

where $e^{-\gamma s_t}$ captures climate damages, modeled as a function of the atmospheric carbon concentration s_t. 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 $v_{1,t}$ and $v_{2,t}$ are auxiliary terms given in the Appendix. The price setting problem gives rise to a standard Phillips curve, which in log-linear form reads $\hat{\pi}_t = \kappa \widehat{mc}_t + \beta E_t \hat{\pi}_{t+1}$, where hatted variables denote log-deviations from steady state.

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.

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

$$\lambda \omega_{H,t} = \mu \tau_t p_{e,t} e_t \quad \text{and} \quad (1-\lambda) \omega_{S,t} = (1-\mu) \tau_t p_{e,t} e_t. \tag{12}$$

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. I calibrate the model using macro and micro moments from the data and drawing on values previously used in the literature. I discuss the calibration in detail in Appendix D.7.

Model evaluation. The impulse responses to a carbon policy shock, normalized to match the estimated peak energy price response, are shown in Figure 9. In what follows, we focus on the peak responses, as the model is not designed to match the hump-shaped responses in the data. We can see that 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.

The monetary response turns out to be an important factor for the transmission of carbon policy shocks. Recall that we assume that monetary policy targets headline inflation and thus leans against the inflationary pressures emerging from the increase in carbon prices, consistent with the monetary policy response estimated in the data. If we assume that monetary policy targets core inflation instead, the effects of carbon policy are attenuated. Household heterogeneity acts as a further amplifying channel through the unequal income incidence of the shock linked to the heterogeneity in MPCs. Without these demand channels, it is difficult to match the empirical magnitudes unless the energy share is set to implausibly high levels, see Appendix B.4.

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 9 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).

We can see that 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 Appendix B.4. The intuition is that the redistribution 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 reduc-

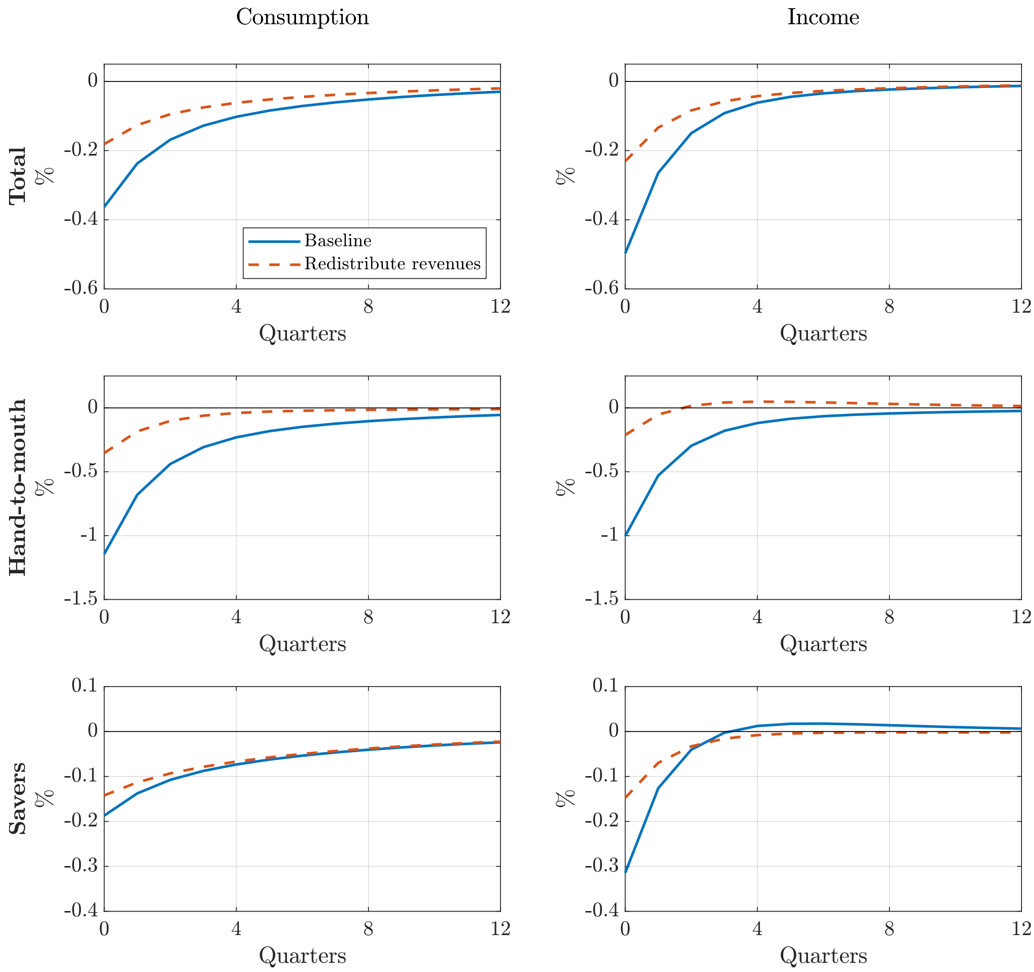

Figure 9: Model Responses for Consumption and Income

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 match the peak energy price response in the data. 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.

tion in consumption inequality. Emissions on the other hand change by less as low-income households' energy demand is particularly 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.

6. Beyond the Short Term

We have seen that carbon pricing is successful in reducing emissions but this comes at an economic cost, at least in the short term. This section aims to shed light on some of the longer-term implications, specifically the impact of carbon pricing on public attitudes towards climate policy and the effects on green innovation.

Attitudes toward climate policies. An important argument for cushioning the distributional impact is that a successful transition to a low-carbon economy requires public support. If certain groups feel left behind, this could undermine the success of climate policy as the yellow vest movement in France, which started as a protest against higher fuel taxes, has shown for instance (Knittel, 2014).

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 and is an important barometer of public attitudes in the UK. To proxy attitudes towards climate policy, I rely on a question that elicits the approval rate for environmentally-motivated fuel taxes (see Appendix B.3 for more information).

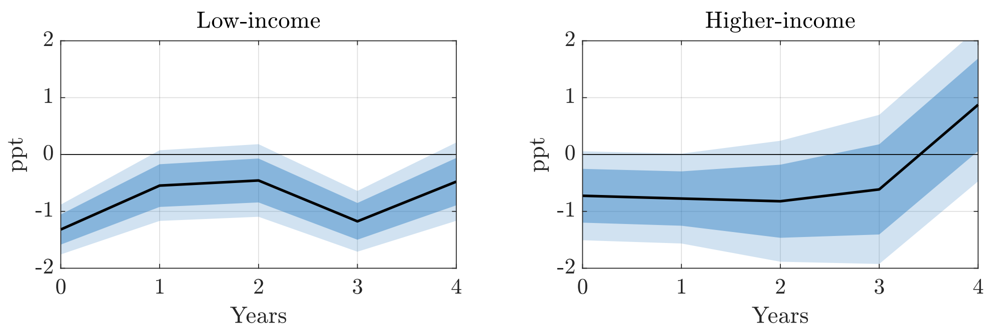

Figure 10: Effect on Attitude Towards Climate Policy

Notes: Impulse responses of public attitude towards climate policy for low- and higher-income groups. The public attitude towards climate policy is proxied by the share of households in the British social attitudes survey that express support for environmentally-motivated fuel taxes. Low-income correspond to the bottom 25 percent and higher-income to the other 75 percent of the income distribution.

Figure 10 shows the response of the approval rate of environmentally-motivated tax policies to a carbon policy shock across income groups. While the response of higher-income households is barely significant and even turns positive at longer horizons, low-income households display a significant and persistent fall in the support of climate policies. Recall, these households are also the ones that are most adversely affected by carbon policy shocks. These results suggest that compensating the most affected households may help increase the public support of climate change mitigation policies – consistent with recent evidence by Anderson, Marinescu, and Shor (2019) and Dechezleprêtre et al. (2022).

The impact on green innovation. A key motivation behind carbon pricing is to create an incentive for directed technical change. In fact, part of the vision for the EU ETS is to promote investment in clean, low-carbon technologies (European Comission, 2020*a*). Innovation in low-carbon technologies will be crucial to sustain emission reductions without permanently lowering output.

To analyze this channel empirically, I study how the patenting activity in climate change mitigation technologies changes in response to carbon policy shocks. I use data on patent applications from the European Patent Office (EPO), which has developed specific classification tags for patents in climate change mitigation technologies.

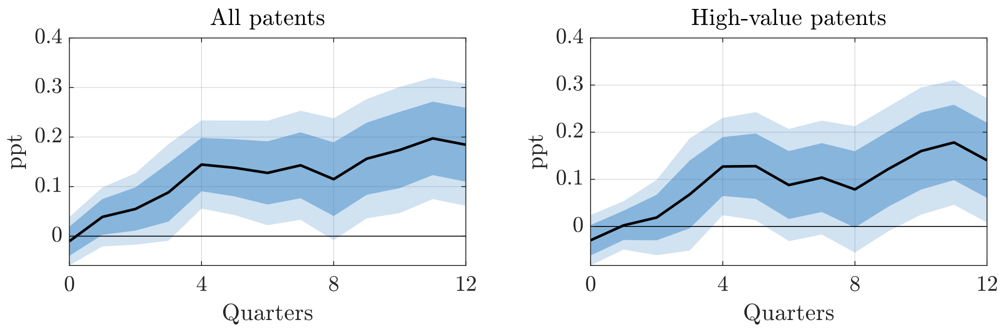

Figure 11: Patenting in Climate Change Mitigation Technologies

Notes: Impulse responses of patenting activity in climate change mitigation technologies, as measured by the number of climate change mitigation patents as a share of all patents filed at the EPO. The left panel displays the share based on all patents while the right panel focuses on high-value patents, i.e. patents filed at multiple patent offices.

The results are shown in Figure 11. We can see that the shock leads to a significant increase in low-carbon patenting, and this is robust to focusing on high-value patents. The effect is also economically significant as the average share of climate change mitigation patents is around 10 percent. Thus, carbon pricing appears to be successful in stimulating green innovation. These results support the findings of Calel and Dechezleprêtre (2016), who employ a quasi-experimental design exploiting inclusion criteria at the installations level to estimate the causal impact of the EU ETS on firms' patenting.

7. Conclusion

Fighting climate change is one of the greatest challenges of our time. While it has proved to be difficult to make progress at the global level, several national carbon pricing policies have been put in place. However, still little is known about the effects of these policies on emissions and the economy. This paper provides new evidence from the largest carbon market in the world, the EU ETS. I show that tightening the carbon pricing regime leads to a significant increase in energy prices, a persistent fall in emissions and an uptick in green innovation. This comes at the cost of temporarily lower economic activity and higher inflation. Importantly, these costs are borne unequally across society. Poorer households lower their consumption significantly and are driving the aggregate response while richer households are less affected. Not only are these households more exposed to carbon pricing because of their higher energy expenditure share, they also experience a larger fall in their income. These indirect effects via income and employment turn out to be quantitatively important. My results suggest that redistributing some of the carbon revenues to the most affected groups can reduce the economic costs of carbon pricing and may help strengthen the public support of the policy. In future work, it would be interesting to better understand how climate, fiscal and monetary policy can be coordinated to organize a successful transition to a low-carbon economy.

References

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.

Annicchiarico, Barbara and Fabio Di Dio. 2015. "Environmental policy and macroeconomic dynamics in a new Keynesian model." *Journal of Environmental Economics and Management*, 69: 1–21.

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

Baqaee, David Rezza and Emmanuel Farhi. 2019. "The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem." *Econometrica*, 87(4): 1155–1203.

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.

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

Bernanke, Ben S., Mark Gertler, and Mark Watson. 1997. "Systematic monetary policy and the effects of oil price shocks." *Brookings Papers on Economic Activity*, 1997(1): 91–157.

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."

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

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

Bilbiie, Florin O., Diego R. Känzig, and Paolo Surico. 2021. "Capital and income inequality: An aggregate-demand complementarity."

Blanchard, Olivier J. and Roberto Perotti. 2002. "An empirical characterization of the dynamic effects of changes in government spending and taxes on output." *The Quarterly Journal of Economics*, 117(4): 1329–1368.

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

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."

Cloyne, James. 2013. "Discretionary tax changes and the macroeconomy: new narrative evidence from the United Kingdom." *American Economic Review*, 103(4): 1507–28.

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.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2022. "Fighting climate change: International attitudes toward climate policies." National Bureau of Economic Research.

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

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

European Comission. 2020*b*. "Report on the functioning of the European carbon market." *https://ec.europa.eu/clima/sites/clima/files/news/docs/com_2020_740_en.pdf*, [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.

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 and Marc Hafstead. 2018. *Confronting the Climate Challenge.* Columbia University Press.

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.

Hamilton, James D. 2008. "Oil and the macroeconomy." *The New Palgrave Dictionary of Economics*, 2.

Hamilton, James D. 2009. "Daily changes in fed funds futures prices." *Journal of Money, credit and Banking*, 41(4): 567–582.

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.

Jappelli, Tullio and Luigi Pistaferri. 2014. "Fiscal policy and MPC heterogeneity." *American Economic Journal: Macroeconomics*, 6(4): 107–136.

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.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): 1589–1610.

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

Kaplan, Greg and Giovanni L. Violante. 2014. "A model of the consumption response to fiscal stimulus payments." *Econometrica*, 82(4): 1199–1239.

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

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.

Konradt, Maximilian and Beatrice Weder di Mauro. 2021. "Carbon Taxation and Inflation: Evidence from the European and Canadian Experience."

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.

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

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.

McKibbin, Warwick J., Adele C. Morris, Augustus Panton, and Peter J. Wilcoxen. 2017. "Climate change and monetary policy: Dealing with disruption."

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*a*. "The Macroeconomic Impact of Europe's Carbon Taxes." *NBER Working Paper*.

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

Miranda-Agrippino, Silvia and Giovanni Ricco. 2018. "Identification with external instruments in structural VARs under partial invertibility."

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. 2020. "Local Projection Inference is Simpler and More Robust Than You Think."

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.

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

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.

Patterson, Christina. 2021. "The matching multiplier and the amplification of recessions."

Piazzesi, Monika and Eric T. Swanson. 2008. "Futures prices as risk-adjusted forecasts of monetary policy." *Journal of Monetary Economics*, 55(4): 677–691.

Plagborg-Møller, Mikkel and Christian K. Wolf. 2019. "Local projections and VARs estimate the same impulse responses."

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

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. 2016. "Macroeconomic shocks and their propagation." In *Handbook of Macroeconomics*. Vol. 2, 71–162. Elsevier.

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.

Romer, Christina D. and David H. Romer. 2010. "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks." *American Economic Review*, 100(3): 763–801.

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.