Carbon Pricing and Inequality: A Normative Perspective*

Saki Bigio
UCLA

Diego R. Känzig
Northwestern University

Pablo Sanchez
Northwestern University

Conor Walsh
Columbia University

March, 2025

Abstract

Despite broad acceptance among economists, carbon taxes face persistent public resistance. We measure the sources and distribution of welfare losses from unexpected European carbon price changes by estimating their impact on consumer prices, labor income, financial wealth, and government transfers. A 1% carbon-policy-induced increase in energy prices yields an average welfare loss of about 0.5% of annual consumption, primarily driven by indirect labor-income effects. Younger, poorer, and less educated households, especially in Southern and Eastern Europe, bear a disproportionate burden. These findings suggest public opposition to carbon taxes stems from legitimate distributional concerns.

JEL classification: D31, H23, Q58

Keywords: Carbon pricing, inequality, welfare

*Bigio: UCLA and NBER (email: sbigio@econ.ucla.edu); Känzig: Northwestern University, CEPR and NBER (email: dkaenzig@northwestern.edu); Sanchez: Northwestern University (email: pablo.sanchez@u.northwestern.edu); Walsh: Columbia University (email: conor.walsh@columbia.edu). We thank Alberto Bisin for helpful comments and suggestions.

1. Introduction

Climate change is perhaps humanity's greatest challenge. It is also a manmade challenge, driven by carbon emission externalities that can be undone. Economics offers a century-old fix for climate change: carbon taxes. While widely accepted in policy circles, practical attempts to implement carbon taxes have been consistently met with fierce resistance.[1] Why does a tax meant to solve one of humanity's greatest threats provoke such push-back? Are the average welfare losses from this specific tax unusually large, or does it disproportionately burden certain groups?

This paper attempts to quantify both the sources and the distribution of the welfare impacts of a carbon tax. By providing quantitative answers to the above questions, it also suggests strategies to mitigate resistance to carbon taxation.

To measure the sources and distribution of carbon-tax welfare effects, we adopt the "feasible set approach" developed in Del Canto et al. (2023). The main insight behind this technique is to apply the Envelope Theorem on a grand scale: the first-order impact of an unexpected tax change on any household is summarized by the change in the discounted present value of their future budgets. In turn, the effects on these budgets can be measured by estimating how the policy alters the consumption-basket prices, labor income, financial wealth, and government transfers. By measuring these budget changes, we obtain a "money-metric" welfare loss for each demographic of interest. A key advantage of this metric is that it is interpretable because it is expressed in income units. Moreover, it remains independent of specific preference assumptions, encompassing any theory in which we only require agents to optimize subject to budget constraints.

Although applying this framework is straightforward, these steps demand artisanal empirical work. First, we must gather historical panel data on consumption-basket indices, labor income, financial wealth, and transfers across different demographics. Next, we estimate how *unanticipated* changes in carbon prices affect these variables over time. That step requires finding a valid instrument for the tax and then computing the impulse response functions (IRFs) for each variable for each demographic group. The last step requires identifying an appropriate discount factor to convert those impulse responses into a welfare loss by income group.

We focus our measurements on Europe because, to date, it is the region of the globe that has advanced the furthest in putting a price on carbon. Specifically, we compile

[1] Notable examples of quick repeals of carbon taxes include Australia in the early 2010s, the repeal in France in 2018 after the "yellow vest" protests, the failure of the Waxman-Markey bill in the United States, and, more recently, the resistance in three Canadian provinces to federal carbon taxes.

a new dataset combining rich information from three major European household surveys from Eurostat and the ECB. We then use this data to construct IRFs for the different variables for various demographic groups—four age groups, two education levels, three income brackets, and four geographic regions. As an instrument for unexpected carbon-tax changes, we follow Känzig (2023). That paper exploits high-frequency carbon futures price shifts around regulatory changes in the European carbon market to identify unexpected shocks to carbon prices. With those shocks, we estimate the IRFs of our variables of interest using local projections.

Our main findings are as follows: First, the overall welfare losses are unexpectedly large: a carbon-policy induced increase in energy prices by 1% yields an average welfare loss of 80 euros, which is about 0.5% of one year's consumption.

Second, when we decompose the sources of these losses, the direct effect—from higher consumption-basket prices—amounts to 0.2%. Among the indirect effects (labor income, asset prices, and transfers), the respective impacts are 0.45%, -0.05%, and -0.1% euros, making labor income the dominant contributor.

Third, regarding distribution, direct effects on consumption-basket prices appear similar across income groups. However, the indirect channels disproportionately affect younger and poorer households, mainly through forgone labor income—similarly, those without a college education shoulder the highest burden. By contrast, retirees gain from carbon taxes, likely due to concurrent asset-price appreciations.

We also document substantial regional heterogeneity, with Southern and Eastern Europe shouldering the most significant welfare losses. The indirect effects via labor markets become the key driving force, consistent with the notion that labor markets in this region tend to be more rigid due to stronger employment protection laws and less flexible wage-setting institutions.

Related literature. This project contributes to a burgeoning literature on the impacts of carbon taxes and pricing. The literature has studied the effects of carbon pricing on emissions (Martin, De Preux, and Wagner, 2014; Andersson, 2019, among others), economic activity (Metcalf, 2019; Bernard and Kichian, 2021; Metcalf and Stock, 2023; Andersson, 2019; Känzig, 2023), and inflation (Konradt and Weder di Mauro, 2021; Bettarelli et al., 2025).

Our paper focuses squarely on inequality, as does a smaller subset of papers motivated by the same concerns. Andersson and Atkinson (2020), for example, studies the distributional effects of a sequence of carbon tax increases in transport fuel across different income groups. Beznoska, Cludius, and Steiner (2012) use input-output tables to project increases in carbon taxes to electricity prices to different industry prices and present incidence measures on different consumers using German survey data. The approach in those papers is to consider estimating the direct effect on consumption-basket prices and, following Poterba (1989), to approximate lifetime effects using current expenditures.

Our methodology differs for three reasons: first, we consider surprise increases by instrumenting for shocks. Second, we also consider indirect effects (on income, asset prices, and transfers). Finally, we compute dynamic effects directly. Thus, a first contribution relative to this literature is the computation of IRF functions to unexpected carbon price shocks of various budget-relevant variables for various demographics. Whereas the methodology to compute impulse responses is borrowed from Känzig (2023), we extend the estimates in that paper, which focuses on the UK, to encompass all European countries here.

An important novel finding from our approach is that carbon price increases predominantly impact lower-income households, but predominantly through the indirect income channel, which is not measured in studies that consider only direct effects. This finding is important because it corroborates recent survey evidence (see e.g. Drews and Van den Bergh, 2016; Dechezleprêtre et al., 2022).[2] Our quantitative finding suggests that surveys capture more than just perceptions and that the people who are most affected by the policies to mitigate climate change are the poor. In that sense, we also contribute to a broader literature in public finance other Pigouvian taxes (e.g. Allcott, Lockwood, and Taubinsky, 2019).

A second contribution is translating our IRFs into welfare metrics for carbon price shocks. We apply the feasible set approach in Del Canto et al. (2023), leveraging the estimation approach in Känzig (2023). This approach relates to several recent papers studying the welfare effects of price movements (Baqaee and Burstein, 2023; Dávila and Schaab, 2022). Baqaee, Burstein, and Mori (2022) and Jaravel and Lashkari (2022) provide methods to estimate the changes in consumer money-metric welfare over time under nonhomothetic preferences. These studies focus on long-run welfare changes. The approach here focuses on short- to medium-term impacts on money-metric welfare. To our knowledge, we are the first to empirically estimate the incidence of carbon pricing on welfare. While we focus on Europe, our findings are highly relevant to other developed countries.

Our IRFs and welfare loss estimates can be used as calibration inputs in various quantitative applications. For example, they should be used in quantitative public finance analysis. The theoretical foundation of carbon taxes rests on targeting principle (Sandmo,

[2]However, survey evidence from Sweden indicates lesser distributional effects (Ewald, Sterner, and Sterner, 2022).

1975; Dixit, 1985). This principle states that one should always correct an externality by targeting its direct sources. This result extends to settings where distributional effects are present regardless of the available tax instruments (Kopczuk, 2003).[3] Yet, public finance does emphasize that any undesirable distributional consequences should be offset, but to find the proper compensation, we need estimates such as the ones provided in this paper.

A second use for the estimates here is as moment targets for quantitative-structural climate change models. The IRFs we report here can be used to discipline a structural model that can speak to the longer-term impacts of carbon pricing and perform counterfactual analysis (e.g., Belfiori and Macera, 2025).

2. A Feasible Set Approach to Welfare

Our methodology builds on the feasible set approach put forward in Del Canto et al. (2023). The key idea is that, up to first order, the welfare impact of a shock on households is summarized by the induced movements in their feasible sets: their budget constraint and borrowing constraints. We outline the approach below leaving the formal setup and derivations to Appendix A.

The starting point is the household's expenditure and savings problem.[4] Households maximize their lifetime utility over consumption, labor, and asset holdings. They consume a basket of J consumption goods, with expenditure $p_{jt}c_{j,t}$ at time t for good j. They earn (after-tax) labor income W_tL_t, deciding on L_t. In addition, they may receive some transfer income T_t. Lastly, they have a dynamic portfolio of K assets, chosen optimally, each paying a dividend D_{kt} and fetching a price Q_{kt} at time t. These assets are varied: equities, debt, housing, cash, or durable goods, for example. Some of these assets could directly generate utility in addition to yields. One asset is a risk-free one-period debt, which is used to discount future utility streams.

Our notion of welfare, following a long literature in public finance, is money-metric utility. We are interested in how many euros each individual would be willing to pay not to be exposed to the fundamental carbon price shock. Up to the first order, the money-metric welfare gain, denoted dV, is given by:

[3] Bovenberg and De Mooij (1994), Bovenberg and Goulder (1996) and Fullerton (1997) consider environmental taxes in general equilibrium contexts but under restricted sets of instruments.

[4] To simplify the notation, we suppress indexing of the budget set variables by individual. However, everything below should be understood to vary across individuals; people who differ by age, income, and education will have different consumption plans, asset accumulation paths, and labor income processes.

$$dV = \sum_t R_{0\to t}^{-1}\Bigg(\underbrace{-\sum_j p_{jt}c_{jt}\Psi_t^{p,j}}_{\text{Consumption Price Changes}} + \underbrace{W_tL_t\Psi_t^W}_{\text{Labor Income Changes}} + \sum_k \Big[\underbrace{N_{kt-1}D_{kt}\Psi_t^{D,k}}_{\text{Asset Income Changes}} - \underbrace{Q_{kt}\Delta N_{kt}\Psi_t^{Q,k}}_{\text{Asset Price Changes}} \Big] + \underbrace{T_t\Psi_t^T}_{\text{Transfer Income Changes}} \Bigg), \tag{1}$$

where $R_{0\to t}$ denotes the cumulative risk-free interest rate from time 0 to t, and $\mathbf{\Psi}_t^i$ is the impulse response of price $i \in \{p, W, D, Q, T\}$ to the carbon policy shock in. Namely, each element of this vector represents the log deviation of a variable from its pre-shock trend in period t after a shock to the carbon price in period 0.

The welfare impact of a shock to the carbon price consists of five effects captured by each term inside the parenthesis. The first term captures the rise in the prices of the consumption goods in the agent's basket from the carbon tax. The welfare impact of this price rise, in money-metric terms, weights the percentage change in the price of each good by the nominal expenditure on the good in a world absent the shock—this term has been the predominant focus of the literature following Poterba (1989). The second term captures how the indirect effect on wages or labor income stemming from general equilibrium effects. Again, the welfare loss of these changes is summarized by weighting the impulse response of wages by total nominal labor income absent the shock. The third is the analog on capital income. This term includes changes in dividends or interest rates. It also includes asset price changes. As pointed out by Fagereng et al. (2024), the welfare impact of this change only depends on whether the agent is a *net buyer* or a *net seller* of the asset in a particular period.[5] The fourth term captures changes in transfer schemes, directly or indirectly, through indexed pension payments.

We can obtain empirical counterparts to this decomposition following two steps: First, we estimate the response of the product and factor prices to shocks to the carbon price, using the methodology from Känzig (2023). We call the vector of these impulse response functions $\mathbf{\Psi}_t$. These are combined with life-cycle statistics on consumption expenditure ($p_{jt}c_{jt}$), labor income profiles (W_tL_t), asset income ($N_{kt-1}D_{kt}$) and asset accumulation profiles ($Q_{kt}\Delta N_{kt}$) and transfer profiles (T_t).

[5]The logic of this can be understood with an example. Suppose an older Spanish household is planning to sell down their equity holdings to fund their retirement. When the carbon price rises, equity prices fall. This causes a welfare loss in the period they were planning to sell their holdings, as they receive less cash and can purchase fewer consumption goods. The opposite logic holds for a younger household in the accumulation phase of life.

The advantage of this semi-structural approach is that we can be agnostic about many of the primitives of the household's problem. The result, however, leverages the envelope theorem and is thus only accurate for relatively small movements in carbon prices. However, given that the shocks we consider are relatively small, this approach yields valuable insights. It is also important to note that up to a first order, behavioral responses are irrelevant for welfare, which is why no behavioral responses appear in (1). Lastly, the formula is only accurate in a world of both small aggregate and idiosyncratic risk. Extending it to handle general idiosyncratic risk is straightforward, and we discuss in the Appendix.

Lastly, we note that the *benefits* of reduced carbon emissions appear nowhere in the welfare calculation. This is because we are only interested in the overall and distributional pecuniary costs of carbon taxes. The benefits of a better climate or disaster risk reductions do not appear here. The formula, however, may be useful to test if carbon taxes are set appropriately, provided that we have a measure of these benefits obtained elsewhere. Whether carbon taxes are set appropriately can be thought of in the following way: on the margin, an additional euro of a carbon tax must balance of the direct economic welfare effects in (1), with the additional reduction in CO2, multiplied by the marginal welfare gains of this reduction.

3. Data and Identification

We focus on measuring the welfare costs of carbon prices on European households. The EU is widely recognized as a leader in climate policy due to its comprehensive and ambitious strategies. The cornerstone of the EU's climate policy is the European Union Emission Trading System (EU ETS). It is the world's first major carbon market and one of the largest markets. Established in 2005, it covers more than 11,000 heavy energy-using installations, accounting for around 40 percent of the EU's greenhouse gas emissions.

Identification. Carbon prices in the EU ETS are market prices, determined by the supply of emissions allowances and the demand for emissions. As with any demand and supply system, identifying shocks to supply and demand is subject to identification challenges. Thus, projecting carbon prices onto our outcome variables of interest will lead to biased estimates of the effects coming from policies that limit emissions. For example, carbon prices can decline in recessions, precisely when household income is falling. Thus, to estimate the dynamic causal effects, i.e., our IRF inputs, we need to isolate plausibly exogenous variation in carbon prices—specifically, changes in carbon prices driven by policy changes affecting the supply of emission allowances.

Figure 1: Carbon policy shocks

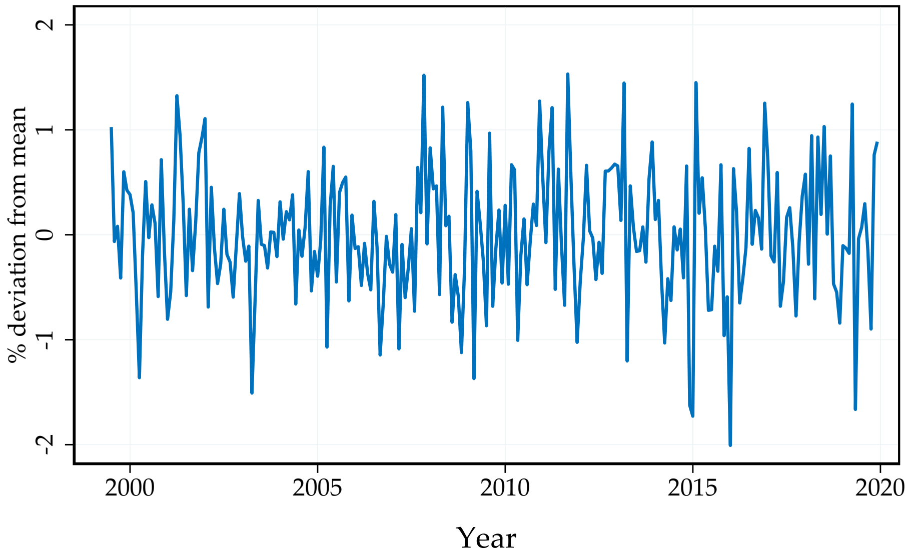

Notes: The monthly carbon policy shock series from Känzig (2023), estimated based an external instruments VAR for the European economy using the carbon policy surprise series as an instrument.

However, climate policy changes are not determined in a vacuum. Policymakers take economic conditions into account when formulating reforms of the carbon market. Thus, to adequately identify the exogenous variation in policies, we employ the approach by Känzig (2023) exploiting institutional features of the carbon market combined with high-frequency information on carbon futures prices. The idea is to construct a series of carbon policy surprises by measuring how carbon futures prices change in a tight window around regulatory policy events affecting the supply or allocation of emission allowances. These surprises are then used as an instrument in a structural vector autoregression (VAR) model of the European economy to identify a monthly series of structural carbon policy shocks, $CPShock_t$. Figure 1 shows the identified shocks over our sample from 1999 to 2019.

Local projections. To analyze the effects of a carbon policy shock on prices and household incomes, we use local projections à la Jordà (2005). For each outcome variable of interest, y_i, we run the following set of regressions for $h = 0, \ldots, H$:

$$y_{i,t+h} = \beta^i_{h,0} + \psi^i_h CPShock_t + \beta^i_{h,1} y_{i,t-1} + \ldots + \beta^i_{h,p} y_{i,t-p} + \xi_{i,t,h}, \tag{2}$$

where ψ^i_h is the effect on variable i at horizon h. This approach follows Cloyne et al. (2023) and Drechsel (2023), and correctly recovers the dynamic causal effects of interest, provided the carbon policy shock is properly identified. Some of our outcome variables are only available at quarterly frequency. However, we can still employ this approach: we aggregate the shock $CPShock_t$ to a quarterly series by summing over the respective months before running the local projections. We normalize the IRFs to correspond to a carbon-policy-induced increase in the HICP energy component by 1% on impact. Confidence bands are computed using the lag-augmentation approach (Montiel Olea and Plagborg-Møller, 2021).

Data. To implement the feasible set approach, we have the following data requirements. First, we require time-series data on consumer prices, and asset prices, all of which are readily available from national statistics. Second, we require life-cycle data on consumption expenditures by good, labor income profiles, asset income and accumulation profiles, and transfer profiles. These can be sourced from large-scale household surveys. Our formula for the money-metric welfare change (1) calls for the life-cycle variables to be evaluated at the deterministic steady state. We approximate this by assuming that 2015 represents a stochastic steady state (i.e., is absent aggregate shocks), and further assuming that the stochastic steady state is "close" in the sense of household choice variables to the deterministic steady state.

Our sample includes countries in the euro area, all of which are subject to the European carbon market, spanning 1999 to 2019.[6] We collect time-series data on consumer prices from Eurostat. Specifically, the Harmonised Index of Consumer Prices (HICP) for the 12 main consumption categories according to the Classification of Individual Consumption by Purpose (COICOP). For asset income and price changes, we collect stock price, dividend, bond price, and house price data from the ECB. We assume that stock and bond markets are integrated and we use euro area-wide data, while for house prices, we assume that markets are local and we collect country-specific data.

We complement this time-series data with household survey data covering the required parts of the household budget constraint. The challenge is that there is no single survey in Europe that spans all the information we need. Therefore, we combine granular information from a selection of European household surveys. Specifically, we use the Household Budget Survey (HBS) to obtain detailed information on household consumption baskets by demographic, the Statistics on Income and Living Conditions survey (SILC) for information on income, and the Household Finance and Consumption Survey (HFCS) for household portfolios. The former two surveys are available through

[6] Because we do not have all the required survey data for the Netherlands and Austria, these two countries are excluded from the analysis.

Eurostat, while the latter is administered by the ECB. The SILC survey has coverage back to 2004 and includes rich information on labor and transfer income. We construct time series of labor income by different household groups to estimate the impulse responses based on the approach outlined above. Details on the data sources and construction are in Appendix B.

4. The Dynamic Causal Effects of Carbon Policy Shocks

This section presents the estimated impulse responses to a carbon policy shock. Throughout, we report point estimates alongside 68% and 90% confidence bands.

Consumer prices. Figure 2 presents the responses of the energy component of the HICP, as well as HICP indices for housing and utilities, transport, and food. The identified shock leads to a significant increase in energy prices, which rise by 1% on impact and remain elevated for approximately a year before gradually returning to baseline. The response is consistent with the effect estimated in Känzig (2023), reflecting a strong pass-through from carbon prices to energy prices.

Higher energy prices also spillover to an increase in other consumer prices, particularly in categories that directly include energy costs or are heavily dependent on energy. For instance, housing and utilities experience a significant and persistent price increase, with the effect still at approximately 0.5% three years after the shock. A similar pattern emerges for transport prices, which increase by 0.5% on impact but revert to baseline within a year. Although food and beverages do not directly include energy-related goods, they are heavily influenced by energy costs. Accordingly, food prices also exhibit a significant increase, although to a lesser extent.

Figure 3 displays the impact response as well as the response after one year for the consumer prices for all 12 COICOP categories that we consider. In addition to the energy-related categories discussed earlier, we also observe significant price spillovers in restaurants and hotels, recreation and culture, and clothing. However, one year after the shock, most effects become insignificant. The main exceptions are housing and utilities, where prices remain significantly higher, and education and communication, where prices actually fall. This decline can be attributed to adverse general equilibrium effects via wages and employment, which put downward pressure on these prices.

Figure 2: Responses of Selected Consumer Prices

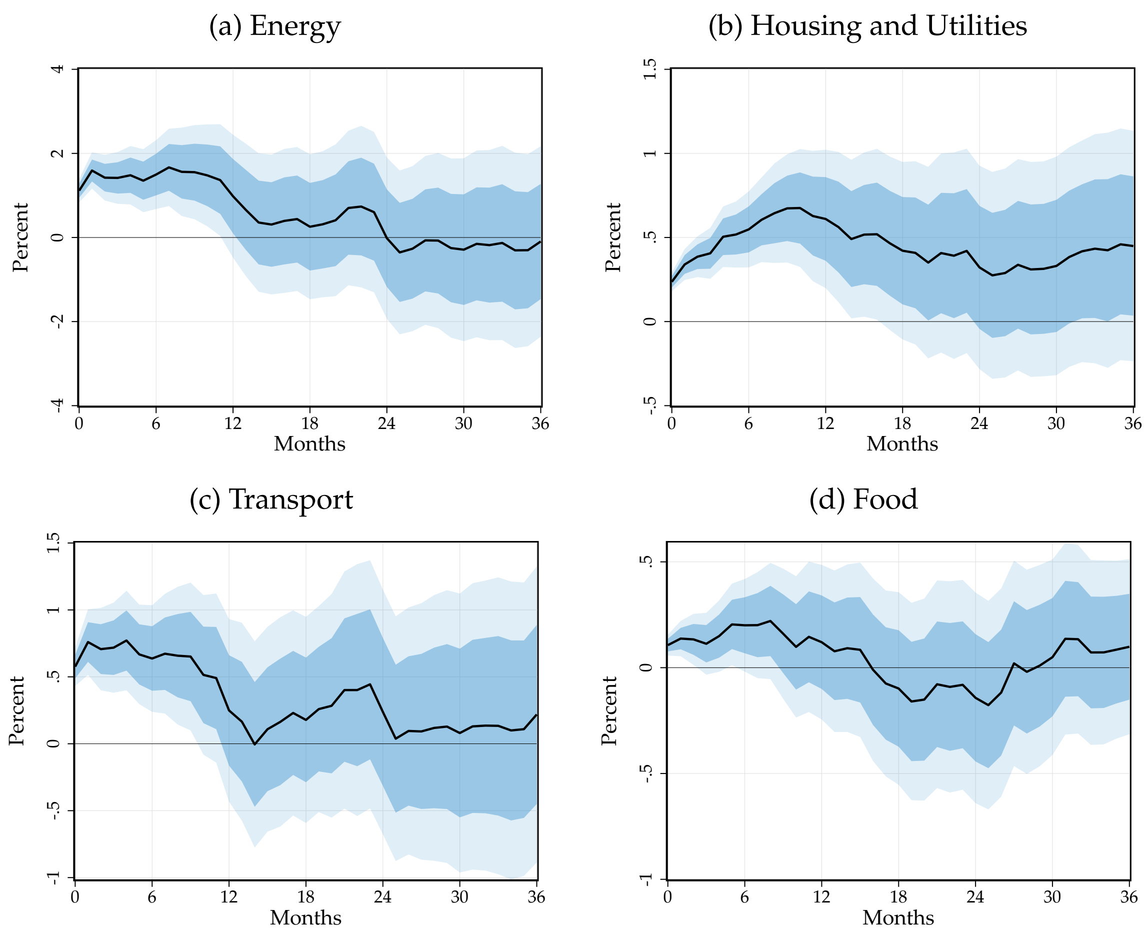

Notes: Impulse responses of selected HICP components, estimated using local projections (2) on the carbon policy shock from Känzig (2023). The responses are normalized to HICP energy by 1% on impact. Controls: 6 lags of outcome variable. The solid line is the point estimate and the dark and light shaded areas are 68% and 90% confidence bands. Standard errors computed based on lag-augmentation approach.

Figure 3: Consumer Price Responses of all COICOP Categories

(a) Impact Response

(b) Response After 12 Months

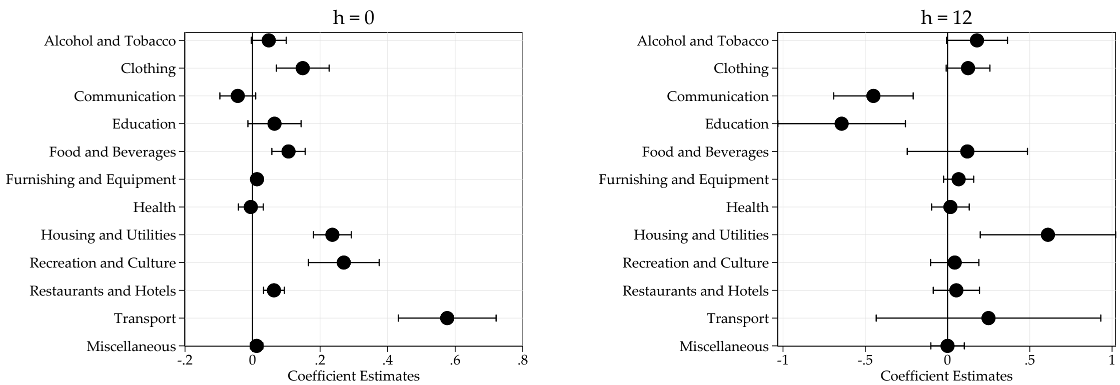

Notes: Impulse responses of all 12 HICP components per COICOP classification at $h = 0$ and $h = 12$, estimated using local projections (2) on the carbon policy shock from Känzig (2023). Controls: 6 lags of outcome variable. The responses are normalized to HICP energy by 1% on impact. Shown is the point estimate with 90% confidence bands.

Labor income. Figure 4 shows the average response of labor income. Labor income falls substantially after the shock, with a peak effect of about 1%. The response also features a considerable degree of persistence.

However, this average response masks substantial heterogeneity across households. To explore potential heterogeneity, we group households by educational attainment (college vs. non-college educated) and income level (low income, defined as the lower quar-

Figure 4: Average Labor Income Response

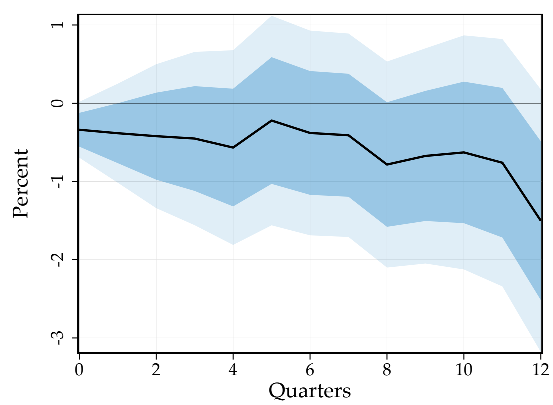

Notes: Impulse responses of labor income, estimated using local projections (2) on the carbon policy shock from Känzig (2023) aggregated to quarterly frequency. The responses are normalized to HICP energy by 1% on impact. Controls: 2 lags of outcome variable. The solid line is the point estimate and the dark and light shaded areas are 68% and 90% confidence bands. Standard errors computed based on lag-augmentation approach.

Figure 5: Labor Income Responses by Education and Income Groups

Education:

(a) Non-College Educated

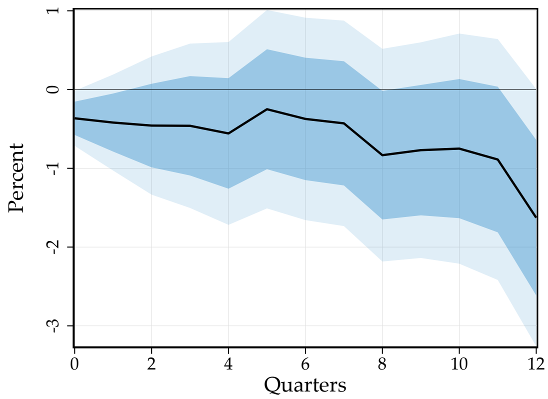

(b) College Educated

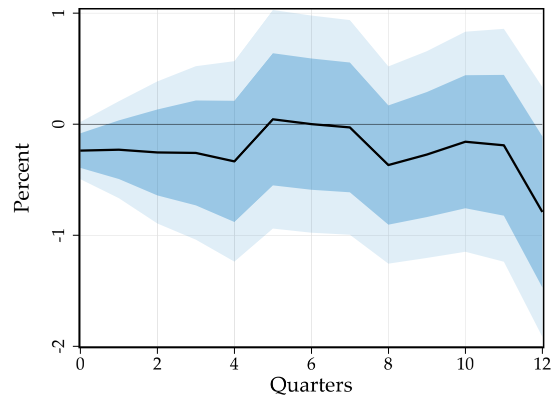

Income:

(c) Low Income

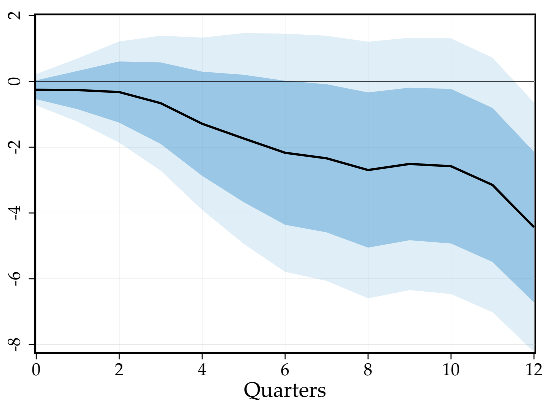

(d) High Income

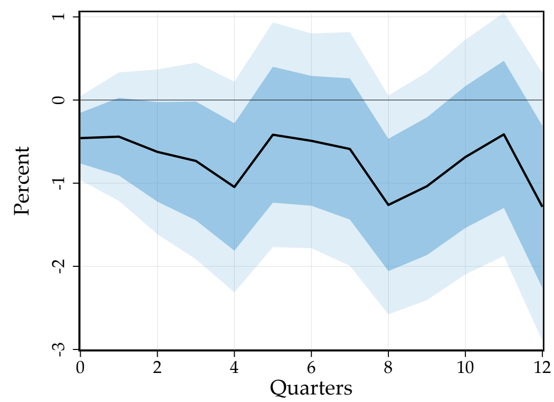

(e) Mid Income

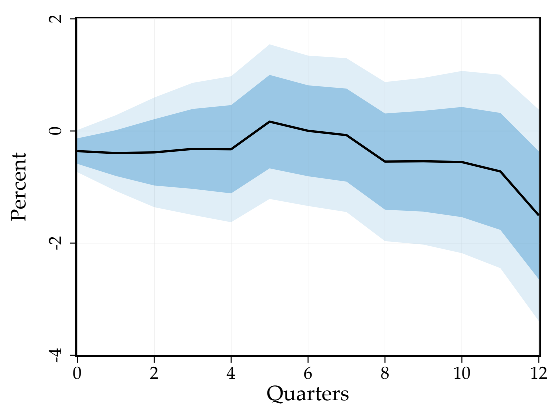

Notes: Impulse responses of labor income by education and income level, estimated using local projections (2) on the carbon policy shock from Känzig (2023) aggregated to quarterly frequency. The responses are normalized to HICP energy by 1% on impact. Controls: 2 lags of outcome variable. The solid line is the point estimate and the dark and light shaded areas are 68% and 90% confidence bands. Standard errors computed based on lag-augmentation approach.

tile of the income distribution; middle income, defined as the middle 50%; and high income, defined as the upper quartile).

Figure 5 shows the income responses by education level. We can detect meaningful differences: The response of non-college educated displays a larger fall, with a peak effect in excess of 1%. The response looks fairly similar to the average labor income response, consistent with the fact that the majority of households do not have a college degree. By contrast, the response of college educated is more muted and less statistically significant.

The heterogeneity is even starker by income levels. Figure 5 displays the IRFs for the three income groups we consider. We find the strongest effect for low-income households, where incomes fall by around 2% to 4% over the three years of the estimated response. Interestingly, we also document a non-negligible fall for high-income households. Their labor income falls by about 1%. Middle-income households, on the other hand, appear to be more insulated. Their labor income falls by around 0.5-1%. These findings are consistent with results in Känzig (2023) for a sample of British households.

Compared to the results in Del Canto et al. (2023), which studies the impacts of an oil price shock on US households, we find more pronounced labor income responses and a more substantial degree of heterogeneity. This may be attributed to the relatively more rigid labor markets in Europe, which could amplify the effects of such shocks on labor income.

Financial income. Figure 6 reports the responses of asset prices to a carbon policy shock. Stock prices fall significantly, reaching a peak effect of about 5%. In turn, dividends respond more gradually, starting to fall significantly after about a year, with a decline of around 3%. The effect on bond prices is weaker, with bond prices declining in the first year, followed by a reversal. Finally, house prices decline substantially, but with a considerable lag.

From an asset price perspective, equity positions lose most from the carbon policy shock because they receive lower dividend income. Importantly, however, the carbon policy shock is beneficial to those who were planning to accumulate equity, because doing so is cheaper after a shock. Thus, the strength of the portfolio channel for carbon policy shocks depends critically on who currently holds and who is planning to hold equities. Similarly, for house prices, those who were planning to buy homes benefit, as they can acquire these at a lower prices after a shock.

Figure 6: Asset Price Responses

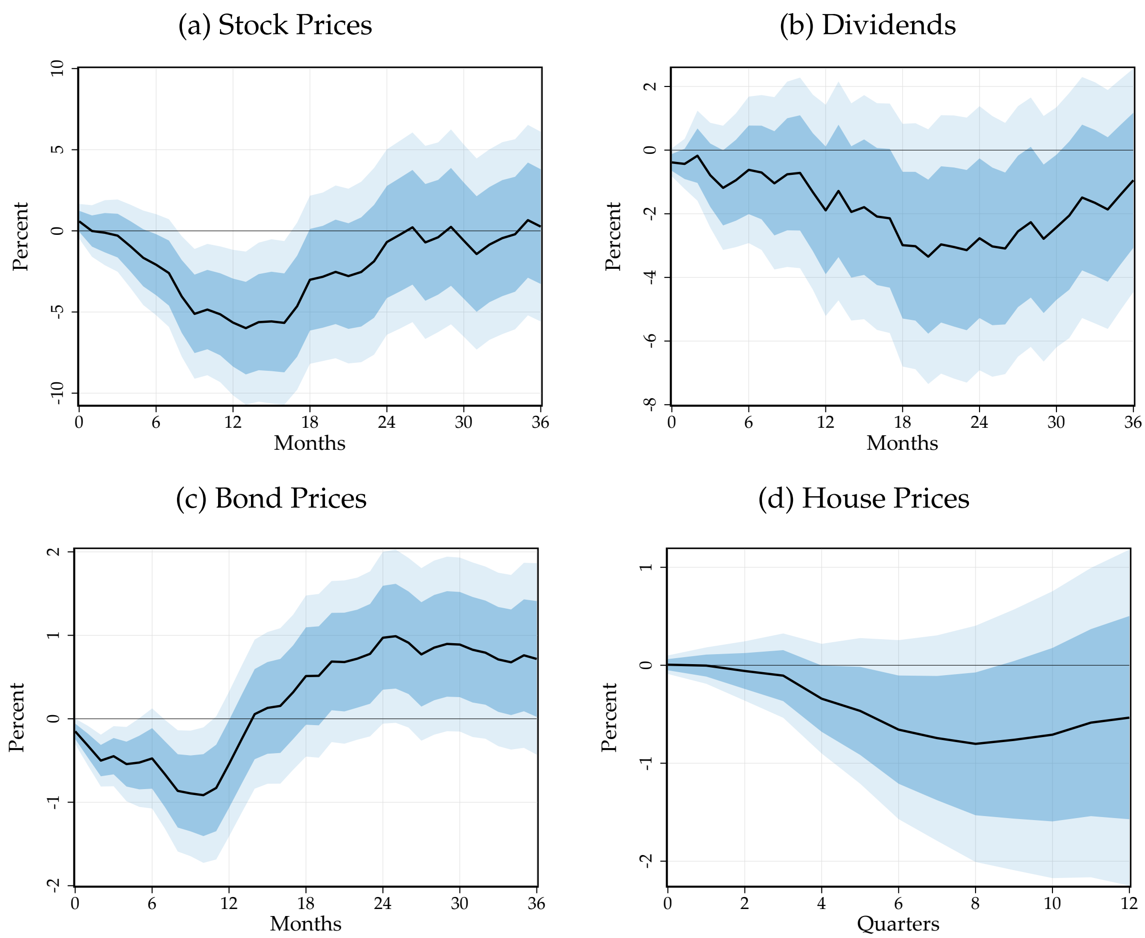

Notes: Impulse responses of stock prices, dividends, bond prices, and house prices, estimated using local projections (2) on the carbon policy shock from Känzig (2023). The responses are normalized to HICP energy by 1% on impact. Controls: 6 lags of outcome variable for monthly variables, 2 lags for quarterly variables. The solid line is the point estimate, and the dark and light shaded areas are 68% and 90% confidence bands. Standard errors are computed based on the lag-augmentation approach.

5. The Cross-Section of EU Households

We now turn to the second part of our money-metric welfare calculation: the cross-sectional moments of households income, expenditure, and portfolio positions. We are particularly interested in differences along the income distribution.

Table 1: Descriptive Statistics on Households by Income Group

	Overall	By income group		
		Low-income	Mid-income	High-income
Income				
Labor income	15,823.00	5,350.63	15,198.66	25,635.54
Transfer Income	875.50	1,335.35	742.97	679.24
Pension Income	5,382.96	3,964.65	6,024.67	5,522.47
Expenditure				
Housing and utilities	5,595.43	5,116.29	5,184.38	7,482.59
Transport	1,979.21	1,393.55	1,906.48	3,367.75
Goods	4,201.84	3,813.62	3,935.82	5,587.68
Other services	3,238.23	2,428.64	3,072.06	5,297.91
Wealth				
Housing wealth	45,830.20	24,904.92	44,652.19	84,069.98
Other assets	32,121.66	11,646.56	29,503.44	80,730.31
Household characteristics				
Age	52	54	53	50
Education (% with college degree)	27.9%	14.8%	24.1%	48.6%

Notes: Descriptive statistics on household characteristics: Average household income and expenditure, average wealth, average age as well as the share with college degree. All monetary terms are expressed in 2015 Euros. Income and expenditure figures are annual estimates. Goods expenditures includes food and beverages, alcohol and tobacco, clothing and food and equipment. Other services include communication, education, health, recreation and culture, and restaurants and hotels. Age and education characteristics were taken from the EU-SILC dataset separately.

Table 1 presents summary statistics on (after-tax) labor and transfer income, major expenditure categories, housing wealth, and assets, as well as their variation across the three income groups we consider. As expected, low-income households, on average, have much lower incomes but receive relatively more in transfer income. They also spend more than they earn, which could reflect borrowing, dissaving, or unreported income. Relative to their total expenditures, they allocate a larger share to housing and utilities. Additionally, they own far less wealth than richer households.

In terms of demographics, age distributions across income groups are relatively similar. However, as expected, the share of college-educated individuals is much lower among low-income households—only about 15%, compared to nearly 50% in the high-income group.

In Appendix Tables B.1-B.3, we provide further descriptive statistics by education, region, and age. The heterogeneity by education is qualitatively similar as for income. The wealthiest region is the West, followed by Northern, Southern and then Eastern European countries. Interestingly, Eastern European countries spend a smaller share of their income on energy-intensive goods such as housing and utilities, and transport.

6. Money-metric Welfare Calculations

We now report the welfare cost of carbon-policy shock calculations. We construct the money-metric utility losses associated with the carbon policy shocks, the data counterpart of (1). We group households along three dimensions: age, education, and income. For the age groups, we consider 25-34, 35-49, 50-64, and over 65 old. Throughout, we evaluate the welfare consequences of a carbon policy shock that increases HICP energy by 1% on impact. We limit our focus to short- to medium-run impacts over a three-year horizon, as we lose statistical power if we estimate the effects further out. The welfare changes are expressed as a share of total three-year consumption.

The average household experiences a substantial welfare impact: a carbon policy shock increasing energy prices by 1% leads to a welfare loss of about 0.5%. This translates into an annual loss of about 80 euros. This, however, again masks substantial heterogeneity in these impacts.

Heterogeneity by education. We start by studying how the welfare costs vary by educational attainment. The top panel in Figure 7 shows the total welfare loss for different education groups by age. Two main findings stand out: First, households with a non-college degree display substantially higher losses, between 0.6-0.7%, compared to 0.3-0.4% for the college-educated. Second, this pattern holds across all ages, except among the retired. Indeed, for households over 65 years old, we find much smaller welfare losses of around 0.2%. Thus, while older households have been largely responsible for climate change, they are least affected by climate change mitigation policies. On the other hand, they also benefit the least from such policies, as the benefits will only materialize relatively far in the future.

Next, we decompose these losses into the channels operating through different parts of the budget constraint. The lower panels in Figure 7 show how the consumption, labor income, portfolio, and transfer income channel contribute to these welfare losses. The

Figure 7: Welfare Change by Education Level

(a) Total Welfare

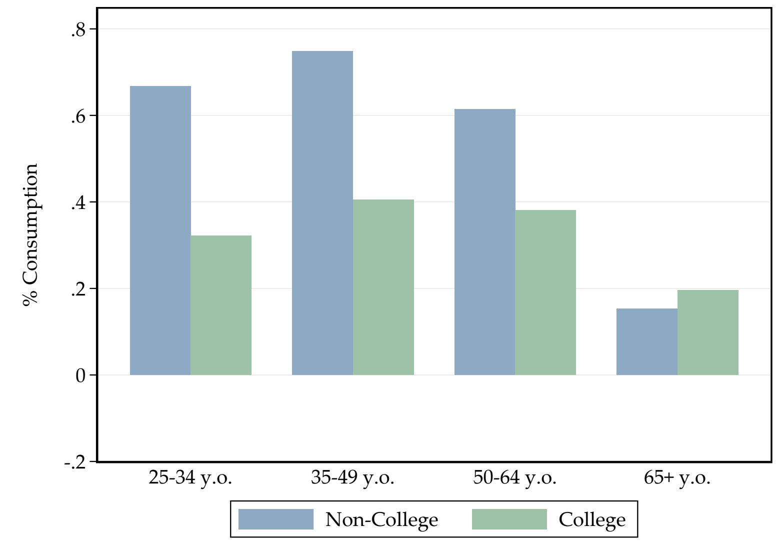

(b) Consumption Channel

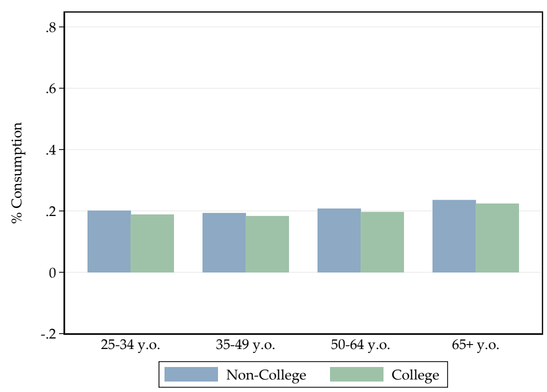

(c) Labor Income Channel

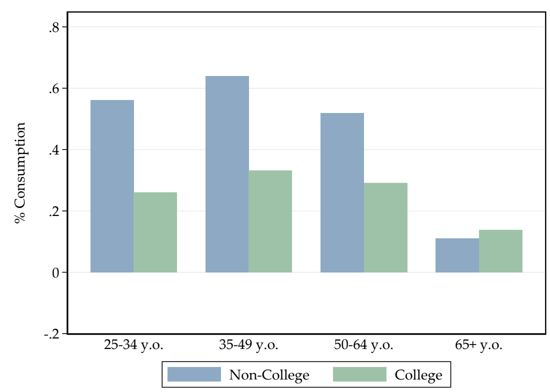

(d) Portfolio Channel

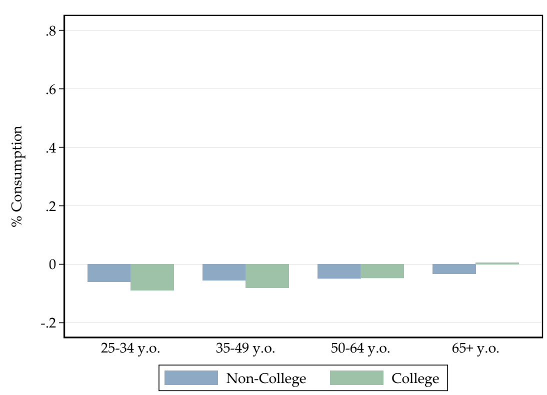

(e) Transfer Channel

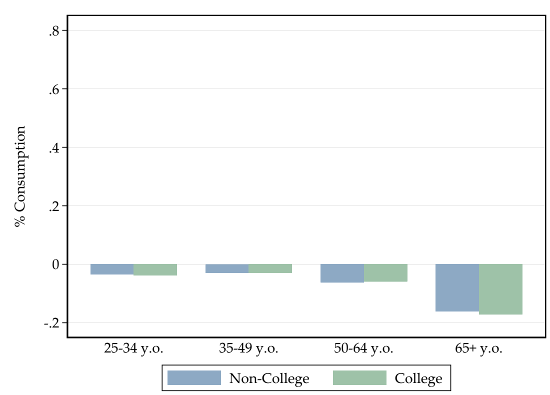

Notes: Money-metric welfare impacts, estimated based on (1), by age and education groups. The welfare change is expressed as a share of total three-year consumption by age and group based on 2015 data. A positive number represents a welfare loss, a negative number a welfare gain.

main insights are the following: First, there is not much heterogeneity in the consumption channel. Non-college-educated and older households lose the most as a share of consumption. However, the differences are minimal. The effects also turn out to be rather modest, varying from slightly below 0.2% to around 0.25%. This is consistent with the fact that the energy-intensive portion of households' consumption basket is not high and does not vary greatly across households.

The results for the labor income channel are vastly different. We observe stark heterogeneity in the welfare losses through labor income, with the non-college-educated households displaying substantially larger losses (around 0.6%) compared to college-educated (about 0.3%). As expected, we only observe modest impacts for households aged 65 years or older, as most of this group is retired from the labor market.

The portfolio channel, on the other hand, leads to broad-based welfare gains. These are most pronounced among younger college-educated households. The magnitude of this channel is much smaller, though, with average welfare gains of around 0.04%. Older households gain less or may even lose through the portfolio channel, as they are more likely to be net sellers of assets rather than buyers.

The transfer income channel also leads to welfare gains. Interestingly, we do not find big differences across different levels of educational attainment. The transfer channel is most pronounced for older households, mainly because rents are indexed to inflation in many European countries and thus increase with the inflationary shock.

Heterogeneity by income. While education is a common proxy for lifetime income, it does not correlate perfectly. Hence, we also study how the effects on welfare vary by income.

The top panel in Figure 8 shows the total welfare loss for our three income groups, low-income, middle-income, and high-income households, by age. Carbon taxes lead to an increase in inequality. The most affected households are poor and young households that display significant welfare losses exceeding 1%. Interestingly, high-income households also experience significant losses. The middle-income and older households are more shielded from carbon policy shocks. Thus, unexpectedly, welfare losses by income are u-shaped.

By dividing the total effect into different welfare channels, we can corroborate that the heterogeneity in effects is due mainly to the labor income channel. Importantly, these effects are substantially larger than what Del Canto et al. (2023) estimate for oil price shocks in the United States. One explanation is that labor markets are more rigid in Europe (Blanchard and Wolfers, 2000), leading to worse labor market outcomes in response to shocks.

Figure 8: Total Welfare Change by Income Groups

(a) Total Welfare

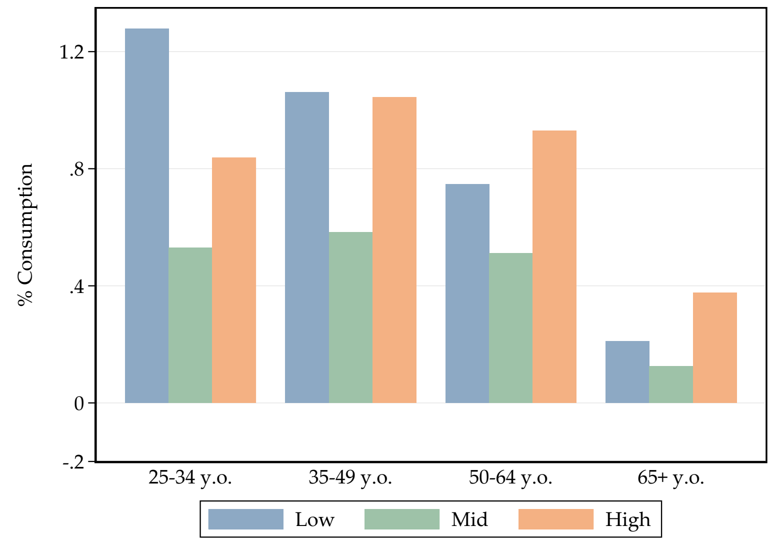

(b) Consumption Channel

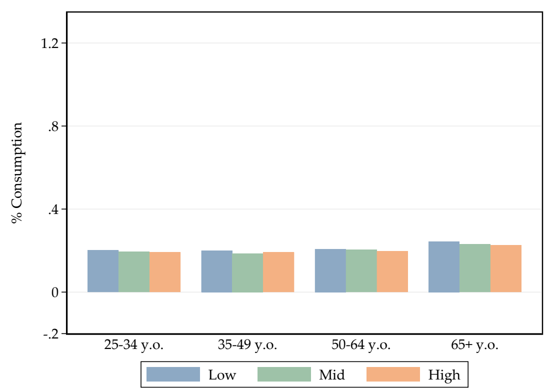

(c) Labor Income Channel

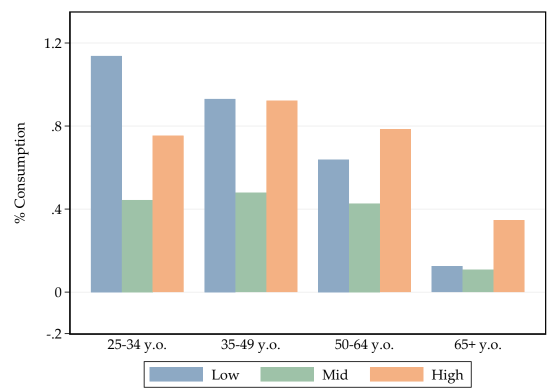

(d) Portfolio Channel

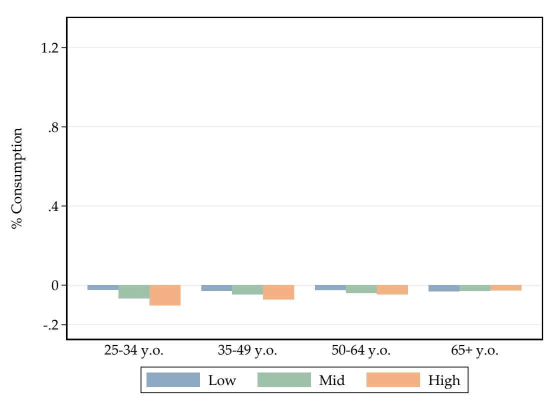

(e) Transfer Channel

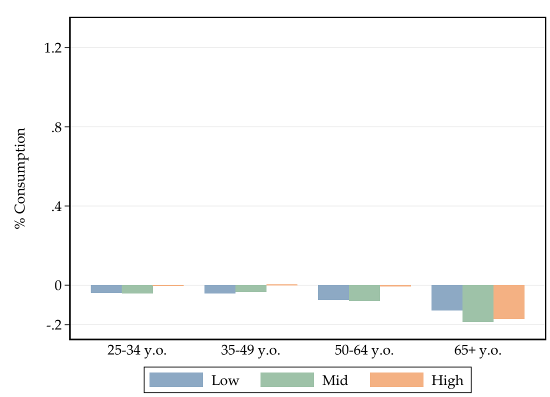

Notes: Money-metric welfare impacts, estimated based on (1), by age and income groups. The welfare change is expressed as a share of total three-year consumption by age and group based on 2015 data. A positive number represents a welfare loss, a negative number a welfare gain.

The portfolio channel cushions some of these losses, particularly for high-income and younger households, but is quantitatively much less important than the labor income or the consumption channel.

Regional heterogeneity. We have documented substantial heterogeneity by education and income. Another important dimension is regional heterogeneity. Next, we examine how the welfare impacts of carbon policy shocks differ across four European regions: Southern, Western, Northern, and Eastern Europe.[7] Because our main sample of euro area countries includes only a few Eastern European countries, we extend the analysis where possible to include all EU-27 countries. This allows for a more comprehensive assessment of regional differences. However, the balance sheet data is only available for countries in the euro area. Thus, we assume that the balance sheet positions of euro-area countries within each region broadly represent all countries in that region.

The top panel in Figure 9 shows the total welfare impacts by age group across different regions. We observe a very stark regional heterogeneity. The most affected region is Southern Europe, where the welfare impact ranges between 0.5% to 0.8% for the working-age population. The welfare losses are also substantial in Eastern Europe and to a somewhat lesser extent in Western Europe, standing at around 0.2% to 0.3%. The most insulated region appears to be Northern Europe, which displays welfare losses below 0.2%.

Unpacking the sources of these welfare impacts, we again find a stark heterogeneity in all the channels: Southern and Western European countries display a high welfare cost through the consumption channel, standing at around 0.2%. The effect is much weaker for Northern and Eastern European countries. This is mainly driven by a lower pass-through of carbon to energy prices in these countries. Northern European countries are more insulated because they source a large share of their energy through renewables. Eastern Europe, on the other hand, was allocated a disproportionate share of free ETS allowances over our sample, which could help explain the limited price impact in that region. This is consistent with the results in Känzig and Konradt (2023).

The labor-income channel represents the largest share of losses, particularly in Southern and Eastern European countries. The effect is most pronounced in Southern Europe, consistent with the notion that labor markets in this region is more rigid due to stronger employment protection laws and less flexible wage-setting institutions. Both of these regions benefit disproportionately from the portfolio channel but not by enough to offset

[7]We group European countries into these four regions according to the UN geoscheme.

Figure 9: Welfare Change by Regions

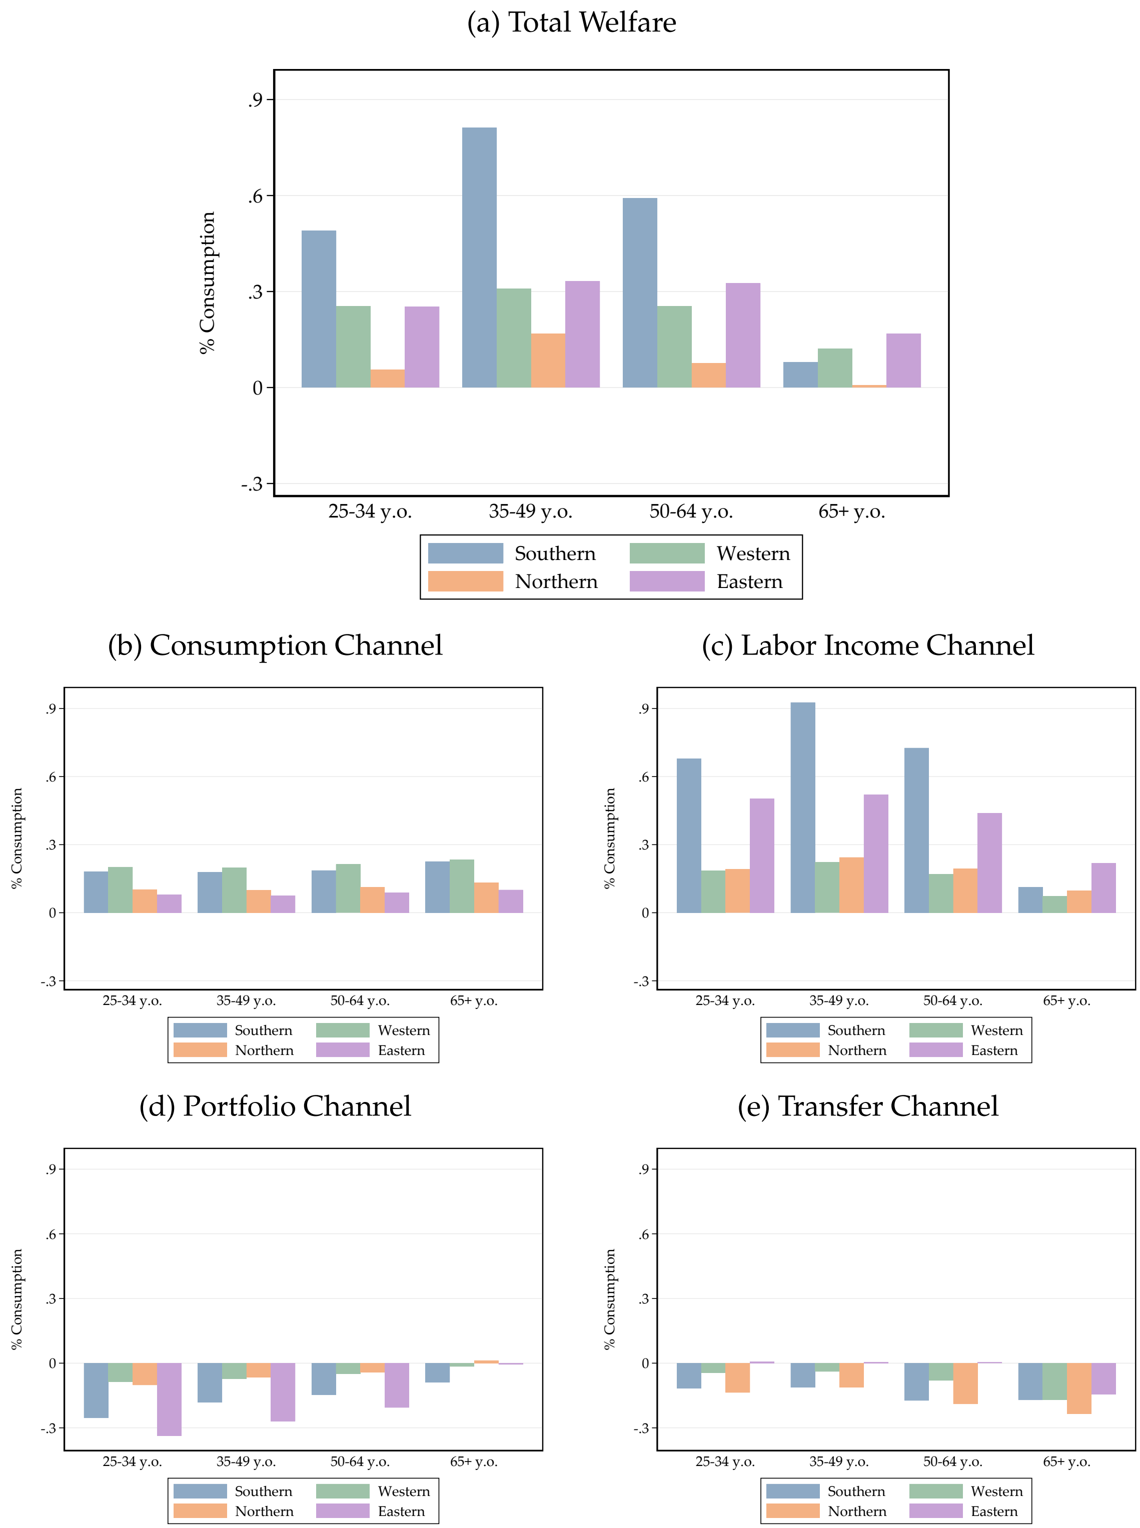

Notes: Money-metric welfare impacts, estimated based on (1), by age and region. The welfare change is expressed as a share of total three-year consumption by age and group based on 2015 data. A positive number represents a welfare loss, a negative number a welfare gain.

the severe consumption and labor income impacts.

Another interesting finding concerns the impact of the transfer channel: We observe welfare gains through the transfer channel in most regions, particularly in Northern and Southern Europe. This pattern is consistent with the fact that these regions have relatively more generous welfare programs. In contrast, the transfer channel is absent in Eastern Europe, at least among the working-age population. This can be attributed to the generally lower level of social protection and weaker automatic stabilizers in Eastern European countries, where transfer systems are less extensive and less effective in countering economic shocks.

7. Conclusion

The curse of economics is that the importance of a question is inversely related to the precision with which it can be answered. Here, we attempt to answer an important question: the distributional impact of carbon taxes. Given its importance, the answers need further refinement. Improving data collection and understanding the budget constraints of different consumers better is one possible way forward. Another promising route is to consider higher-order effects on welfare beyond those studied here: effects on unemployment hazards or borrowing costs induced by carbon taxes.

Despite the many limitations, the estimates here, especially in light of corroborating survey evidence, offer insights into reducing opposition against carbon taxation. We highlight three takeaways:

1. The sizable welfare losses suggest that public resistance may stem from genuine economic concerns, not simply distrust of taxes or an ideological divide. On average, an increase in carbon taxes that raises energy prices 1% causes a $\approx 0.5\%$ decrease in money-metric welfare.

2. Because the main driver is indirect macroeconomic feedback through the income channel, measures beyond conventional public finance approaches—such as compensation schemes that focus solely on prices—may be beneficial. For example, the results suggest that expansionary (green) monetary policy could help ease the income burden.

3. The disproportionate burden on working-age households relative to retirees highlights the right direction of targeted redistribution: from the old to the young.

We hope this exercise inspires complementary research, leading to a shared consensus. Such consensus is essential for formulating a more comprehensive climate-change policy, one that also accounts for inequality.

References

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky (2019). "Regressive sin taxes, with an application to the optimal soda tax". *The Quarterly Journal of Economics* 134.3, pp. 1557–1626.

Andersson, Julius and Giles Atkinson (2020). "The distributional effects of a carbon tax: The role of income inequality".

Andersson, Julius J (2019). "Carbon taxes and CO2 emissions: Sweden as a case study". *American Economic Journal: Economic Policy* 11.4, pp. 1–30.

Baqaee, David, Ariel T Burstein, and Yasutaka Mori (2022). "A new method for measuring welfare with income effects using cross-sectional data". *NBER Working Paper* w30549.

Baqaee, David R and Ariel Burstein (2023). "Welfare and output with income effects and taste shocks". *The Quarterly Journal of Economics* 138.2, pp. 769–834.

Belfiori, Elisa and Manuel Macera (2025). "Optimal Climate Policy with Demographic Transitions". UTDT Working Paper.

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.

Bettarelli, Luca, Davide Furceri, Loredana Pisano, and Pietro Pizzuto (2025). "Greenflation: Empirical Evidence using Macro, Regional and Sectoral Data". *European Economic Review*, p. 104983.

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".

Blanchard, Olivier and Justin Wolfers (2000). "The role of shocks and institutions in the rise of European unemployment: the aggregate evidence". *The economic journal* 110.462, pp. 1–33.

Bovenberg, A Lans and Ruud A De Mooij (1994). "Environmental levies and distortionary taxation". *The American Economic Review* 84.4, pp. 1085–1089.

Bovenberg, A Lans and Lawrence H Goulder (1996). "Optimal environmental taxation in the presence of other taxes: general-equilibrium analyses". *The American Economic Review* 86.4, pp. 985–1000.

Cloyne, James, Clodomiro Ferreira, Maren Froemel, and Paolo Surico (2023). "Monetary policy, corporate finance, and investment". *Journal of the European Economic Association* 21.6, pp. 2586–2634.

Dávila, Eduardo and Andreas Schaab (2022). *Welfare assessments with heterogeneous individuals*. Tech. rep. National Bureau of Economic Research.

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

Del Canto, Felipe N, John R Grigsby, Eric Qian, and Conor Walsh (2023). *Are Inflationary Shocks Regressive? A Feasible Set Approach*. Tech. rep. National Bureau of Economic Research.

Dixit, Avinash (1985). "Tax policy in open economies". *Handbook of public economics*. Vol. 1. Elsevier, pp. 313–374.

Drechsel, Thomas (2023). "Earnings-based borrowing constraints and macroeconomic fluctuations". *American Economic Journal: Macroeconomics* 15.2, pp. 1–34.

Drews, Stefan and Jeroen CJM Van den Bergh (2016). "What explains public support for climate policies? A review of empirical and experimental studies". *Climate policy* 16.7, pp. 855–876.

Ewald, Jens, Thomas Sterner, and Erik Sterner (2022). "Understanding the resistance to carbon taxes: Drivers and barriers among the general public and fuel-tax protesters". *Resource and Energy Economics* 70, p. 101331. ISSN: 0928-7655.

Fagereng, Andreas, Matthieu Gomez, Emilien Gouin-Bonenfant, Martin Holm, Benjamin Moll, and Gisle Natvik (2024). "Asset-price redistribution".

Fullerton, Don (1997). "Environmental levies and distortionary taxation: comment". *The American Economic Review* 87.1, pp. 245–251.

Jaravel, Xavier and Danial Lashkari (2022). "Nonparametric measurement of long-run growth in consumer welfare".

Jordà, Òscar (2005). "Estimation and inference of impulse responses by local projections". *American economic review* 95.1, pp. 161–182.

Känzig, Diego R (2023). *The unequal economic consequences of carbon pricing*. Tech. rep. National Bureau of Economic Research.

Känzig, Diego R and Maximilian Konradt (2023). *Climate policy and the economy: Evidence from Europe's carbón pricing initiatives*. Tech. rep. National Bureau of Economic Research.

Konradt, Maximilian and Beatrice Weder di Mauro (2021). *Carbon taxation and inflation: evidence from the European and Canadian experience*. Tech. rep. Graduate Institute of International and Development Studies Working Paper.

Kopczuk, Wojciech (2003). "A note on optimal taxation in the presence of externalities". *Economics Letters* 80.1, pp. 81–86.

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, pp. 1–14.

Metcalf, Gilbert E (2019). "On the economics of a carbon tax for the United States". *Brookings Papers on Economic Activity* 2019.1, pp. 405–484.

Metcalf, Gilbert E and James H Stock (2023). "The macroeconomic impact of Europe's carbon taxes". *American Economic Journal: Macroeconomics* 15.3, pp. 265–286.

Montiel Olea, José Luis and Mikkel Plagborg-Møller (2021). "Local projection inference is simpler and more robust than you think". *Econometrica* 89.4, pp. 1789–1823.

Poterba, James M. (May 1989). "Lifetime Incidence and the Distributional Burden of Excise Taxes". *American Economic Review* 79.2, pp. 325–330.

Sandmo, Agnar (1975). "Optimal taxation in the presence of externalities". *The Swedish Journal of Economics*, pp. 86–98.

Online Appendix

Carbon Pricing and Inequality: A Normative Perspective

Saki Bigio
Diego R. Känzig
Pablo Sanchez
Conor Walsh

Contents

A. Framework

Time is discrete and indexed by t. There is a continuum of households indexed by i. There is both aggregate and idiosyncratic uncertainty; let s_t denote a history of realizations of states of the world, including idiosyncratic states, up to period t. We will call this history a state for convenience, but it should be read to include a sequence for previous realizations of stochastic variables before t. There are J consumption goods, indexed by $j \in \{1, \ldots, J\}$, with good j having price $p_{jt}(s_t)$ in period t.

There are $K+1$ long-lived assets, indexed by $k \in \{0, 1, \ldots, K\}$, available for trading in each period. Asset k pays a nominal dividend $D_{kt}(s_t)$ and may be traded at a price $Q_{kt}(s_t)$ in state s_t. We assume that asset $k = 0$ is a one-period nominal bond which pays one unit in all states s_t. We define the cumulative return on buying a sequence of these bonds from period 0 to t as $R_{0\to t}(s_t) \equiv \prod_0^t Q_\tau(s_\tau)^{-1}$. Lastly, asset $k = 1$, which we term "money", serves as the numeraire in this economy, pays a zero dividend forever, and is completely durable.

The economy is populated by a finite set of G different household types with overlapping generations. Let a denote the age of a household at some reference time $t = 0$, which we call the household's "initial age". A household type is determined by a combination of their initial age a and their group g. They die at group- and age-dependent rates, and we denote the cumulative survival rate of a cohort of initial age a by time t as δ_t^{ag}. Note that this nests both the canonical infinitely lived household with $\delta_t^{ag} = 1$, constant death rates, and finitely-lived overlapping generations structures with realistic death probabilities.

Let $N_{kt}^{ag}(s_t)$ denote the amount of asset k held by group g of initial age a at time t given a realization of s_t, where a negative value for N_{kt}^{ag} represents borrowing. We let Δ represent the first-difference operator so that $\Delta X_t^{ag} \equiv X_t^{ag} - X_{t-1}^{ag}$. Assets are subject to convex adjustment costs $\chi_k^{ag}(\Delta N_{kt})$. The one-period bond is assumed to not be subject to adjustment costs.

Let $T_t^{ag}(s_t)$ denote government transfers (or taxes, if negative) to households of group g and initial age a in period t given a realization of s_t.

Households have time-separable preferences with subjective discount factor $\beta_t^{ag} \in (0, 1)$. This is read as a household of initial age a, in group g, discounting period t from period zero, where the rate at which they discount is potentially non-constant. The household has preferences over consumption, labor, and asset holdings. We assume that each household type derives utility from consumption via an aggregator of goods

$$C_t^{ag} = C_t^{ag}(\{c_{jt}^{ag}\}_{j=1}^J), \tag{1}$$

where c^{ag}_{jt} is the consumption of good j chosen in period t by household g that is of initial age a. We assume that $\mathcal{C}^{ag}_t(\cdot)$ is increasing and continuously differentiable in all its arguments.

Households of type ag's preferences may be summarized by the differentiable utility function
$U^{ag}(C^{ag}_t, \{N^{ag}_{kt}\}^K_{k=1}, L^{ag}_t)$, where L^{ag}_t is the labor supplied by households of initial age a at time t. We assume that $U^{ag}(\cdot)$ is weakly increasing and concave in its first two arguments, and weakly decreasing and convex in labor. Note we assume that bonds do not enter the utility function, but money or other assets might. Excluding the quantity of one-period bonds from being in the utility function directly allows us to conveniently characterize an Euler equation for the households in terms of expected marginal utilities of consumption and the return on the bond.

Labor income for individual i is given by the product of three terms: $W^{ag}_t(s_t)e^i_t(s_t)L^{ag}_t(s_t)$. W^{ag}_t is an aggregate component of wages that varies with the aggregate state of the economy, as well as the age and group of the individual. $e^i_t(s_t)$ is an idiosyncratic stochastic process for efficiency units of labor, which has support on $\mathbb{R}_+$ and satisfies $\mathbb{E}_0[e^i_t] = 1$. This captures both general transitory and permanent fluctuations to idiosyncratic labor income. The stochastic process for e^i_t may be different for different household groups g and initial ages a. The state s_t contains realizations of both aggregate and idiosyncratic processes. We assume aggregate and idiosyncratic risks are independent so that we can partition the state into an aggregate component, s^A_t, and an individual component s^i_t: $s_t = \{s^A_t, s^i_t\}$ and $\pi_t(s_t) = \pi_t(s^A_t)\pi_t(s^i_t)$.

A representative type g household of initial age a takes as given its initial stock of asset holdings $\{N_{k,-1}\}_k$ and is a price taker. It solves the following utility-maximization problem:

$$V^{ag}(\{N_{k,-1}\}_k) = \max_{\{\{c^{ag}_{jt}(s_t)\}_j, L^{ag}_t(s_t), \{N^{ag}_{kt}(s_t)\}_k\}^{\infty}_{t=0,s}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^{ag}_t \delta^{ag}_t U^{ag}(C^{ag}_t(s_t), \{N^{ag}_{kt}(s_t)\}^K_{k=1}, L^{ag}_t(s_t)), \tag{2}$$

subject to state-by-state budget constraints for all t,

$$\sum_j p_{jt}(s_t)c^{ag}_{jt}(s_t) = \sum_k \left[N^{ag}_{kt-1}D_{kt}(s_t) - Q_{kt}(s_t)(\Delta N^{ag}_{kt}(s_t)) - \chi^{ag}_k(\Delta N^{ag}_{k,t}(s_t))\right] + W^{ag}_t(s_t)e^i_t(s_t)L^{ag}_t(s_t) + T^{ag}_t(s_t),$$

the consumption aggregator (1), and a series of no-Ponzi conditions

$$\lim_{T\to\infty} \mathbb{E}_0[R_{0\to T}^{-1} N_{kT}^{ag} Q_{kT}] \geq 0, \qquad \forall k \in \{0, \ldots, K\}. \tag{3}$$

Stochastic Structure. We suppose that the (general equilibrium) law of motion for the prices of the economy admits a VAR representation. Concretely, we assume that dividends, asset prices, goods prices, wages, and transfers are stochastic, and take the form

$$D_{kt} = \bar{D}_{kt} \exp(v_{kt}^{D})^{\sigma}, \qquad Q_{kt} = \bar{Q}_{kt} \exp(v_{kt}^{Q})^{\sigma}, \qquad p_{jt} = \bar{p}_{jt} \exp(v_{jt}^{p})^{\sigma},$$
$$W_t^{ag} = \bar{W}_t^{ag} \exp(v_t^{W^{ag}})^{\sigma}, \qquad T_t^{ag} = \bar{T}_t^{ag} \exp(v_t^{T^{ag}})^{\sigma}, \tag{4}$$

where $\sigma > 0$ is a parameter that scales the variance of the aggregate stochastic processes. These variables depend on a deterministic time component, denoted with a bar (e.g. $\bar{D}_{kt}$), and a stationary shock process (e.g. v_{kt}^{D}). We assume that the shock processes are functions of current and lagged values of a structural shock vector $\boldsymbol{\epsilon}_t$, such that

$$v_{kt}^{D} = \theta_k^{D}(L)\boldsymbol{\epsilon}_t, \qquad v_{kt}^{Q} = \theta_k^{Q}(L)\boldsymbol{\epsilon}_t, \qquad v_{jt}^{p} = \theta_j^{p}(L)\boldsymbol{\epsilon}_t,$$
$$v_t^{W^{ag}} = \theta^{W^{ag}}(L)\boldsymbol{\epsilon}_t, \qquad v_t^{T^{ag}} = \theta^{T^{ag}}(L)\boldsymbol{\epsilon}_t,$$

where each $\theta(L)$ is a lag operator matrix of finite dimension, and the elements of $\boldsymbol{\epsilon}_t$ are mutually uncorrelated. We collect these v_t^x into a vector $\mathbf{v}_t$. We further assume that the structural shocks $\mathbf{v}_t$ have no direct effect on household utility functions; we therefore rule out preference shocks, such as discount rate shocks. Finally, we assume that these structural aggregate shocks are independent from the idiosyncratic income process for each individual.

We leave the production structure of the economy unspecified, as long as equilibrium can be written in this fashion. Importantly, the aggregate economy need not be efficient.

Following **Stock2018**, we define the vector of structural impulse responses of the collection of variables affecting households' budget constraint $\mathbf{v}_t \equiv (\{v_{jt}^{p}\}_j, \{v_{kt}^{Q}, v_{kt}^{D}\}_k, \{v_t^{W^{ag}}, v_t^{T^{ag}}\}_{ag})$ at time t to the n^{th} entry of the structural shock vector $\boldsymbol{\epsilon}$ at time $t = 0$ as

$$\blacksquare_{n,t} \equiv \mathbb{E}_0[\mathbf{v}_t | \epsilon_0^n = 1] - \mathbb{E}_0[\mathbf{v}_t | \epsilon_0^n = 0].$$

Elements of this vector are denoted with superscripts, such as $\Psi_{n,t}^{p,j}$ for the consumption price of good j, and $\Psi_{n,t}^{Q,k}$ for the asset price of asset k.

B. Data Appendix

This appendix describes the data used in our analysis. It describes the datasets used: the Household Budget Survey (HBS), the EU statistics on income and living conditions (EU-SILC), and the Household Finance and Consumption Survey (HFCS) as well as the methodology used for cleaning the data. We also provide a description on the additional macroeconomic price data used.

B.1. Household Budget Survey (HBS)

The Household Budget Survey is comprised of national surveys focusing mainly on household expenditure on goods and services in the EU. This contains administrative data and is provided by Eurostat. They were launched in most euro states in the 1960s. The surveys have been conducted by Eurostat in all EU countries every 5 years, starting in 1988. The most recent waves are 2015 and 2020. This dataset contains granular information on households' expenditure at the annual level. Its main aim is to measure the consumption weights for the HICP. We use data from the 2015 wave, which is our base year.

We use the HICP 12 main components: alcoholic beverages and tobacco, clothing and footwear, communication, education, food and non-alcoholic beverages, furnishings, household equipment, health, housing and utilities, recreation and culture, restaurants and hotels, transport, and miscellaneous goods and services.

We focus on households aged 25 or older and at most 75 years old. We calculate per-capita household consumption by dividing household total consumption by household equivalent size (using the OECD scale, already provided in the survey). We group households age groups and another grouping category: education level, income group or regions. Education level is defined by the household head (already defined in the survey) and income groups are calculated as the top 25% (high), bottom 25% (low) and mid 50% (mid) using household net income. We then calculate the average expenditure for each one of these consumption categories for each group and age group. This results in the life-cycle consumption component.

We also calculate the three-year consumption profile for each group and age group. To do this, we first minimize jumps in consumption patterns caused by measurement error by running a Locally Weighted Scatterplot Smoothing (LOWESS) for each of the categories and each of the groups by age (without grouping by age groups yet). We then calculate the three-year consumption profile for a given group by assuming a constant life-cycle profile of consumption. We interpolate consumption by age for a given group and then calculate the total three-year consumption profile. We then group by age-groups and get three-year consumption profiles by each demographic group.

B.2. EU Statistics on Income and Living Conditions (EU-SILC)

The EU-SILC collects cross-sectional and longitudinal data on income, poverty, social exclusion, and living conditions. This survey is conducted annually across the EU. This survey is used to construct both life-cycle and time-series income statistics (both labor and transfers) for specific household groups. It is available from 2004 for some countries and 2005 for most euro area countries.

We construct labor income at the household level by adding cash and non-cash employment income. We build transfers by adding regular transfers and pensions. Regular transfers contain unemployment, sickness, disability and education benefits. On the other hand, pensions contain old age and survivors benefits. We use the household head education level to group houses. The household head is not defined in the survey, so we define it as the member who earns the highest income. We group houses into income groups by using net income at the household level in a similar way as with the consumption data.

We then proceed in two steps. We first calculate the life-cycle component by taking the weighted average labor and transfer income for each demographic group in our base year—2015. As a second step, we derive annual income time series for each group—education level, income group, or region—without grouping by age to reduce noise in the impulse response functions. This allows us to study the heterogeneous effect of carbon pricing on different groups, as in Känzig (2023). These time series are built by taking the weighted averages for each group. We only calculate the time series for labor income and regular transfers. We exclude pensions, as most are indexed to inflation, and we can use the response of HICP to carbon prices to estimate the effect on pensions.

The annual income time series are then converted to a quarterly frequency using a Chow-lin methodology. We use euro area wage data from the ECB to interpolate labor income and the unemployment rate to interpolate transfer income. For regional labor income, we directly use regional wage data from the ECB. The regional wage data is estimated by taking a weighted average of country-level wages, weighted by 2015 GDP, our baseline year.

B.3. Household Finance and Consumption Survey (HFCS)

The HFCS collects household-level data on household finances and consumption up to 2021. This is conducted every 4 years throughout the euro area by the ECB. It includes granular information related to households' balance sheet, income and consumption. We use this dataset to calculate life-cycle household portfolio.

We group assets into bonds, equities, and real estate. Bond holdings contains both government and corporate bonds. We don't separate into each type of bond holding as we don't have enough data. On the other hand, equities include mutual funds, directly held equities, and managed accounts. Finally, real estate is composed of housing and any other type of real estate. We also include adjustable-rate mortgages to estimate the effect of a change in mortgage yields caused by a carbon pricing shock. We focus on these rather than all mortgages as the yield for fixed-rate mortgages will not be affected by definition.

Similar to consumption, we also restrict attention to households of age between 25 and 75. We remove the top and bottom 1% of households by net worth. Demographic groups are built in the same way as with consumption and labor. We take the level of household head education to group households by education. We also group households by income groups using the household income level.

We are interested in both asset accumulation and asset holdings for the previous quarters. To do so, we use a similar approach to Del Canto et al., 2023. We first construct asset accumulation by interpolating asset holdings by year of age for a given demographic group. We then define asset accumulation as the difference in asset holding between age a and age $a - 1/4$. Our approach assumes that the accumulation profile of assets is constant over time. Using this approach, we also get the asset holdings for the previous quarter.

B.4. Additional Data Sources

We complement the household survey data with some aggregate price data from the ECB. In particular, we use the HICP and HICP components at the Euro level and at the country level. We also use house price indices as well as mortgage rates. We seasonally adjust HICP and house price data using the Census X11 approach.

The portfolio data is mainly sourced from Bloomberg. We use data for Eurostoxx 50 index and dividend yield to estimate the effect of carbon prices on the price of equities and their dividends. We use the ICE BofA corporate bond index to estimate the effect on bond prices.

B.5. Summary Statistics by Demographic Groups

Tables B.1-B.3 show descriptive statistics by education, region and age, respectively. We can see that those with college education have a larger labor income than those with non-college. The transfer component is larger for non-college education. Both consumption and wealth is higher for college educated. In particular, they spend a larger share in goods and other services, whereas non-college educated spend a larger proportion in housing and utilities. Those with college education have a larger share of other assets in their portfolio, whereas most of non-college educated wealth comes from housing.

When looking at regions, we can see that western and northern regions are the ones that earn and spend the most, followed by the southern region and eastern. The western region is the wealthiest while the northern and southern regions have similar wealth. The eastern region has lower income, consumption and wealth.

The data by age groups is consistent with a traditional life-cycle model. Labor income peaks has an inverse u-shape, peaking at mid age, whereas pension income peaks at old age for retirees. Consumption on housing increases by age but transportation decreases by age. Goods and services remain somewhat constant throughout the life-cycle. Finally, wealth increases throughout the life-cycle, where the proportion of housing assets increases as people age.

Table B.1: Descriptive Statistics on Households by Education Level

	By Education Level	
	Non-college	College
Income		
Labor income	12,193.67	23,985.40
Transfer Income	936.34	718.24
Pension Income	5,559.40	4,926.84
Expenditure		
Housing and utilities	5,320.13	6,386.48
Transport	1,742.17	2,660.34
Goods	3,987.80	4,816.87
Other services	2,780.02	4,554.88
Wealth		
Housing wealth	43,394.82	68,035.34
Other assets	29,256.45	62,173.68

Notes: Descriptive statistics on household characteristics. Average household income and expenditure, average wealth. All monetary terms are expressed in 2015 Euros. Income and expenditure figures are annual estimates. Goods expenditures includes food and beverages, alcohol and tobacco, clothing and food and equipment. Other services include communication, education, health, recreation and culture, and restaurants and hotels.

Table B.2: Descriptive Statistics on Households by Region

	By Region			
	Southern	Western	Northern	Eastern
Income				
Labor income	9,551.62	21,070.77	15,032.54	5,351.38
Transfer Income	556.28	1,027.09	1,723.61	193.36
Pension Income	3,883.53	6,620.47	3,442.19	1,656.92
Expenditure				
Housing and utilities	5,310.01	5,997.99	5,239.32	2,007.38
Transport	1,605.46	2,334.97	2,130.39	588.55
Goods	4,087.63	4,433.64	3,677.91	1,863.24
Other services	2,940.76	3,572.73	3,183.68	1,103.32
Wealth				
Housing wealth	55,182.78	54,816.49	48,327.50	24,994.05
Other assets	32,976.06	49,240.89	34,990.50	11,279.20

Notes: Descriptive statistics on household characteristics. Average household income and expenditure, average wealth. All monetary terms are expressed in 2015 Euros. Income and expenditure figures are annual estimates. Goods expenditures includes food and beverages, alcohol and tobacco, clothing and food and equipment. Other services include communication, education, health, recreation and culture, and restaurants and hotels.

Table B.3: Descriptive Statistics on Households by Age Groups

	By Region			
	25-34 y.o.	35-49 y.o.	50-64 y.o.	65+ y.o.
Income				
Labor income	13,506.84	15,794.80	14,173.82	4,023.78
Transfer Income	952.99	802.44	1,346.68	244.66
Pension Income	788.88	595.69	2,825.53	14,202.86
Expenditure				
Housing and utilities	4,880.41	4,806.89	5,618.42	6,706.17
Transport	2,353.41	2,073.76	2,201.13	1,493.86
Goods	3,816.62	3,905.89	4,437.15	4,434.65
Other services	3,440.27	3,217.42	3,323.53	3,087.69
Wealth				
Housing wealth	18,370.69	34,933.33	55,307.09	68,982.82
Other assets	20,940.40	31,698.95	44,253.08	47,502.05

Notes: Descriptive statistics on household characteristics. Average household income and expenditure, average wealth. All monetary terms are expressed in 2015 Euros. Income and expenditure figures are annual estimates. Goods expenditures includes food and beverages, alcohol and tobacco, clothing and food and equipment. Other services include communication, education, health, recreation and culture, and restaurants and hotels.

References Appendix

Del Canto, Felipe N, John R Grigsby, Eric Qian, and Conor Walsh (2023). *Are Inflationary Shocks Regressive? A Feasible Set Approach*. Tech. rep. National Bureau of Economic Research.

Känzig, Diego R (2023). *The unequal economic consequences of carbon pricing*. Tech. rep. National Bureau of Economic Research.