

Carbon tax incidence over the life cycle

Stanley Barrell-Kane and Hakki Yazici*

February 20, 2026

This paper studies the distributional consequences of carbon taxation over the life cycle. Using an environmentally-extended input-output table linked with the US Consumer Expenditure Survey, we show that the carbon intensity of expenditure declines with income but rises with age, largely due to differences in energy expenditure shares. We develop a heterogeneous-agent overlapping-generations model with non-homothetic preferences and age-dependent energy needs that matches these patterns. We evaluate a carbon tax path consistent with the U.S. net-zero target under the Paris Agreement. We find that, although higher energy prices disproportionately affect poorer and older households in the short run, lifetime welfare costs are broadly similar across income groups and display a U-shape over age, with middle-aged cohorts most affected. Crucially, these distributional outcomes depend on the timing of carbon taxes: front-loaded paths are more regressive and hurt older households more, while back-loaded paths shift the burden toward middle-aged cohorts and richer households.

JEL classification: E25, J31.

Keywords: Climate change, life cycle, carbon taxation, redistribution.

*University of Bristol. 12 Priory Rd, Bristol BS8 1TU, UK. Emails: stanley.barrell-kane@bristol.ac.uk, hakki.yazici@bristol.ac.uk.

We thank seminar participants at Bilkent University and the University of Bristol for valuable comments.

I Introduction

Climate change poses an existential threat to humanity. There is a growing consensus about the urgency to act. In line with this, an overwhelming majority of nations have pledged to the United Nations Framework Convention on Climate Change (UNFCCC) to reduce their carbon footprints as part of the Paris Agreement, an international treaty on climate change with an overarching goal to limit the extent of global warming by the end of 21st century.¹ At the same time, there is growing concern that the economic burden of climate policies falls unequally across households. Understanding the distributional consequences of climate policies is of utmost importance not only for equality per se but also because the political feasibility of implementing climate policies at the national level may depend on their distributional impacts.

In this paper, we choose the US economy as a case study to investigate the uneven impacts of climate policies across the population of the world's largest developed economy. In addition to analyzing the differential impact of climate policies across poor vs. rich households, a novelty of our analysis is to investigate how such policies affect households depending on their stage in the life cycle. We begin our analysis by documenting a key empirical regularity: the expenditure of older households is associated with a higher level of carbon emissions per dollar spent. We also confirm a key finding of the prior literature that the carbon emissions associated with a dollar spent is higher for poorer households (Chancel et al., 2022). Furthermore, we find that these emission patterns are mainly driven by energy expenditure shares. That is, poorer and older households spend disproportionately more on energy (predominantly gasoline, electricity and natural gas), which is by far the most emission-intensive category. This empirical characterization of energy spending patterns is important because the direct impacts of climate policies, which generally work by increasing the price of energy, will be felt differently according to households' spending behavior. At the same time, climate policies

¹The Paris Agreement was adopted by 195 Parties at the UN Climate Change Conference (COP21) in Paris, France, on 12 December 2015. It entered into force on 4 November 2016. It establishes a goal of holding the increase in the global average temperature to well below 2°C above pre-industrial levels and pursuing efforts to limit the temperature to 1.5°C.

are large-scale interventions that affect households' welfare indirectly by changing their wages and investment returns. To gauge the cumulative impact of these policies, we construct an empirically-disciplined, quantitative macroeconomic model which provides a laboratory to evaluate the short- and long-run consequences of climate reforms taking into account behavioral responses and general equilibrium effects. Specifically, we build a heterogeneous-agent overlapping generations model with incomplete markets that is able to match the observed energy expenditure patterns across households as well as some other key aggregate and distributional moments of the US economy. Our model does not take into account the environmental benefits from carbon tax policies; rather, we focus on the non-environmental welfare consequences.

In our baseline exercise we consider the specific climate reform that attains the emissions pathway pledged by the US in their latest nationally determined contribution (NDC) to the UNFCCC.² We use the model to back out the carbon tax pathway that delivers this emissions target, taking existing estimates of the time-varying costs of abatement technology as given. We find that the unique carbon tax pathway that implements the US emissions pathway to net-zero by 2050, the target net-zero date according to their NDC, starts from \$1 per ton of carbon emissions in 2020, our initial reform year, and increases rapidly to about \$500 by 2050, after which it gradually falls to about \$400 by 2100 and eventually to \$0 by 2250, reflecting the fact that staying at zero emissions requires lower taxes over time as abatement technology improves. The induced change in the price of energy follows the carbon tax, reaching a peak of 50% (relative to the 2019 price) by 2050. Wage and interest rates fall until 2050 after which they recover gradually. In the baseline reform, we assume that the government uses carbon tax revenues to reduce labor income taxes which implies average labor taxes also fall until 2050 and recover afterwards. Aggregate output increases by 1% until 2050 after which it gradually falls but remains

²For each country, a pledge to UNFCCC, known as nationally determined contribution (NDC), takes the form of a commitment to an emissions pathway to net-zero by a certain date and a particular plan of actions that will be taken to achieve this pathway. The emissions pathway pledged by each country's NDC is part of the broader Paris Agreement between countries to limit warming. Latest NDC was in 2024 by the Biden administration. Although since the submission of this NDC to United Nations in 2024 Trump administration has vowed to break away from the Paris Agreement, we think that evaluating the consequences of US's pathway to net-zero is still valuable as the US case provides a good laboratory and US may go back to Paris Agreement in the near future.

above the pre-reform level. Average household welfare falls however, primarily because higher output is achieved in part by higher labor supply.

Regarding the incidence over the life cycle, the welfare consequences of the reform have a U-shape, with those around 50 years old at the time of the reform suffering the most and the retired, and especially those older than 70, being hurt the least. At first glance, this may seem at odds with the energy spending patterns: older households spend disproportionately more on energy, and hence should suffer more from rising energy prices. This is in fact true in the short run, say within a decade or two after the reform. A different picture emerges, however, when we take into account longer-term effects. Because carbon taxes start low, older people die without facing particularly high energy prices. Those who are currently middle-aged become old and need energy the most by around 2050, when carbon taxes are at their highest levels. The young, on the other hand, are hurt less because they face high but not peak energy prices in old age, since for most of them retirement begins after 2050, and they face these taxes in a more distant future, which is less costly due to discounting. In addition to these direct effects, the indirect effects of the reform, which imply welfare gains through lower labor taxes and losses via lower interest rates, also favor younger people as labor income constitutes a relatively larger share of their total income.

The burden of the baseline reform is shared quite evenly between poor vs. rich households. Again, this may look at odds with the energy spending patterns: the poor spend disproportionately more on energy, and hence should suffer more from rising energy prices, echoing the consensus in the literature (see, for example, Kanzig, 2023). This mechanism is present in our framework. However, this static, short-term thinking falls short of explaining the full welfare consequences: the young and the retired, who are overrepresented in the low income quintiles, suffer less when we account for the dynamic impacts of the reform. Furthermore, the indirect welfare effects of the reform disfavor higher income groups as they benefit less from the decline in labor income taxes and lose more from falling interest rates, since a larger share of their total income stems from

capital income. The combination of these effects implies the rather uniform impact of the baseline reform along the income distribution.

Two general insights follow from our analysis. First, focusing on the short-term effects is likely to be misleading in gauging the full distributional impact of reforms. Second, the full distributional impacts of a climate reform depend crucially on the pathway of carbon tax rates it imposes and how this pathway coincides with people’s energy needs, which in turn depends on where they are in their life cycle.

We also consider alternative climate reforms. For comparability, we analyze a set of reforms that limit global warming by 2100 to the same level as in the baseline reform but differ in terms of how front-loaded carbon taxes are. We find that front-loaded paths are more regressive and hurt older households more, while back-loaded paths shift the burden toward middle-aged cohorts and richer households.

Literature review Seminal contributions to the literature on the economics of climate change include Nordhaus (1994), Nordhaus (2007) and Nordhaus (2008). Golosov et al. (2014) and Barrage (2019) investigate optimal carbon taxes in calibrated, representative agent macroeconomic climate models while Leach (2009), Rausch (2013), Carbone et al. (2013) and Kotlikoff et al. (2021) use representative-agent overlapping generations models to explore the differential welfare impacts of climate policies across generations.

More recently, a new generation of papers that analyze carbon taxation in the presence of heterogeneity have emerged³. Among these papers perhaps most closely related to the current paper are Fried, Novan and Peterman (2018) and Fried, Novan and Peterman (2024) which also analyze the consequences of carbon taxation using rich, quantitative OLG models with idiosyncratic productivity shocks and incomplete markets. Importantly, we differ from these papers in that we model the heterogeneity in energy needs over the life cycle and focus on actual policy proposals which involve carbon policy pathways that achieve net-zero emissions.⁴

³See Belfiori, Carroll and Hur (2024), Douenne et al. (2024), Kuhn and Schlattmann (2024)

⁴Benmir and Roman (2022) also evaluate the net-zero reform according to Paris Agreement; however, unlike us, they do not model energy consumption and life cycle.

The rest of the paper is organized as follows. Section II describes the empirical findings while Section III lays out the model. Section IV explains how we parameterize the model and Section V discusses the main quantitative results. Section VI concludes. An online appendix contains, among other things, the details of our empirical calculations, a discussion of the welfare gains decomposition, and the details of alternative calibrations of the model.

II Empirical Findings

Data construction We follow Levinson and O’Brien (2019) and Sager (2019) among others in constructing consumption-based accounts of household CO₂ emissions. These accounts combine input-output tables with production-based emission estimates to assign to each product the emissions ‘embodied’ in a dollar’s worth of its consumption. This includes the emissions produced at every stage of the production process. Assigning emissions to households in this way is generally regarded to account for between 60%-75% of a country’s total emissions, with the rest being attributable to investment or government spending.

We use publicly available data from EXIOBASE 3 for the years 1997-2019. EXIOBASE provides a multi-region environmentally-extended input-output table, which means that it also accounts for imported inputs in production processes which may have different carbon intensities from domestic inputs. As in Sager (2019), we augment these estimates with data from the Energy Information Administration (EIA) on tail-pipe emissions for electricity, natural gas and gasoline, to also assign to households the emissions generated from the final use of energy.

EXIOBASE provides carbon intensity estimates for 200 product categories, which is relatively granular compared to many earlier studies. Nonetheless, like all previous work which employs these consumption-based estimates, we are unable to account for differences in emission intensity *within* products, for instance between more and less expensive versions of the same good, as all goods are assigned the average emission

intensity in their product category. The interpretation of carbon intensity per-dollar of expenditure as structural is also not without limitations. In particular, it assumes that doubling expenditure on a good will double the emissions generated, regardless of whether the price or the quantity have doubled. This is unlikely to be a problem for homogeneous goods like gasoline, but may not hold for all goods. A breakdown between price and quantity is also not available in the consumer surveys we use, and as such the per-dollar of expenditure formulation is a necessary approximation.

We combine these estimated carbon intensities with the Interview Files of the Consumer Expenditure Survey (CEX) over the same time period. Households are interviewed quarterly about their expenditures, with their spending categorised into one of around 850 UCC codes. We aggregate these files to an annual frequency, and drop all households that do not participate in all waves or who have incomplete income information. Following Levinson and O’Brien (2019), we drop households with top and bottom 1% after-tax income due to topcoding, and also drop households with nursing home expenditure. We further restrict the sample to households with heads aged 20 to 80, who report positive total expenditure and energy expenditure. Households with no information about their state of residence are also dropped. The linking procedure between the EXIOBASE product categories and the CEX UCC codes is discussed in Appendix A. Our criteria leaves us with a final sample of 81,525 households across the 23 years.

Estimation By combining estimated carbon intensities with consumption data, we are able to construct total annual CO₂ emissions per household. We then define our measure of the carbon intensity of consumption as being the kilograms of CO₂ emissions generated per dollar of expenditure. With this as our outcome variable, we estimate the following regression:

$$y_i = I_{d,i} + I_{h,i} + \mathbf{X}_i' \Gamma + I_{s,i} + I_{t,i} + \epsilon_i$$

where I_d, I_h, I_s and I_t are fixed effects for after-tax income decile, 10-year age bin, state of residence and year respectively. \mathbf{X} is a vector of controls including the number of vehicles

owned, marital status, sex, race and education level dummies, an urban-rural dummy, and linear and quadratic terms for family size. This allows us to assess heterogeneity in the carbon intensity of expenditure across income and age without imposing restrictions on the shape of the relationship, and to control for other demographic confounders which will not be present in our quantitative model.

Results Figure 1 shows the predicted carbon intensity of expenditure across income deciles and age groups, holding all other variables at their sample average level. This shows that a household being in a higher income decile is associated with a lower carbon intensity of expenditure. The differences are quantitatively substantial, with households in the top income decile seeing an almost 25% decline in the carbon intensity of their consumption. The pattern is the opposite across age groups, where the emission intensity of expenditure rises monotonically until the oldest age group, where it drops slightly. The differences are somewhat less pronounced than across income, but remain significant – the spending of the 60-70 age group is associated with around 12% more carbon emissions than the 20-30 age group.

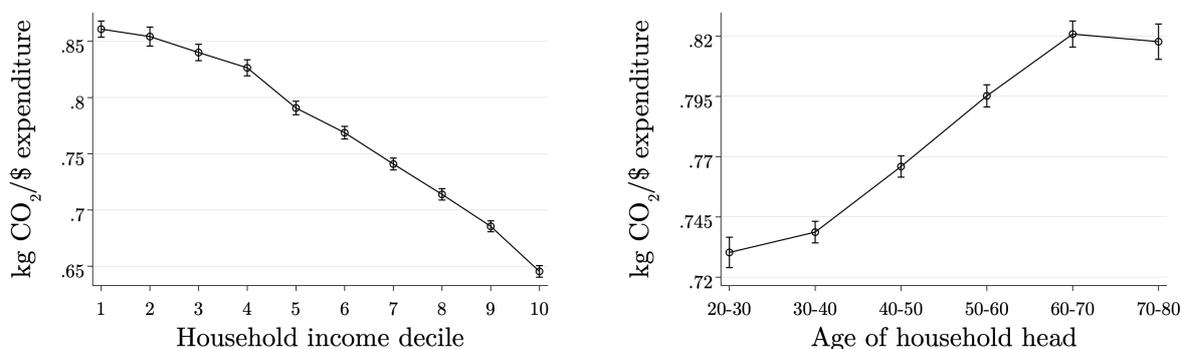


Figure 1: CO₂/\$ of expenditure

As previously noted, within-product variation in emission intensity is not captured by our consumption-based accounting due to limitations in the granularity of the data. Consequently, these differences in emission intensity cannot be the result of more or less environmentally-friendly spending within product categories, such as different food

choices. Instead, all variation along these dimensions must be the result of differences in the chosen product categories in household consumption baskets.

There are naturally many ways to disaggregate household expenditure to analyse emissions patterns. We find that a simple division into energy and non-energy expenditure is surprisingly parsimonious. We classify energy expenditure as spending on gasoline, electricity, natural gas, and other heating fuels. These are also the consumption categories for which we add tail-pipe emissions estimates, with the exception of electricity, which is a scope 2 emitter.

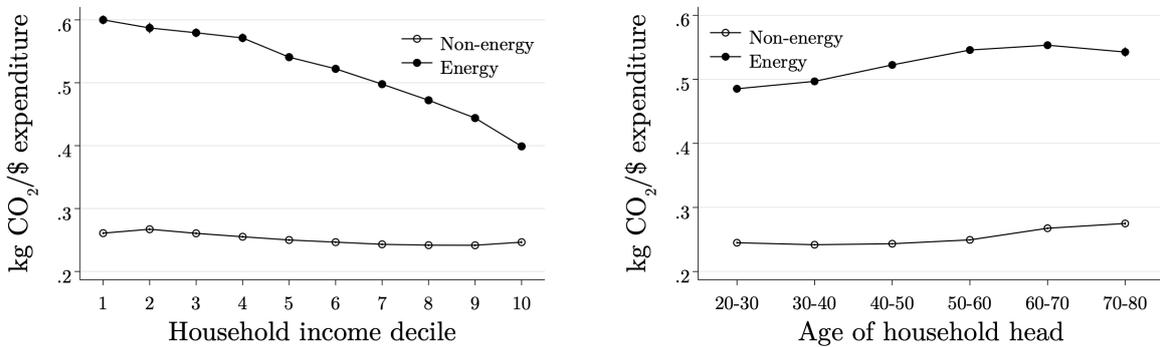


Figure 2: CO₂/\$ of expenditure, energy vs. non-energy

Figure 2 shows that the bulk of the variation in emission intensity is driven by energy emissions, which also account for a substantial majority of total household emissions. Energy emissions per dollar decline by 0.2 kg between the bottom and top income decile, compared with 0.215 kg for total emissions. Across age groups, the rise in CO₂ emissions per dollar of expenditure is around 0.06 kg for energy emissions, compared with around 0.09 kg for total emissions.

There are two possible explanations for the patterns we see in Figure 2. It could be that richer households switch to less carbon-intense energy goods, or it could be that energy makes up a smaller share of their overall expenditure – and likewise for younger households. Figure 3 shows that the latter effect dominates as households energy emissions per dollar almost exactly tracks their energy expenditure shares.

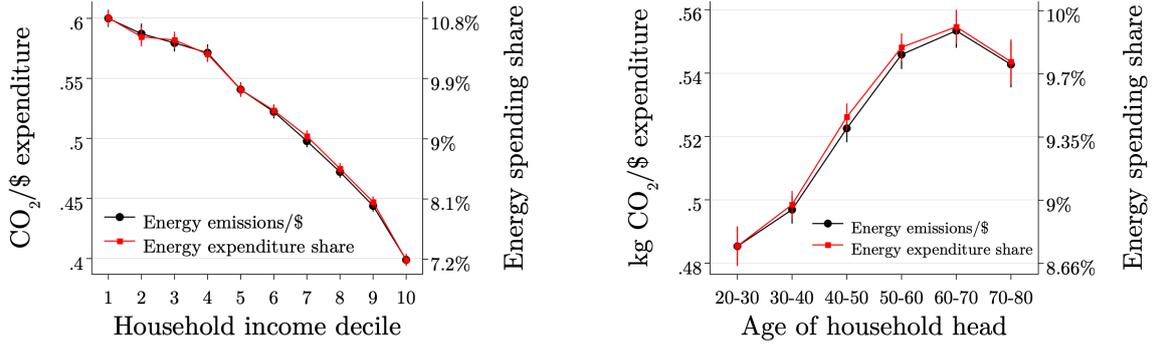


Figure 3: CO₂/\$ of expenditure from energy vs. energy share of expenditure

These figures demonstrate the stylized facts that motivate our structural model. We find that the carbon intensity of expenditure falls with income, and rises with age. These patterns are driven largely by differences in energy consumption, specifically the energy share of overall expenditure across these groups. We also find that while more complex patterns are present in non-energy expenditure (see Appendix B.1), their quantitative significance is limited, and we therefore omit this channel in the model that follows.

The next section lays out a model that rationalizes empirical patterns of energy vs. non-energy spending across income and age. We follow the literature and model the variation in energy spending across income via non-homothetic preferences between energy vs. non-energy consumption. More specifically, we assume the presence of a subsistence level in energy consumption. There are many possible mechanisms that could explain the observed pattern of energy consumption across age. In Appendix B.2, we find that a key force is the increasing need for residential heating as people age. Motivated by this finding, we model the observed energy pattern over the life cycle by assuming age-dependent subsistence levels in energy consumption. We bear in mind that this choice may have consequences for policy exercises we carry out later on.

III Model

The economy consists of a unit measure of households, two representative firms, and a government.

Households Each model household lives for a maximum of 60 periods, corresponding to a real life cycle of ages 20 to 80. Life prior to labor market entry is not modelled and everyone enters the economy with zero assets. Between periods of 1 and 40, people make consumption, saving and labor supply decisions. In period 40, they retire and start receiving a pension after which every period they only choose consumption and saving. Households face a known, age-dependent probability of death in every period, and die with certainty in period 60. The probability of surviving from period h to $h + 1$ is given by S_h . We assume that the economy starts from a stationary age distribution, which can be calculated using survival probabilities. Let the fraction of the population at age h be denoted by π_h . The assets of deceased people are distributed among the survivors in proportion to the survivors' wealth. This assumption is equivalent to assuming that people can buy an actuarially fair life insurance policy. The population grows at rate n .

Preferences Households consume two distinct goods: e stands in for energy good and c stands in for a composite non-energy good. Without loss of generality, non-energy consumption good is assigned to be the numeraire good, with a price that equals to 1 each period. Preferences over the consumption of non-energy and energy goods and labor in a period is defined according to a utility function $u_h(c, e, l)$, where the subscript h allows for preferences to depend on age, h . People discount utility across periods by $\beta \in (0, 1)$.

Production There are 2 sectors that produce non-energy and energy goods. The representative firm in sector $s = 1$ uses capital, labor and energy to produce the non-energy good. In any period t , taking rental and wage rates, r_t and w_t , as given, this firm solves:

$$(1) \quad \max_{K_{1,t}, E_{1,t}, L_{1,t}} F_{1,t}(K_{1,t}, E_{1,t}, L_{1,t}) - p_t E_{1,t} - r_t K_{1,t} - w_t L_{1,t},$$

Both capital rent and wages are paid in terms of the numeraire good. Capital depreciates at rate δ .

The representative firm in sector $s = 2$ uses capital and labor to produce the energy good. In any period t , taking rental and wage rates, this firm solves:

$$(2) \quad \max_{K_{2,t}, L_{2,t}, \mu_t} (p_t - \tau_{e,t}(1 - \mu_t))F_{2,t}(K_{2,t}, L_{2,t}) - r_t K_{2,t} - w_t L_{2,t} - \Theta_t(\mu_t E_t),$$

where μ_t is the share of abatement in energy production, $\tau_{e,t}$ is an excise tax that applies to the unabated share of energy production, and $\Theta_t(\mu_t E_t)$ is the cost of abating $\mu_t E_t$ amount of energy production in period t .

Wages Workers face idiosyncratic labor productivity shocks over time. The productivity shock, denoted by z , follows a Markov chain with states $z \in \mathcal{Z} = \{z_1, \dots, z_I\}$ and transitions $\Pi(z'|z)$. The productivity shock for labor market entrants (age 20) is drawn from a distinct entrants' distribution. There is also a life-cycle profile of wages denoted by κ_h . When a worker of age h draws productivity level z and works l units in a period, she generates $l \cdot \kappa_h \cdot z$ units of effective labor. Her wage per unit of time is $w_t \cdot \kappa_h \cdot z$.

Government In any period t , the government levies an excise tax on energy production at rate $\tau_{e,t}$. The government collects revenue also via non-linear labor income taxes, $T_t(y)$, where y is labor income, a linear consumption tax $\tau_{c,t}$ and linear tax τ_t applied to capital income net of depreciation. It is irrelevant for our analysis whether capital income is taxed at the consumer or at the corporate level. We assume without loss of generality that all capital income taxes are paid at the consumer level. We follow Heathcote, Storesletten and Violante (2017) and assume that tax liability given labor income y is defined as:

$$T_t(y) = \bar{y} \left[\frac{y}{\bar{y}} - \lambda_t \left(\frac{y}{\bar{y}} \right)^{1-\tau_l} \right],$$

where \bar{y} is the mean labor income in the economy, $1 - \lambda_t$ is the average tax rate of a mean income individual and τ_l controls the progressivity of the tax code. When $\tau_l > 0$, labor taxes are progressive and the tax function implies transfers to people with sufficiently low income. The government uses taxes to finance a stream of expenditure $\{G_t\}_{t=0}^{\infty}$,

repay government debt $\{D_t\}_{t=0}^\infty$, finance the pension system $\{P_t(\cdot)\}_{t=0}^\infty$. Pensions are also subject to the labor tax.

Asset market structure Government debt is the only financial asset in the economy. It has a one period maturity and return R_t in period t . Consumers can also save through capital. In the absence of aggregate shocks, both assets must yield the same after-tax return in equilibrium, $R_t = 1 + (r_t - \delta)(1 - \tau_t)$. As a result, one does not need to distinguish between savings via different types of assets in the consumer's problem. Consumers' (total) asset holdings will be denoted by a and $\mathcal{A} = [0, \infty)$ denotes the set of possible asset levels that agents can hold. Our assumptions imply that, in every period, the total savings of consumers must be equal to the total borrowing of the government plus the total capital stock in the economy.

Worker's problem In period t , an individual of model age $h \leq 40$ with productivity shock and asset level (z_t, a_t) solves:

$$\begin{aligned}
v_{h,t}(z_t, a_t) = & \max_{(c_t, e_t, l_t, a_{t+1}) \geq 0} u_h(c_t, e_t, l_t) + \beta S_h \sum_{z \in \mathcal{Z}} \Pi(z_{t+1}|z_t) v_{h+1,t+1}(z_{t+1}, a_{t+1}) \quad \text{s.t.} \\
& (c_t + p_t e_t)(1 + \tau_{c,t}) + a_{t+1} \leq w_t z_t \kappa_h l_t - T_t(w_t z_t \kappa_h l_t) + R_t a_t (1 + \xi_t), \\
(3) \quad & c_t, e_t, a_{t+1} \geq 0 \text{ and } l_t \in (0, 1).
\end{aligned}$$

The fact that assets of the deceased are distributed among the survivors in proportion to their wealth is captured by the fact that the return to a_t is enhanced by ξ_t in the budget constraint above where

$$\xi_t = \frac{\sum_{h=1}^{60} (1 - S_h) \pi_h \int_{\mathcal{Z} \times \mathcal{A}} a_{h+1,t+1}(z, a) d\Lambda_{h,t}(z, a)}{\sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} a_{h+1,t+1}(z, a) \cdot d\Lambda_{h,t}(z, a)}.$$

Retiree's problem In period t , an individual of age $h > 40$ with last productivity shock and current asset level (z^{40}, a_t) solves:

$$\begin{aligned}
v_{h,t}(z^{40}, a_t) &= \max_{(c_t, e_t, a_{t+1}) \geq 0} u_h(c_t, e_t, 0) + \beta S_h v_{h+1,t+1}(z^{40}, a_{t+1}) \quad \text{s.t.} \\
(c_t + p_t e_t)(1 + \tau_{c,t}) + a_{t+1} &\leq P_t(z^{40}) - T_t(P_t(z)) + R_t a_t(1 + \xi_t), \\
(4) \qquad \qquad \qquad c_t, e_t, a_{t+1} &\geq 0.
\end{aligned}$$

Competitive equilibrium Before we provide a formal definition of equilibrium, it is useful to introduce some concepts and notation. The state of a worker of age h in period t is fully described by their productivity and asset holdings. Let $(z, a) \in \mathcal{Z} \times \mathcal{A}$ denote this state. Let $\Lambda_{h,t}(z, a)$ denote the distribution of workers of age h across productivities and assets. The initial, $t = 0$, distributions are given exogenously.

Definition: Given initial conditions, a recursive competitive equilibrium is a government policy $(T_t(\cdot), P_t(\cdot), \tau_t, \tau_{c,t}, \tau_{e,t}, D_t, G_t)_{t=0}^\infty$, production plans for firms in each sector, $(K_{1,t}, E_{1,t}, L_{1,t})_{t=0}^\infty$ and $(K_{2,t}, L_{2,t}, \mu_t)_{t=0}^\infty$, value and policy functions for individuals, that is for $h = 1, \dots, 60$, $(v_{h,t}(z, a), c_{h,t}(z, a), e_{h,t}(z, a), l_{h,t}(z, a), a_{h+1,t+1}(z, a))$, a price system $(p_t, w_t, R_t)_{t=0}^\infty$ and distributions over individual states, $(\Lambda_{h,t}(z, a))_{t=0, h=1, \dots, 60}^\infty$, such that:

1. In each period $t \geq 0$, taking prices as given, $(K_{1,t}, E_{1,t}, L_{1,t})$ and $(K_{2,t}, L_{2,t}, \mu_t)$ solve firms' problems in sectors 1, 2 given by (1) and (2).
2. Given government policy and the price system, the policy functions solve individual problems given by (3) and (4), and the solutions of these problems define the value functions.
3. The evolution of distributions of agents across productivities and assets over time is consistent with agent choices. That is, for all $t \geq 0$, $h = 1, \dots, 59$ and $(z', a') \in \mathcal{Z} \times \mathcal{A}$:

$$\Lambda_{h+1,t+1}(z', a') = \sum_{z \in \mathcal{Z}} \Pi(z'|z) \int_{\{a: a_{h+1,t+1}(z, a) \leq a'\}} d\Lambda_{h,t}(z, a),$$

where $(\Lambda_{h,0}(z, a))_{h=1, \dots, 60}$ is given and the distribution of initial entrants in any period $t \geq 0$, that is $\Lambda_{1,t}$, is exogenously specified and puts a mass point at $a = 0$.

4. The stationary fraction of population at age h , denote by π_h , is defined recursively according to $\pi_{h+1} = \pi_h S_h$ for all $h = 1, \dots, 59$, which implies $\pi_1 = \frac{1}{1 + \sum_{i=1}^{59} \frac{1}{\prod_{j=1}^i S_j}}$.
5. Markets for assets, labor and goods clear: for all $t \geq 0$,

$$\begin{aligned} K_{1,t} + K_{2,t} + D_t &= \sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} a \cdot d\Lambda_{h,t}(z, a)(1 + \xi_t), \\ L_{1,t} + L_{2,t} &= \sum_{h=1}^{40} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} l_{h,t}(z, a) z d\Lambda_{h,t}(z, a), \\ C_t + X_t + G_t &= Y_{1,t} - \Theta_t(\mu_t E_t), \\ E_t &= E_{1,t} + E_{c,t}, \end{aligned}$$

where $\Theta(\cdot)$ is per-capita abatement cost, and per-capita consumption, investment and household energy consumption in period t are given by

$$\begin{aligned} C_t &= \sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} c_{h,t}(z, a) d\Lambda_{h,t}(z, a), \\ X_t &= \sum_{s=1}^2 (K_{s,t+1} - (1 - \delta)K_{s,t}), \\ E_{c,t} &= \sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} e_{h,t}(z, a) d\Lambda_{h,t}(z, a). \end{aligned}$$

6. The government's budget constraint is satisfied every period: for all $t \geq 0$,

$$G_t + R_t D_t + P_t^{agg} - D_{t+1} = \tau_{e,t}(1 - \mu_t)E_t + \tau_{c,t}(C_t + p_t E_{c,t}) + \tau_t(r_t - \delta) \sum_{s=1}^2 K_{s,t} + T_t^{agg},$$

where

$$T_t^{agg} = \sum_{h=1}^{40} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} T_t(l_{h,t}(z, a) w_t z) d\Lambda_{h,t}(z, a) + \sum_{h=41}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} T_t(P_t(z^{40})) d\Lambda_{40,t-(h-40)}(z = z^{40}, a)$$

and

$$P_t^{agg} = \sum_{h=41}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} P_t(z^{40}) d\Lambda_{40,t-(h-40)}(z = z^{40}, a)$$

are aggregate labor income tax collection and pension payments, respectively.

Equilibrium emissions The total level of carbon an economy emits over a year is a linear function of its energy use in that year. In line with this, we calculate the aggregate level of emissions in the model economy as $Em_t = (1 - \mu_t)E_tM$, where M is tons of carbon emitted per unit use of energy.

IV Parameterization

This section explains how we parameterize the model to make it aligned with the U.S. economy around 2019 along several key dimensions. We do so in two steps. First, in the external calibration stage, we set the values of a number of parameters to values from the data or from the literature. This stage includes estimating parameters that control demand for non-energy vs. energy goods using CEX data. The results of the external calibration stage are summarized in Table 1. Second, the remaining parameters are internally calibrated so that the competitive equilibrium of the model economy matches the U.S. economy around 2019 along selected dimensions that are key for our investigation. We assume that the U.S. economy is on a balanced growth path during this time in which the aggregate variables grow at the assumed population rate of 1% while the per-capita variables are stationary.⁵ The internal calibration procedure is summarized in Table 2. Whenever data is not available until 2019 for some variable, we use the most recent data. The details and definitions of the data are included in Appendix A.

⁵The existence of balanced growth path equilibrium requires the assumption that government expenditure and debt grow at the rate of population growth and taxes do not change over time.

Preferences and demographics One period in the model corresponds to one year.

The period utility function at age h is assumed to have the form:

$$u_h(c, e, l) = \frac{(c^\eta(e - \underline{e}_h)^{(1-\eta)})^{1-\sigma}}{1-\sigma} - \Psi \frac{l^{1+1/\gamma}}{1+1/\gamma},$$

where \underline{e}_h denotes subsistence level of energy consumption at age h . We set the coefficient of relative risk aversion, σ , to 2, which is within the range of values considered in the literature. The parameter that controls the Frisch elasticity of labor supply, γ , is set to 0.82, following Blundell, Pistaferri and Saporta-Eksten (2016). The age-dependent survival probability, S_h , is taken from Human Mortality Database 2019. The discount factor β and the disutility parameter Ψ are calibrated internally as discussed later. Next we explain how we choose the parameter that controls the relative weights of energy and non-energy consumption in overall consumption, η , and the age-dependent subsistence levels in energy consumption, \underline{e}_h .

Estimating preferences for energy Our Stone-Geary utility function generates a simple demand function for energy spending that we can use to estimate η and \underline{e}_h . If we denote total expenditure as x , the implied energy expenditure is $p \cdot e = \eta \cdot p \cdot \underline{e}_h + (1 - \eta)x$. The corresponding estimating equation is

$$p \cdot e = I_h + (1 - \eta)x + \mathbf{X}'\Gamma + I_s + I_t + \epsilon,$$

where I_h is a dummy variable representing the 10-year age bin an individual is in, \mathbf{X}' is the same set of demographic controls used in Section II, and I_s and I_t are state and year fixed effects.⁶ We estimate this equation using CEX data. We follow Aguiar and Bils (2015) in instrumenting for total expenditure with income and income decile.

The estimated coefficient for age bin h measures how energy spending varies by age and identifies \underline{e}_h according to $I_h = \eta \cdot p \cdot \underline{e}_h$. The relative price of the energy good in terms

⁶We estimated an alternative specification in which we allowed for η to vary by age as well. This estimation produces values of η_h that do not vary with h significantly. We, therefore, opted not to model age dependence in η .

Table 1: External Calibration

Parameter	Symbol	Value	Source
<i>Preferences and demographics</i>			
Coefficient of relative risk aversion	σ	2	
Frisch elasticity of labor supply	γ	0.82	BPS
Energy's weight in overall consumption	η	0.964	Own estimate
Age-dependent subsistence level for energy	\underline{e}_h	See Appendix C.3	Own estimate
Age-dependent survival probability	S_h	See Appendix C.3	HMD
Life-cycle profile of wages	κ_h	See Appendix C.3	Hansen (1993)
<i>Production</i>			
Capital's share of output in non-energy sector	α_1	0.3	Barrage (2019)
Energy's share of output in non-energy sector	χ	0.03	Barrage (2019)
Capital's share of output in energy sector	α_2	0.403	Barrage (2019)
Depreciation rate of capital	δ	0.1	Barrage (2019)
Cost of abatement	$\Theta(\cdot)$	see Appendix C.5	Barrage (2024)
<i>Government</i>			
Labor tax progressivity	τ_l	0.1	FN
Linear tax on consumption	τ_c	0.05	MRT
Linear tax rate on capital income	τ_k	0.36	TU
Government consumption	G/Y	0.15	NIPA
Linear tax on carbon	τ_e	0	Barrage (2019)
Pension income	$P(\cdot)$	See Appendix C.3	KK

This table reports the parameters that we take directly from the literature or estimate using the data. The acronyms BPS, FN, KK, MRT and TU stand for Blundell, Pistaferri and Saporta-Eksten (2016), Ferriere and Navarro (2024), Kindermann and Krueger (2022), Mendoza, Razin and Tesar (1994), and Trabandt and Uhlig (2011), respectively. HMD and NIPA stand for Human Mortality Database (2018) and the National Income and Product Accounts of the U.S. Bureau of Economic Analysis (2021), respectively.

of the non-energy good, p , depends on how we define the relative units of these goods; in this sense, the value of p is arbitrary. This means that, although we have an estimate of η from the above regression, this is not enough to identify p and \underline{e}_h separately. We resolve this indeterminacy by normalizing $\underline{e}_1 = 1$, which pins down p and hence the rest of the age-dependent subsistence levels. When employing these values in the model, we fit a smooth quadratic, taking each estimated \underline{e}_h as relating to the middle age in its bracket.

Production The functional form and parameter choices follow Barrage (2019). The functions that summarize production in the non-energy and energy sectors are given, respectively, by

$$\begin{aligned} F_{1,t}(K_{1,t}, E_{1,t}, L_{1,t}) &= Z_{1,t} K_{1,t}^{\alpha_1} L_{1,t}^{1-\alpha_1-\chi} E_{1,t}^{\chi} \quad \text{and} \\ F_{2,t}(K_{2,t}, L_{2,t}) &= Z_{2,t} K_{2,t}^{\alpha_2} L_{1,t}^{1-\alpha_2}. \end{aligned}$$

We set $\alpha_1 = 0.3$, $\alpha_2 = 0.403$, $\chi = 0.03$, and the depreciation rate $\delta = 0.1$, in line with Barrage (2019).⁷ Z_1 and Z_2 are internally calibrated for the initial steady state and assumed to remain constant thereafter. The function that controls the cost of abatement, Θ , is a per-ton approximation of the US abatement costs from RICE, following Barrage (2024). The details of this abatement function and its calibration are discussed in Appendix C.5.

Government The government consumption-to-output ratio is set to 15%, which is roughly in line with National Income and Product Accounts (NIPA) average for the period 2010-2019. In the status quo, the carbon tax, τ_e , is set to 0 following Barrage (2019). We follow Trabandt and Uhlig (2011) and assume that the effective status-quo tax rate on capital income is $\tau_k = 36\%$. We follow Mendoza, Razin and Tesar (1994) and assume that the consumption tax $\tau_c = 5\%$. Regarding labor income taxes, using longitudinal IRS (Internal Revenue Service) data, Ferriere and Navarro (2024) provide annual estimates of

⁷Barrage sets $\chi = 0.03$, whereas GHKT set it to 0.04. Setting it to 0.04 allows us to get energy shares on the consumer end that are below what is implied by CEX (which might be OK as there are reasons to believe that CEX might be exaggerating these shares). On the other hand, setting it to 0.04 helps matching observed allocation of energy production between firm use and consumer use. For now, we decided to keep 0.03, we may change it 0.04 or conduct robustness later.

τ_l until 2012. We use $\tau_l = 0.1$, which is the average value for the period 2010-2012. This is also consistent with Dyrda and Pugsley (2019) who estimate a progressivity parameter of slightly below 0.1 for the same period. The function that represents the pension system, $P(\cdot)$, follows Kindermann and Krueger (2022) and is assumed to be a piecewise linear function of mean earnings. For further details, see Appendix C.3. We calibrate λ , which controls the average labor tax in the economy, to clear the government budget.

Productivity shocks The class of models used in this paper inherently falls short of matching empirical earnings and wealth inequality simultaneously, especially at the top end of the corresponding distributions. Castaneda, Díaz-Giménez and Ríos-Rull (2003) propose adding a superstar individual productivity state to resolve this issue. We follow this approach. Specifically, productivity, z can be either in a normal or a superstar state. In the normal state, z follows an AR(1) process: $\log z_{t+1} = \rho \log z_t + \nu_t$ where ν_t is log-normally distributed with mean 0 and variance $var(\nu)$. which we approximate by a finite number Markov chain using the Rouwenhorst method described in Kopecky and Suen (2010). Productivity transitions to a superstar state at any given time from any normal state with probability ι_2 . When at the superstar state, productivity is ι_1 times larger than the average productivity across normal states. The probability of remaining at the superstar state is ι_3 . When agents return to a normal state, they draw a new labor market ability from the ergodic distribution associated with the AR(1) process. Labor market entrants draw their initial productivity shock from a distribution which is parameterized by ζ as explained in Appendix C.2. In total, the productivity process introduces six parameters to be calibrated.

Internal calibration In addition to the six parameters of the productivity process, we need to calibrate the production function parameters, Z_1, Z_2 , the discount factor β , and the labor disutility parameter Ψ . Although these ten parameter values are chosen simultaneously to align the model economy with the U.S. economy, it is instructive to assign parameters to moments. In this regard, the six parameters that are related to the productivity process are calibrated to match six distributional targets: the earnings Gini,

Table 2: Internal Calibration

Panel A: Moments	Data	Model
<i>Distributional moments</i>		
Earnings Gini	0.65	0.65
Wealth Gini	0.85	0.83
Wealth top 1% share	0.37	0.37
Earnings Top 1%'s share	0.17	0.17
Earnings coeff. of var. < 25	0.83	0.78
Earnings 95-99 share	0.18	0.19
<i>Aggregate moments</i>		
Capital-output ratio	2.07	2.01
Average time spent working	0.35	0.35
Relative size of energy sector in GDP	0.07	0.07
GDP per capita	\$ 65,000	\$ 65,000
Panel B: Parameters	Symbol	Value
<i>Distributional parameters</i>		
Size of superstar earnings shock	ι_1	7.6844
Prob. of transit into superstar state	ι_2	3.000×10^{-4}
Prob. of remain in superstar state	ι_3	0.2014
Persistence of normal state productivity	ρ	0.9747
Variance of productivity shock innovations	$var(\nu)$	0.0467
Age 20 productivity dispersion	ζ	0.5
<i>Aggregate parameters</i>		
Discount factor	β	0.9220
Labor disutility weight	Ψ	2.3375×10^{-4}
TFP non-energy sector	Z_1	4.9865×10^3
TFP energy sector	Z_2	0.5876
Tax function parameter	λ	2.7669

This table reports internal calibration. The model's ability to match calibration targets are reported in Panel A and the calibrated parameter values are reported in Panel B. All distributional data moments correspond to 2019 U.S. economy and are taken from Kuhn and Ríos-Rull (2022). Capital-output ratio is calculated using NIPA for 2017. Relative size of energy sector in GDP is taken from U.S. Energy Information Administration's State Energy Data System. GDP per capita is taken from World Bank database.

wealth Gini, top 1% wealth share, top 1% earnings share, earnings coefficient of variation for households aged less than 25 and the earnings share of 95-99 percentile households. Z_1 and Z_2 are calibrated to match total output per capita and energy's share in total output. β and Ψ are calibrated to match capital-output ratio and average time spent at work. Table 2 reports calibrated parameter values and the model's ability to match data targets.

IV.1 Model Fit

This section provides a further verification of the model's calibration by comparing the calibrated model to the data along a number of non-targeted moments.

Table 3: Non-Targeted Moments: Cross-sectional Inequality

	1st	A: Earnings 2nd	Quintiles 3rd	4th	5th
Data	-0.1%	2.9%	9.8%	19.1%	68.3%
Model	0.0%	2.7%	11.2%	21.4%	64.7%
	1st	B: Income 2nd	Quintiles 3rd	4th	5th
Data	2.9%	6.7%	11.1%	18.2%	61.1%
Model	4.0%	7.1%	11.1%	18.8%	58.9%
	1st	C: Wealth 2nd	Quintiles 3rd	4th	5th
Data	-0.5%	0.6%	2.9%	8.6%	88.3%
Model	0%	0.1%	2.7%	9.1%	88.0%

This table reports the fit of the model with respect to some non-targeted moments of the earnings, income and wealth distributions. Income refers to capital, labor and pension income combined. All data moments correspond to 2019 U.S. economy and are taken from Kuhn and Ríos-Rull (2022).

Cross-sectional inequality moments Table 3 summarizes the performance of the model in terms of generating the empirical cross-sectional distributions of earnings, income and wealth, which are not targeted in our calibration. Panel A reports the earnings share of each earnings quintile both in the model and in the data whereas Panel B and Panel C do the same for income and wealth. The table shows that the model is successful in reproducing the degree of inequality in earnings, income and wealth that is present in the U.S. economy.

Life-cycle moments As Figure 4 displays, the model does well in replicating the overall empirical profile of consumption over the life cycle. As in the data, the model generates a hump-shaped consumption pattern that peaks around the age of fifty although the model implies a counterfactually sharper decline in consumption after retirement.

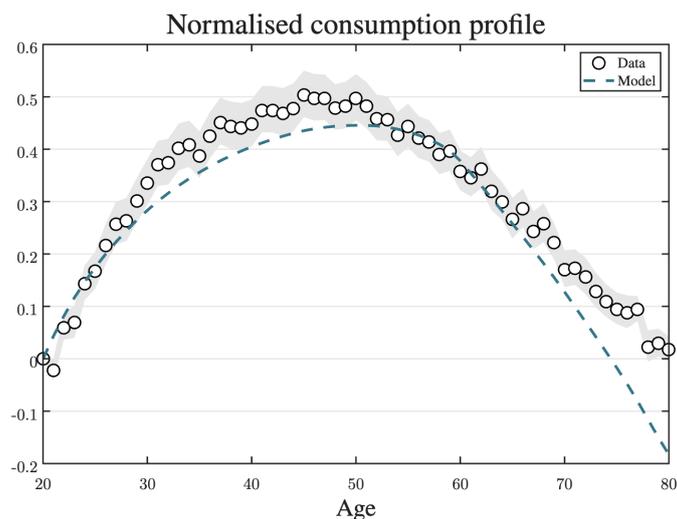


Figure 4: Consumption over the life cycle

Energy share in total expenditure A main channel through which the burden of carbon taxes falls heterogeneously across households is the expenditure channel. In this sense, it is crucial that the model captures well the share of energy spending in total household expenditure across income and age. Figure 5a shows that the model captures qualitatively the declining pattern of energy share by income although it overshoots for the lowest income decile and somewhat undershoots for the highest 8 income deciles. As Figure 5b reports, the model provides a better match to the data along the age dimension.

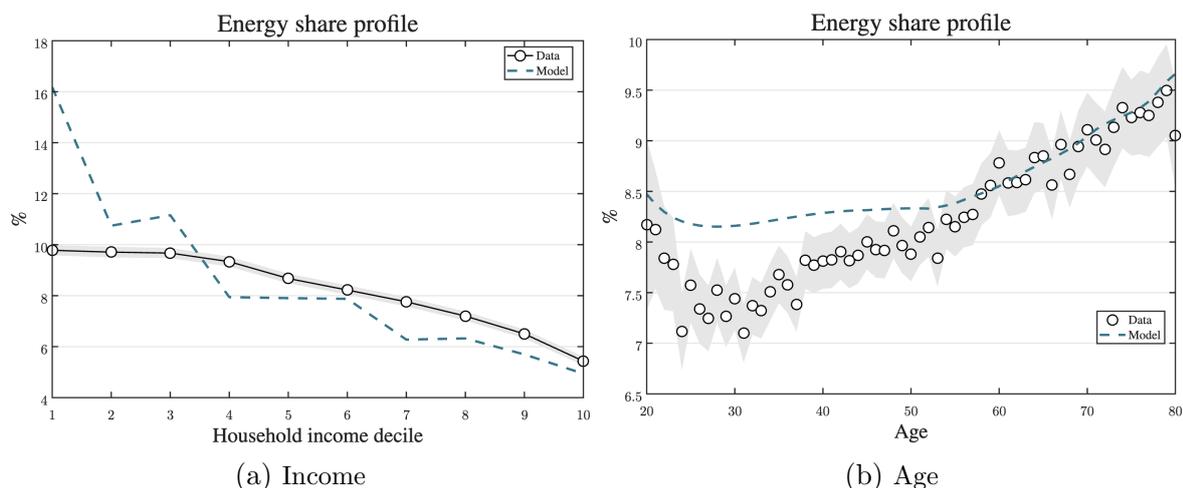


Figure 5: (a) Energy share in total household expenditure by income decile. (b) Energy share in total household expenditure by age.

Aggregate emissions This section evaluates the model's fit to 2019 US economy in terms of matching aggregate moments related to emissions and energy use. Total emis-

Table 4: Aggregate Emissions in 2019

Moment	Data	Model
Total emissions (Em)	1.43 GtC	1.57 GtC
Share of direct emissions by households in total emissions (E_c/E)	43.9%	59.1%
Share of energy emissions in household emissions ($E_c/(E_c + E_1C)$)	65.2%	69.5%
Average energy expenditure share	8.3%	8.5%

This table reports model fit regarding some key aggregate emissions and energy share moments. All data moments correspond to CO₂ emissions in 2019 U.S. economy and are taken or approximated from U.S. Environmental Protection Agency (2024), with the exception of the household energy share which is the average in our CEX sample

sions is calculated according to $Em = EM$, where M is a constant that converts energy expenditure per dollar to tons of CO₂ emissions it generates and is taken from the average emission intensity in our linked EXIOBASE-CEX dataset. The model is generally successful in generating data moments with the exception of share of emissions by households. This moment is directly related to the value of parameter χ , which controls the energy share in non-energy sector production. We take this parameter value from Barrage (2019) which results in counterfactually high emissions coming from household energy use.

V Climate Reforms

There are many climate reforms that can attain the 1.5 °C temperature change restriction put forth by the Paris Agreement. Some reforms begin with low carbon taxes and hence emissions reduction but rapidly increase the curtailment of emissions over time while some others start with high levels of emissions reduction but not increase curtailment sharply over time. In our baseline exercise, we focus on the reform that attains the particular emissions pathway that has been pledged by the US administration in their latest NDC. We then consider alternative reforms.

V.1 Baseline Reform

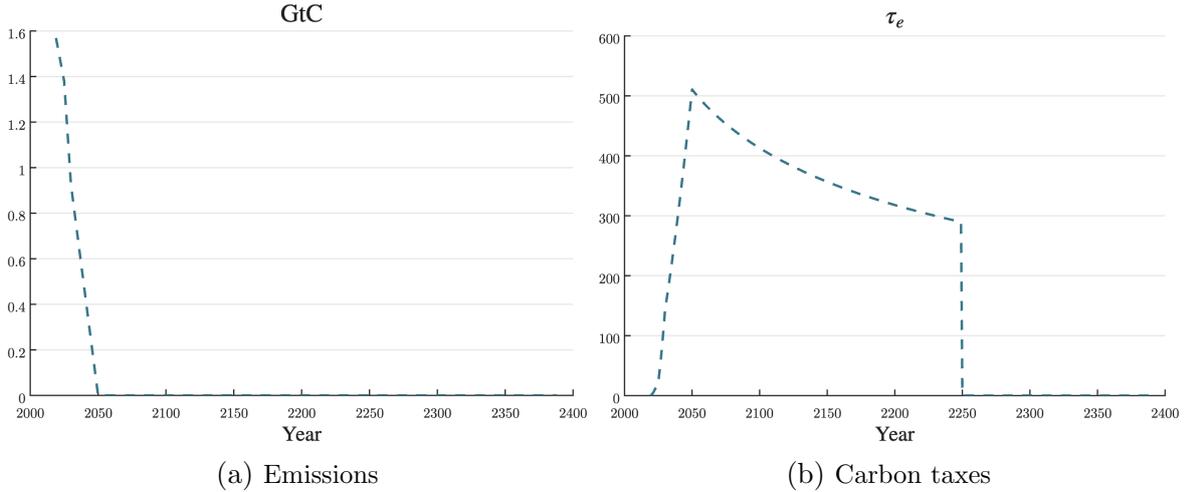


Figure 6: (a) Annual emissions targets that are prescribed in the US net-zero proposal. (b) Carbon tax sequence necessary to achieve these emissions.

Aggregate implications of the reform Figure 6a depicts the net-zero emissions pathway pledged by the US in their latest NDC document. Figure 6b displays the unique sequence of carbon taxes that achieves this emissions pathway in our model. Following Barrage (2019), we assume that carbon taxes are zero in the status quo, which describes business as usual scenario up until 2019. In order to reduce emissions in line with Figure 6a, the government introduces carbon taxes in 2020, the first year of the reform. The requirement that emissions should fall rapidly over time, and hit zero by 2050, implies that the carbon tax rate should also rise rapidly, since emissions reduction requires increasing the share of energy production that is abated. At the same time, recall that the cost of abatement is decreasing over time which means that the same carbon tax results in larger abatement, and therefore, a reduction in emissions as time goes on. This introduces a declining trend in the carbon tax trajectory that achieves net-zero. Figure 6b shows that the first force dominates all the way until 2050: an increasing pathway of carbon tax rates is necessary to generate the emissions trend we see in Figure 6a. As expected, the carbon tax rate required to maintain full abatement falls from 2050 onward and hits 0 in 2250 by assumption, following Nordhaus (2008) and Barrage (2019).

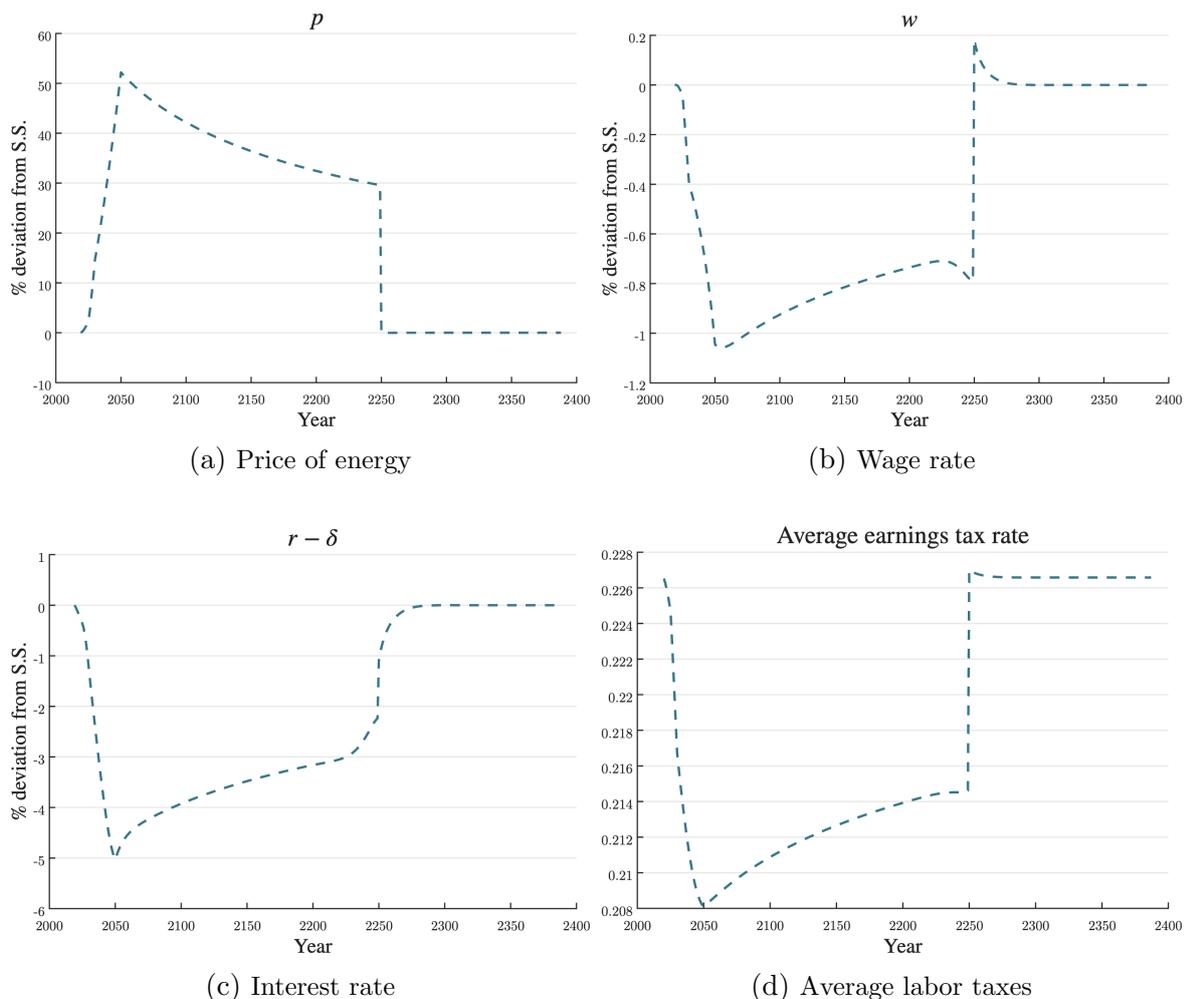
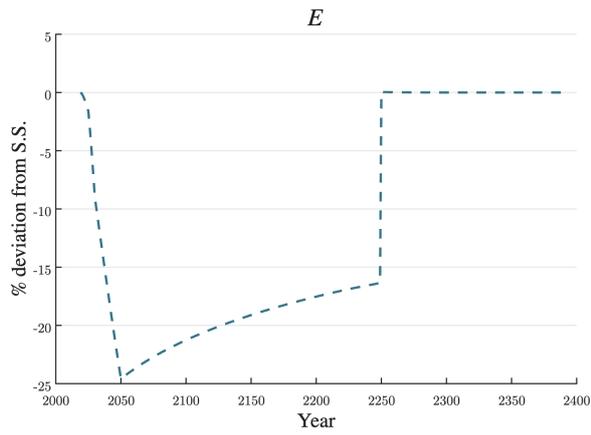


Figure 7: Dynamics of energy prices, wage rates, the net return to capital and average labor income taxes following the reform.

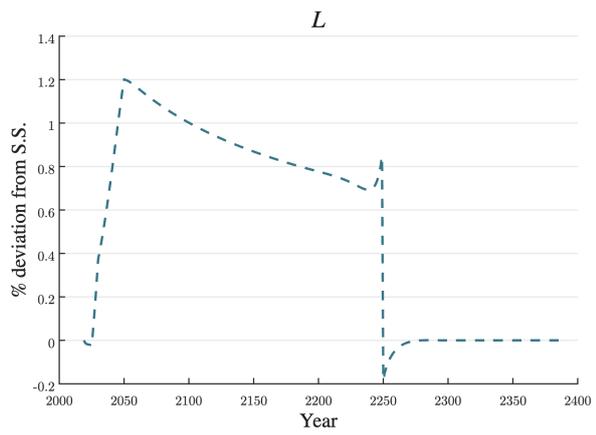
Figure 7a shows that the price of energy tracks the dynamics of the carbon tax closely: higher carbon taxes reduce the supply of energy which raises the energy price. What is remarkable is the magnitude of the response: the model predicts that the energy price will rise by 50% by 2050 in response to the net-zero reform. The dynamics of the wage and the interest rates are depicted in Figures 7b and 7c. When the carbon tax increases, this lowers the demand for capital and labor in the energy sector, reducing equilibrium level of energy, rising its price and use in the non-energy sector, which lowers the demand for capital and labor in that sector as well. As a result, we observe a decline in wage and rental rates. Even though on average energy only takes up about 10% of total consumer expenditures, the fact that the change in energy price is much larger than the change in wage and rental rates indicates that the direct welfare effect of the reform is larger than

the indirect effect. Figure 7d shows that average earnings tax rate, defined as the ratio of total earnings tax revenue to total earnings, follows the opposite pattern of carbon taxes. That is, labor income tax collection as a fraction of GDP falls until 2050, followed by a gradual increase until 2250, after which it jumps back to its initial level completely. This opposite movement of carbon and labor income taxes is intuitive as carbon tax revenue is used to reduce labor income taxes.

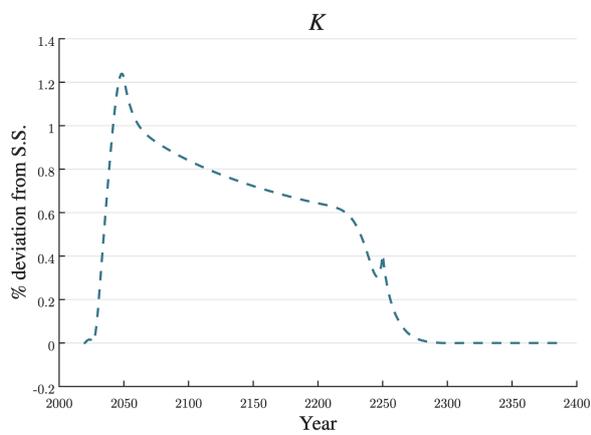
Figure 8 reports the consequences of the reform for key macroeconomic quantities. As expected, energy production falls in response to the introduction of carbon taxes. Figures 8b and 8c show that equilibrium level of labor and capital both fall on impact for a few years, which is in line with the reasoning that their demand falls in response to the decline in energy production. This is soon reversed, however, as the rise in energy prices increases labor supply due to income effects which then induces more capital accumulation. As a result, after an initial decline, we observe a rise in real output. The use of energy as an intermediate good declines substantially more than its use as a final consumption good. This likely reflects the fact that, under the functional form assumptions we make regarding preferences and non-energy good production function, consumers have less room to substitute energy consumption relative to producers.



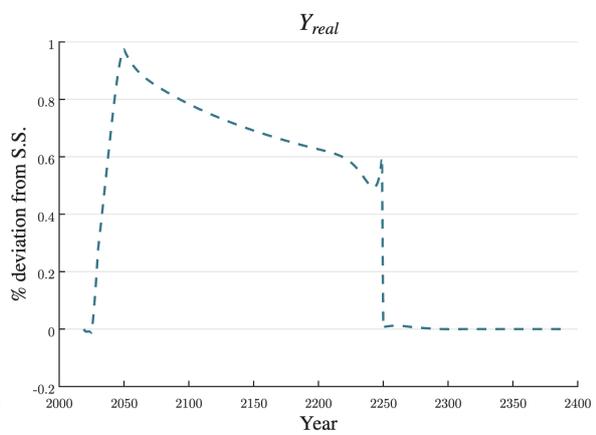
(a) Energy production



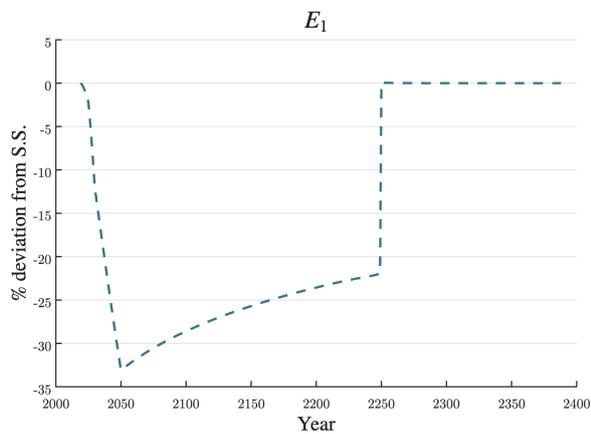
(b) Labor



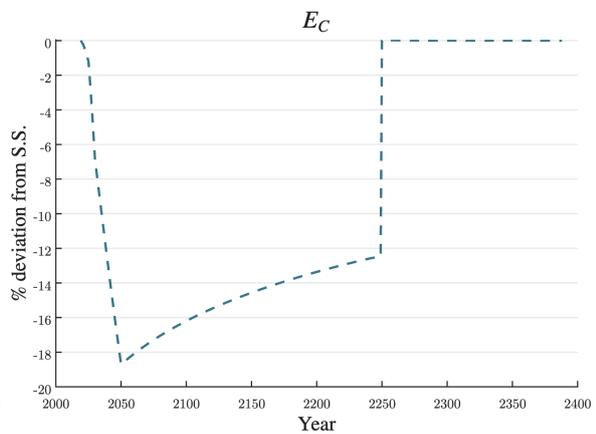
(c) Capital stock



(d) Real GDP



(e) Energy: intermediate good use



(f) Energy: final good use

Figure 8: Dynamics of aggregate quantities following the reform.

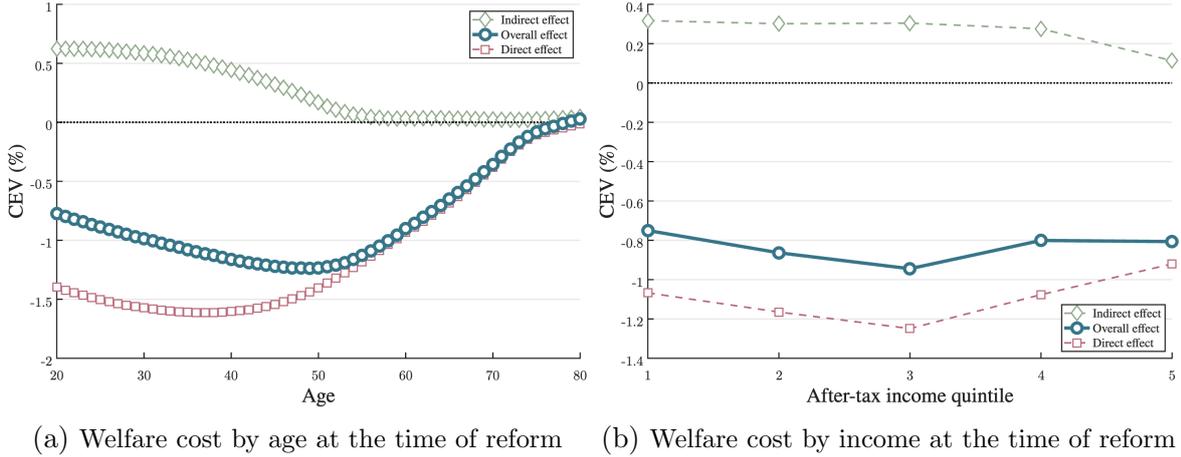


Figure 9: Welfare costs of hitting the US NDC emissions reduction path

Distributional consequences of the reform In this section, we evaluate the distributional implications of the net-zero reform. In Figures 9a and 9b, the blue curves display the discounted lifetime welfare consequences of the net-zero reform in consumption equivalent units. The reform impacts households both directly by increasing their energy expenditures and indirectly by affecting their incomes (via affecting after-tax wages, pensions and returns to investment). The red curves isolate the direct welfare effect of the reform by comparing household welfare in the pre-reform economy to the one in an auxiliary economy in which households face higher energy prices but wages, labor income taxes, pensions and interest rates are kept at the pre-reform levels. The green curves isolate the indirect welfare effects by comparing household welfare between the auxiliary economy and the post-reform economy in which all prices and taxes change. Across both income and age we find that, although the indirect channel mitigates the direct welfare costs of the reform, it is not nearly enough to offset them.

Across cohorts The overall welfare cost of the reform is substantial and has a U-shape, with those aged around 50-years-old at the time of the reform facing the highest burden, and the retired, especially those older than 70 years of age, being affected the least.

The specific timing of carbon taxes implied by the net-zero reform, depicted by Figure 6b, is largely responsible for the U-shaped pattern of the impact of the reform across cohorts. The direct welfare costs are low for the old simply because carbon taxes, and

hence the implied energy prices, are relatively low in the first 20 years of the reform when the old are still alive. These costs are largest for the middle-aged as they become old and need energy the most by around 2050, when energy prices peak. The young face high but not peak energy prices when they need energy the most, since for most of them retirement begins after energy prices start falling in 2050, and in a more distant future relative to others, which reduces the impact on them due to discounting. Regarding the indirect channel, the reform implies welfare gains through lower labor taxes and losses via lower returns to investment. This channel overall implies welfare gains for all and benefits the young more as labor income constitutes a larger share of their total income, which contributes to the U-shaped pattern of overall welfare costs.⁸

Along the income distribution The blue dashed line in Figure 9b shows that the reform hurts households fairly uniformly across income.

This may seem counterintuitive given the fact that, as depicted by Figure 5a, energy expenditure shares decline with income, which should imply that the poorest be hurt more when energy prices increase.⁹ However, at the same time, the young and the retired are hurt relatively less in the net-zero reform (recall the blue curve in Figure 9a) due to the fact that carbon taxes start relatively low in the baseline reform. Given that these demographic groups are overrepresented in low income quintiles, this implies a force for the burden of the reform to fall less on low income groups. Furthermore, as the green curve in Figure 9b shows, the indirect welfare effects of the reform disfavor higher income groups as they benefit less from the decline in labor income taxes and lose more from lower returns to investment (since a larger share of their total income stems from capital income). A combination of these effects implies the rather flat impact of the net-zero reform along the income distribution.

⁸A precise breakdown and further explanation of the indirect effects over the life cycle is relegated to Appendix D.

⁹Similarly, this may seem to be at odds with other papers in the literature which find that carbon taxes are regressive

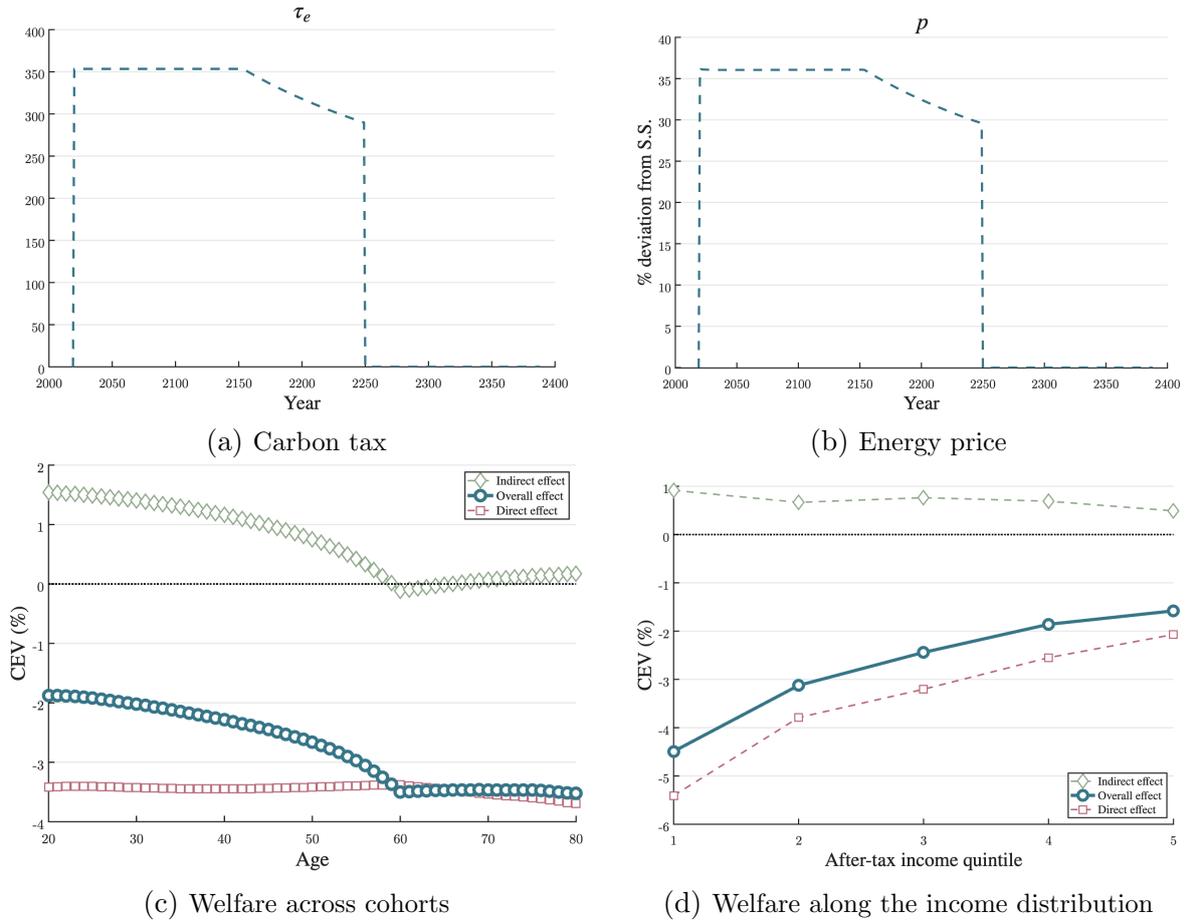


Figure 10: Welfare costs of a once-and-for-all carbon tax

Constant tax reform In order to highlight the importance of the specific timing of carbon taxes imposed by the net-zero reform, we also consider an alternative reform in which, in 2020, the government introduces a carbon tax that stays constant over time. The level of the tax is chosen to achieve the same mitigation in global temperature increase by 2100 as in our baseline exercise.¹⁰ Figure 10a shows that, relative to the baseline reform, the carbon tax starts much higher at \$350 per ton of emissions while it stays lower in later years. The front-loading of carbon taxes in the flat reform is also reflected in the energy price which is higher earlier on and lower later on relative to the baseline reform. With carbon taxes held constant over time, the direct welfare implications of the reform across cohorts, displayed by Figure the 10c, are closely related to the model implied energy spending shares across the life cycle depicted in 5b: those who are old at the time of the reform suffer the most as they have the highest immediate

¹⁰Details on the calculation of temperature change are provided in Appendix C.4

energy needs and the young suffer the least as they have lowest immediate energy needs. The indirect effect of the reform is mainly driven by the rise in after-tax wages and this benefits younger households more, as labor income is a larger part of their overall income, and benefits the retired least as they do not work. Figure 10d shows that carbon taxes are regressive under the constant tax reform. This again reflects the spending patterns displayed in Figure 5a: the poor spend disproportionately more on energy and hence are hurt more when the reform raises energy prices. This finding is in line with the literature on the regressivity of higher energy prices. We also note that the overall welfare costs of the reform on currently alive generations are much larger in the flat reform, which is not surprising as under this reform carbon taxes are tilted toward current generations relative to the net-zero reform.

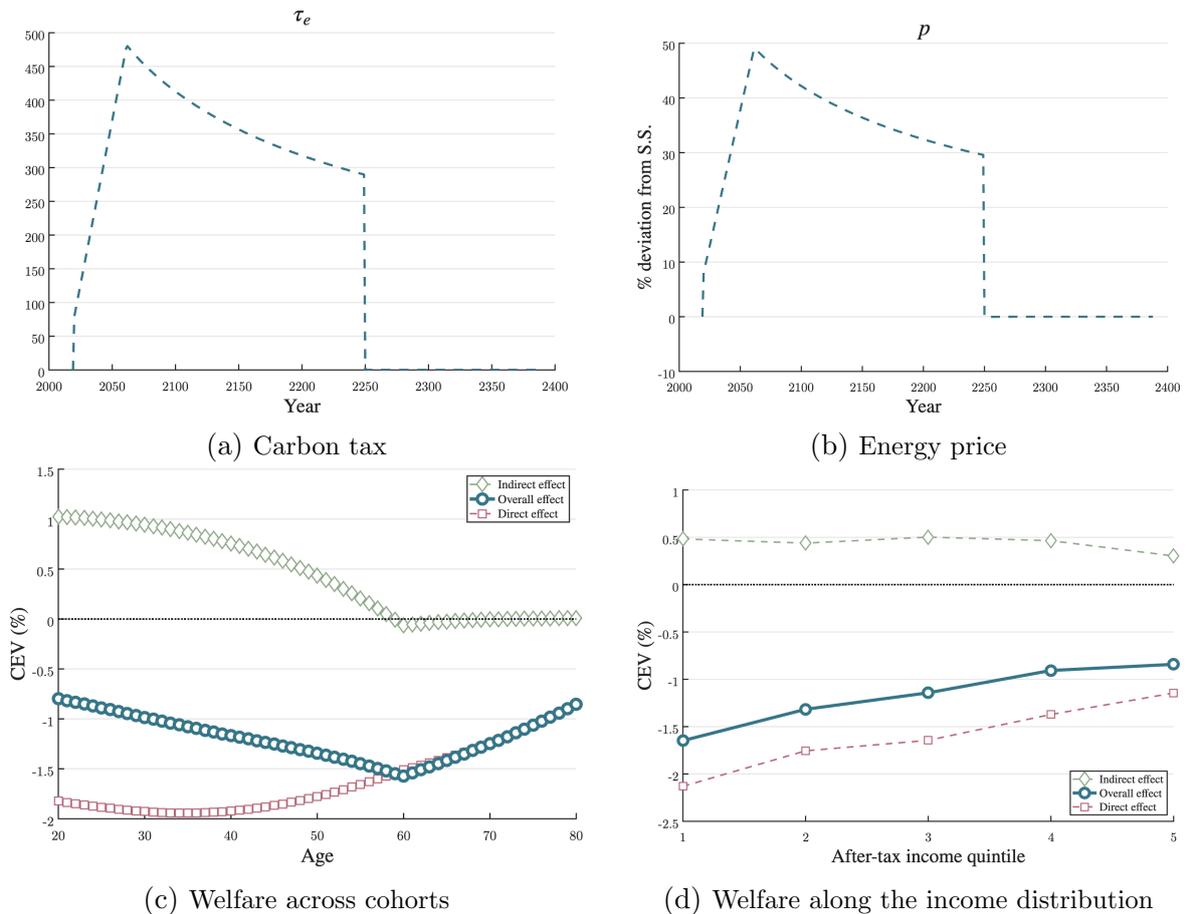


Figure 11: Welfare costs of the optimal reform

Optimal reform In this section, we consider the problem of a government that aims to maximise a social welfare function while keeping temperature at the level implied by

the US NDC (our baseline reform). The specific welfare aggregator we consider is

$$\sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} v_{h,0}(z, a; \psi) d\Lambda_{h,0}(z, a) + \pi_1 \sum_{t=1}^{\infty} \hat{\beta}^t \int_{\mathcal{Z} \times \mathcal{A}} v_{1,t}(z, a; \psi) d\Lambda_{1,t}(z, a).$$

We believe that $\hat{\beta} = \beta$ is a natural case to consider since this amounts to assuming that the social discount factor that applies to utility from consumption in a given period only depends on the period and is independent of the cohort that enjoys that consumption. We restrict ourselves to linear tax paths for computational tractability.

Figures 11a and 11b reveal that carbon taxes and the implied change in energy prices are more front-loaded compared to the baseline reform while they are more back-loaded compared to the flat tax reform. Figure 11c demonstrates that, like in the baseline reform, the direct welfare costs are lowest for the old because carbon taxes, and hence the implied energy prices, are relatively low in the first 20 years of the reform when the old are still alive. They are still higher for this group relative to the baseline since carbon taxes start higher in the optimal reform. The age at which the burden of the reform is felt the most is lower relative to the baseline reform in line with the fact that under the optimal reform carbon taxes peak about a decade later than they do in the baseline reform. As in the prior exercises, the indirect channel effect of the reform works mainly through the rise in after-tax wages which benefits the young more as labor income constitutes a larger share of their total income.

Overall, we see a U-shaped pattern across age of household at the time of reform with 60 year old households being hurt the most. Figure 11d reports the welfare cost of the net-zero reform across the income distribution at the time of the reform. The blue dashed line shows that the reform is regressive, as in the flat tax reform. This follows the direct effect of the reform which is also regressive. The reason for this is that, although as in the baseline reform, the young and the old are hurt the least and this works as a force to lower the welfare burden on the poor, this channel is not quantitatively sufficient to counteract the fact that poorer people use disproportionately more energy.

Importance of the carbon tax pathway A key insight that follows from our exercises is that the distributional impacts of a climate reform depend crucially on the pathway of carbon tax rates it imposes. In particular, reforms that involve flatter (more front-loaded) carbon tax pathways imply front-loaded increases in energy prices, which tend to hurt households that are older and poorer at the time of the reform. On the other hand, reforms that start with lower carbon taxes initially (which requires steeper increase in carbon tax rates over time) induce back-loaded increases in energy prices, which tends to imply a burden of climate policies across cohorts that is more U-shaped with middle aged being hurt the most and is less regressive.

Previous papers that assess the impact of one-time energy price shocks, or compute welfare effects only in the short-run, fail to capture this. While these case studies are a good approximation of the distributional impact of a once-and-for-all carbon tax, they miss important effects from the interaction of income and cohort for evaluating a gradual transition to net zero.

VI Conclusion

This paper studies the distributional consequences of carbon taxation over the life cycle. We document that the carbon intensity of consumption declines with income but rises with age, reflecting systematic differences in energy expenditure shares. We build a heterogeneous-agent overlapping-generations model with non-homothetic preferences and age-dependent energy needs that matches these patterns and use it to evaluate carbon tax paths consistent with a transition to net-zero emissions.

We find that higher energy prices disproportionately affect poorer and older households in the short run. In lifetime terms, however, welfare costs are broadly similar across income groups and follow a U-shaped profile over age, with middle-aged cohorts most affected. A central result is that the timing of carbon taxes matters: front-loaded tax paths are more regressive and hurt the oldest households most, while back-loaded paths

shift the burden toward middle-aged cohorts and richer households, even when long-run emissions targets are the same.

References

- Aguiar, Mark, and Mark Bils.** 2015. “Has Consumption Inequality Mirrored Income Inequality.” *American Economic Review*, 105(9): 2725–2756.
- Barrage, Lint.** 2019. “Optimal Dynamic Carbon Taxes in a Climate–Economy Model with Distortionary Fiscal Policy.” *The Review of Economic Studies*, 87(1): 1–39.
- Barrage, Lint.** 2024. “The Fiscal Costs of Climate Change in the United States.” CER-ETH WP 23-380.
- Belfiori, E., D. Carroll, and S. Hur.** 2024. “Unequal Climate Policy in an Unequal World.” FRB of Dallas FRB of Dallas working paper 447.
- Benmir, G, and J Roman.** 2022. “The Distributional Costs of Net-Zero: A Heterogeneous Agent Perspective.”
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten.** 2016. “Consumption Inequality and Family Labor Supply.” *American Economic Review*, 106(2): 387–435.
- Carbone, Jared C., Richard D. Morgenstern, Roberton C. Williams, and Dallas Burtraw.** 2013. “Deficit Reduction and Carbon Taxes: Budgetary, Economic, and Distributional Impacts Considering a Carbon Tax: A Publication Series from RFF’s Center for Climate and Electricity Policy.”
- Castaneda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull.** 2003. “Accounting for the U.S. Earnings and Wealth Inequality.” *Journal of Political Economy*, 111(4): 818–857.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger.** 2009. “Taxing Capital? Not a Bad Idea after All!” *American Economic Review*, 99(1): 25–48.
- Douenne, Thomas, Sebastian Dyrda, Albert Jan Hummel, and Marcelo Pedroni.** 2024. “Distributional Consequences of Climate Policies.”
- Ferriere, Axelle, and Gaston Navarro.** 2024. “The Heterogeneous Effects of Government Spending: It’s All About Taxes.” *The Review of Economic Studies*, 92(2): 1061–1125.
- Fried, Stephanie, Kevin Novan, and William B. Peterman.** 2018. “The distributional effects of a carbon tax on current and future generations.” *Review of Economic Dynamics*, 30: 30–46.
- Fried, Stephanie, Kevin Novan, and William B. Peterman.** 2024. “Understanding the Inequality and Welfare Impacts of Carbon Tax Policies.” *Journal of the Association of Environmental and Resource Economists*, 11(S1): 231–260.

- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski.** 2014. “Optimal Taxes on Fossil Fuel in General Equilibrium.” *Econometrica*, 82(1): 41–88.
- Hansen, G D.** 1993. “The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked.” *Journal of Applied Econometrics*, 8(1): 71–80.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2017. “Optimal Tax Progressivity: An Analytical Framework.” *The Quarterly Journal of Economics*, 132(4): 1693–1754.
- Kaymak, Barış, David Leung, and Markus Poschke.** 2026. “Accounting for Wealth Concentration in the United States.” *American Economic Journal: Macroeconomics*, 18(1): 187–222.
- Kindermann, Fabian, and Dirk Krueger.** 2022. “High Marginal Tax Rates on the Top 1 Percent? Lessons from a Life-Cycle Model with Idiosyncratic Income Risk.” *American Economic Journal: Macroeconomics*, 14(2): 319–66.
- Kopecky, Karen, and Richard Suen.** 2010. “Finite State Markov-chain Approximations to Highly Persistent Processes.” *Review of Economic Dynamics*, 13(3): 701–714.
- Kotlikoff, Laurence, Felix Kubler, Andrey Polbin, Jeffrey Sachs, and Simon Scheidegger.** 2021. “MAKING CARBON TAXATION A GENERATIONAL WIN WIN.” *International Economic Review*, 62(1): 3–46.
- Kuhn, Moritz, and José-Víctor Ríos-Rull.** 2022. “2019 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics.” Working paper. <https://www.wiwi.uni-bonn.de/kuhn/paper/QR2020update.pdf>.
- Kuhn, Moritz, and Lennard Schlattmann.** 2024. “Distributional Consequences of Climate Policies.” C.E.P.R. Discussion Papers CEPR Discussion Papers 18893.
- Leach, Andrew.** 2009. “The welfare implications of climate change policy.” *Journal of Environmental Economics and Management*, 57(2): 151–165.
- Levinson, Arik, and James O’Brien.** 2019. “Environmental Engel Curves: Indirect Emissions of Common Air Pollutants.” *The Review of Economics and Statistics*, 101(1): 121–133.
- Mendoza, Enrique G., Assaf Razin, and Linda L. Tesar.** 1994. “Effective Tax Rates in Macroeconomics: Cross-Country Estimates of Tax Rates on Factor Incomes and Consumption.” *Journal of Monetary Economics*, 34(3): 297–323.
- Millar, Richard J, Zebedee R Nicholls, Pierre Friedlingstein, and Myles R Allen.** 2017. “A modified impulse-response representation of the global near-surface air temperature and atmospheric concentration response to carbon dioxide emissions.” *Atmospheric Chemistry and Physics*, 17(11): 7213–7228.
- Nordhaus, William.** 1994. *Managing the Commons: The Economics of Climate Change*. The MIT Press.
- Nordhaus, William.** 2007. *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

- Nordhaus, William.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Rausch, Sebastian.** 2013. “Fiscal consolidation and climate policy: An overlapping generations perspective.” *Energy Economics*, 40(S1): 134–148.
- Sager, Lutz.** 2019. “Income Inequality and Carbon Consumption: Evidence from Environmental Engel Curves.” *Energy Economics*, 84(S1).
- Trabandt, Mathias, and Harald Uhlig.** 2011. “The Laffer Curve Revisited.” *Journal of Monetary Economics*, 58(4): 305–327.
- U.S. Bureau of Economic Analysis.** 2021. “National Income Accounts.” <https://www.bea.gov/itable/national-gdp-and-personal-income>. Last accessed: 2021-12-21. For details, see the online appendix and the readme.pdf file in the replication package.
- U.S. Environmental Protection Agency.** 2024. “Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990–2022.”

Appendix

A Appendix: Data Construction

We closely follow Levinson and O'Brien (2019) in linking CEX products to our EE-MRIO data. Levinson and O'Brien provide a concordance between UCC codes (the product classification in the CEX) and the US Input-Output Table. A similar concordance is provided by EXIOBASE for linking their product categories to USEEIO categories. We combine these crosswalks to assign emissions to each UCC good. When a good in the USEEIO data is assigned multiple EXIOBASE goods, we assign its emission intensity as the weighted average of the emission intensity of the EXIOBASE goods, where the weights are the aggregate expenditures on each category. The USEEIO goods can then be mapped cleanly to CEX products using Levinson and O'Brien's crosswalk

B Appendix: Further Empirical Results

B.1 Non-energy expenditure

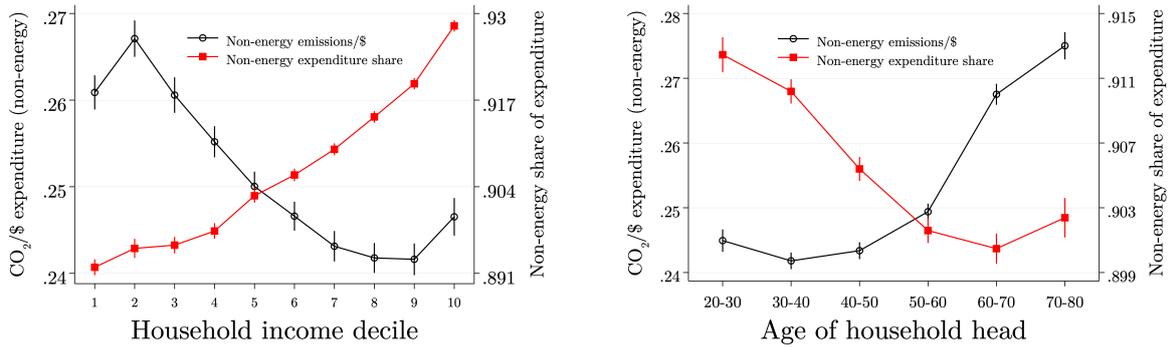


Figure 12: CO₂/\$ of expenditure from non-energy vs. non-energy share of expenditure

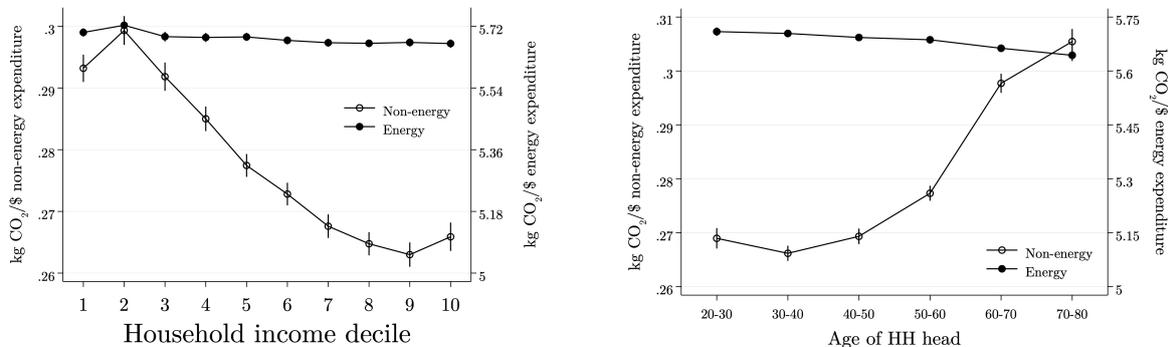


Figure 13: Emission intensity of energy/non-energy expenditure

B.2 Decomposing energy emissions by age

The increasing pattern of energy emissions by age appears to be primarily driven by their increasing expenditure on residential energy (defined here as electricity and natural gas), as shown in Figure 14

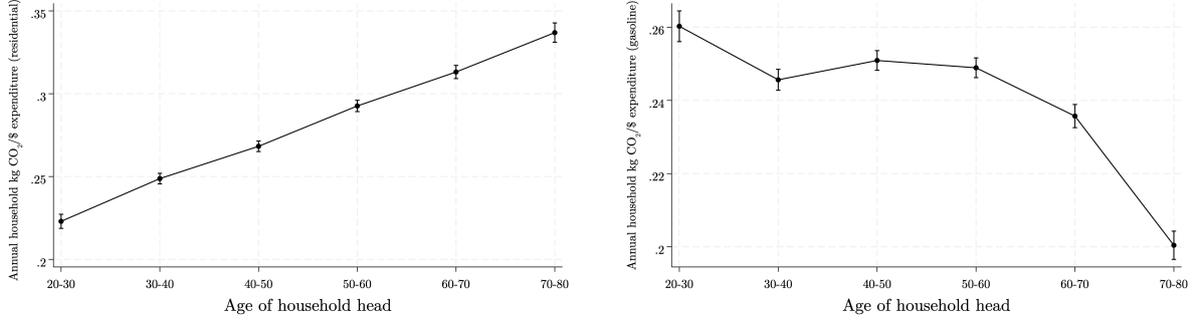


Figure 14: Life-cycle energy emissions decomposed

C Appendix: Calibration

C.1 Life-cycle parameters

We fit a quadratic to the piecewise linear values for κ from Conesa, Kitao and Krueger (2009). We estimate the values of e_h in 10-year age bins and fit a quartic, normalizing e_h at age 20 to 1. The resulting values are displayed in Table 5

C.2 Initial productivity distribution

The parameter ζ determines the productivity dispersion at model age 1 in the following way. If the vector of possible productivity levels at age 1 is \mathbf{z} , where ι_1 is the superstar shock, then the distribution of productivity levels at age 1 is given by Π_1

$$\mathbf{z} = [z_1 \quad z_2 \quad z_3 \quad z_4 \quad z_5 \quad z_6 \quad \iota_1]$$

$$\Pi_1 = [\zeta/4 \quad \zeta/4 \quad (1-\zeta)/2 \quad (1-\zeta)/2 \quad \zeta/4 \quad \zeta/4 \quad 0]$$

i.e., we assume that no one starts life in the superstar state. This is the same way that Kaymak, Leung and Poschke (2026) handle the initial distribution.

C.3 Pension system

Following Kindermann and Krueger (2022), pensions are a piecewise linear function of mean earnings, as with the true social security formula in the United States. Households receive a pension in accordance with their final productivity realization during working age. The levels of pension payments are determined as follows:

$$P(\theta) = \begin{cases} r_1 \bar{y}(\theta) & \text{if } \bar{y}(\theta) \leq b_1 y^{med}, \\ r_1 b_1 y^{med} + r_2 (\bar{y}(\theta) - b_1 y^{med}) & \text{if } b_1 y^{med} \leq \bar{y}(\theta) \leq b_2 y^{med}, \\ \min\{r_1 b_1 y^{med} + r_2 (b_2 - b_1) y^{med} + r_3 (\bar{y}(\theta) - b_2 y^{med}), 0.6 \times y^{med}\} & \text{otherwise.} \end{cases}$$

Age	κ	e_h	S_h
20	0.9986	1.0000	0.9992
21	1.0801	1.0350	0.9991
22	1.1577	1.0706	0.9991
23	1.2312	1.1065	0.9990
24	1.3009	1.1426	0.9990
25	1.3669	1.1786	0.9989
26	1.4291	1.2144	0.9989
27	1.4876	1.2497	0.9989
28	1.5426	1.2844	0.9988
29	1.5940	1.3184	0.9987
30	1.6420	1.3516	0.9987
31	1.6867	1.3836	0.9986
32	1.7280	1.4145	0.9986
33	1.7661	1.4442	0.9985
34	1.8011	1.4724	0.9985
35	1.8329	1.4991	0.9984
36	1.8618	1.5242	0.9983
37	1.8876	1.5477	0.9982
38	1.9106	1.5693	0.9981
39	1.9308	1.5891	0.9980
40	1.9483	1.6071	0.9979
41	1.9630	1.6231	0.9979
42	1.9752	1.6371	0.9978
43	1.9848	1.6490	0.9977
44	1.9920	1.6590	0.9976
45	1.9968	1.6668	0.9974
46	1.9992	1.6726	0.9972
47	1.9994	1.6764	0.9969
48	1.9974	1.6781	0.9967
49	1.9933	1.6777	0.9964
50	1.9872	1.6754	0.9961

Age	κ	e_h	S_h
51	1.9791	1.6710	0.9957
52	1.9690	1.6648	0.9954
53	1.9572	1.6567	0.9950
54	1.9436	1.6468	0.9944
55	1.9283	1.6352	0.9939
56	1.9113	1.6219	0.9934
57	1.8928	1.6071	0.9928
58	1.8728	1.5908	0.9923
59	1.8514	1.5732	0.9916
60	0.0000	1.5543	0.9910
61	0.0000	1.5344	0.9902
62	0.0000	1.5135	0.9894
63	0.0000	1.4918	0.9887
64	0.0000	1.4694	0.9880
65	0.0000	1.4466	0.9872
66	0.0000	1.4234	0.9863
67	0.0000	1.4000	0.9854
68	0.0000	1.3767	0.9845
69	0.0000	1.3537	0.9834
70	0.0000	1.3311	0.9817
71	0.0000	1.3092	0.9800
72	0.0000	1.2881	0.9784
73	0.0000	1.2682	0.9763
74	0.0000	1.2497	0.9743
75	0.0000	1.2327	0.9710
76	0.0000	1.2177	0.9682
77	0.0000	1.2048	0.9652
78	0.0000	1.1944	0.9621
79	0.0000	1.1866	0.9580
80	0.0000	1.1819	0.0000

Table 5: Life-cycle labor efficiency, subsistence energy demand and survival probabilities

where y^{med} is median earnings, $\bar{y}(\theta)$ is the mean earnings among households of productivity type θ , $\{r_1, r_2, r_3\}$ are progressive replacement rates, and $\{b_1, b_2\}$ are ‘bend points’ in the pension system. There is also a cap on pension payments, set as a share of median earnings. The values of these parameters are taken from Kindermann and Krueger (2022).

C.4 Temperature change

C.4.1 Model implied emissions

Our counterfactual exercises rely on matching a certain level of temperature change by 2100. In order to do this, we need to convert our model units of energy, E , into the temperature change they imply. Linking the CEX with emission intensities from EX-IOBASE, as described in Appendix A, we find that one dollar of household expenditure on energy generated 3.83 kg CO₂ on average in 2019. This value multiplied by the model

implied price of energy, p , gives us kg CO₂ per household, per unit of E . We can multiply this by the number of US households in 2019 to calculate the aggregate emissions implied by our model.

C.4.2 Mapping emissions to temperature change

Once aggregate emissions have been calculated, we follow Barrage (2024) in employing a slightly adjusted version of the DFAIR module of DICE-2023, which is based in turn on the FAIR model of Millar et al. (2017). We reproduce the system of equations exactly from Barrage (2024) here for clarity.

In this model, emissions add to four excess carbon ‘reservoirs’

$$R_{i,t} = a_i \tau_i \alpha_t \frac{E_t^M + E_t^{M,ROW}}{\Delta} (1 - \exp\left(\frac{-\Delta}{\alpha_t \tau_i}\right)) + R_{i,t-1} \exp\left(\frac{-\Delta}{\alpha_t \tau_i}\right)$$

These reservoirs increase atmospheric carbon concentration, measured relative to pre-industrial levels

$$MAT_t = MAT_{1765} + \sum_{i,t} R_{i,t}$$

Atmospheric carbon concentration, along with other forcings taken as exogenous (e.g., methane) contribute to two ‘temperature boxes’ which increase global temperature, T_t

$$\begin{aligned} FORC_t &= \frac{F_{2x}}{\ln(2)} \ln\left(\frac{MAT_t}{MAT_{1765}}\right) + FORC_t^{EX} \\ Tbox_{j,t+1} &= Tbox_{j,t} \cdot \exp\left(\frac{-\Delta}{d_j}\right) + teq_j \cdot FORC_{t+1} \left(1 - \exp\left(\frac{-\Delta}{d_j}\right)\right) \quad \text{for } j \in \{1, 2\} \\ T_t &= Tbox_{1,t} + Tbox_{2,t} \end{aligned}$$

The parameter α_t , which affects the transmission of emissions into the carbon reservoirs, is endogenous to previous emissions, carbon concentration and temperature change in the following way:

$$C_{acc,t} = \sum_{\nu=1765}^t (E_{\nu}^M + E_{\nu}^{M,ROW}) - (MAT_t - MAT_{1765})$$

where $C_{acc,t}$ is carbon accumulated outside of the atmosphere

$$iIRF_{100,t} = irf_0 + irf_C C_{acc,t} + irf_T T_t$$

where $iIRF_{100,t}$ is the predicted 100 year integrated impulse response function

$$\hat{\alpha}_t = 0.01374981 \cdot \exp(0.0884138 \times iIRF_{100,t-1})$$

All other parameters are exogenous. We take the initial conditions from Barrage (2024) and iterate forward with a given model implied path of emissions to calculate temperature change by 2100. We leave emissions from the rest of the world, E_{ROW}^M , at their business-as-usual level as defined by RICE-2010.

C.5 Abatement costs

We use the abatement cost function estimated in Barrage (2024). This takes the abatement costs in the RICE model, which are defined in terms of shares of business-as-usual output, and converts them into a cost function defined in dollars per ton of carbon. There is a backstop price, P^{backstop} which equals the marginal cost of abatement when fully abating emissions. This is taken directly from RICE (2010), and is equal to \$1,134 per tonne of carbon in 2005\$. The remaining parameters are selected to minimize the distance to the RICE abatement costs. We follow Barrage in assuming that the backstop cost declines by 0.5% per year and equals zero from 2250 onward, hence the time subscripts on the parameters.

$$\Theta_t(\mu_t E_t) = \frac{\bar{a} P_t^{\text{backstop}}}{1 + a_t \exp(b_{0t} - b_{1t}(\mu_t E_t)^{b_2})} \cdot (\mu_t E_t)^{b_x}$$

\bar{a}	0.0662
P_t^{backstop}	$0.567(1 + \exp(-0.05 \cdot t))$
a_t	$49.8896 + 0.8551 \log(t)$
b_{0t}	$14.3338 - 6.498 \log(t)$
b_{1t}	$15.1937 - 6.6864 \log(t)$
b_2	9.468×10^{-4}
b_x	2.6931

Table 6: Parameters for $\Theta_t(\mu_t E_t)$

Barrage’s cost function is designed to take gigatonnes of carbon emissions over a decade and return costs in thousands of dollars. As our model is in per-capita terms and at annual frequency, we scale the emissions from our model up to the implied aggregate emissions over a decade, and then convert the implied costs back down to our model scale.

D Appendix: Decomposition of Indirect Welfare Gains

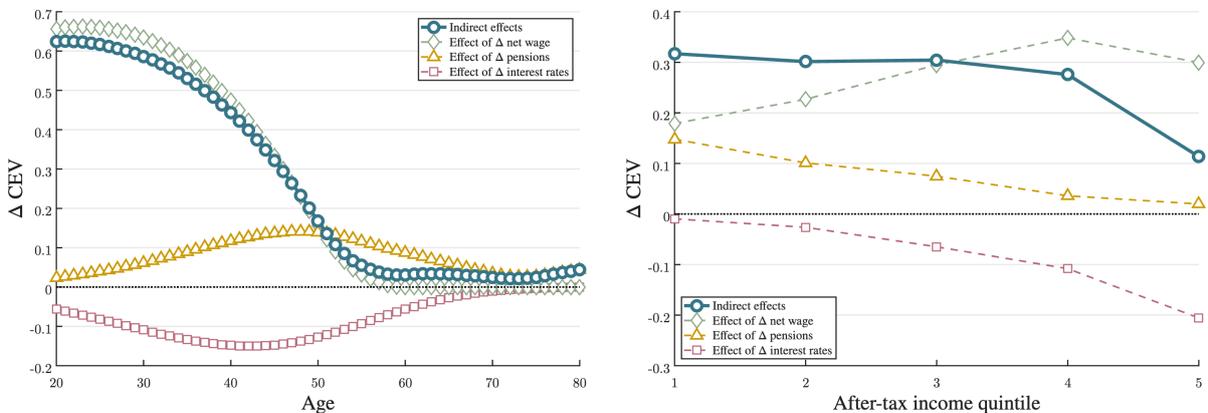


Figure 15: Decomposition of indirect welfare effects