

The Distributional Impacts of Net-Zero Transitions

Stanley Barrell-Kane and Hakki Yazici*

September 29, 2026

This paper studies the distributional consequences of net-zero transitions. Using an environmentally-extended input-output table linked with the U.S. Consumer Expenditure Survey, we show that the carbon intensity of expenditure declines with income, largely due to differences in energy expenditure shares. We develop a heterogeneous-agent overlapping-generations model with non-homothetic preferences, estimated using household-level data, that matches key distributional moments of the U.S. economy. We evaluate a carbon tax path consistent with the U.S. net-zero target under the Paris Agreement. Although the reform disproportionately affects poorer households in the short run, lifetime welfare costs are similar across income groups. Across cohorts, those in their fifties when the reform is introduced experience the largest losses. We find that the time path of carbon taxes impacts the distributional consequences of climate reforms substantially. An optimal reform that maximizes average welfare almost fully coincides with the carbon tax path implied by the Paris Agreement.

JEL classification: E62, H21, H23, D52

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

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

We thank Michal Kejak, Ctirad Slavik and seminar participants at Bilkent University, CERGE-EI, the University of Bristol and the Growth and Public Economics Workshop at the University of Minnesota for valuable comments.

I Introduction

Climate change poses a serious 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 the 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 U.S. economy as a case study to analyze the uneven impacts of climate policies across the population of the world’s largest developed economy. Specifically, we provide an in-depth investigation of how carbon taxation affects households across different cohorts and income groups.² We begin our analysis by confirming a key finding of the prior literature that the carbon footprint associated with a dollar spent is higher for poorer households.³ We also document that the emissions content of a dollar spent does not vary substantially by household age.⁴ Furthermore, we find that these emission patterns are mainly driven by energy expenditure shares. That is, poorer households spend disproportionately more on energy (predominantly gasoline, electricity and natural gas) while there is no significant age gradient in household energy spending shares. 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, are felt by households in line with their energy spending behavior.

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 change to 1.5°C.

²A widespread consensus exists that carbon taxation or pricing, together with green subsidies, should be at the center stage of policies that mitigate climate change. See Stern, Stiglitz and Taylor (2022) and Blanchard, Gollier and Tirole (2023).

³Among others, see Levinson and O’Brien (2019) and Sager (2019).

⁴Throughout the analysis, household age refers to the age of the household head.

At the same time, climate policies 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 other key aggregate and distributional moments of the U.S. 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 U.S. in its latest nationally determined contribution (NDC) to the UNFCCC.⁵ We use the model to back out the carbon tax path that delivers this emissions pathway target, taking existing estimates of the time-varying costs of abatement technology as given. We find that the implied carbon tax path that implements the U.S. emissions trajectory to net-zero by 2050, the target net-zero date according to its NDC, starts from \$1 per tonne of CO₂ emissions in 2020, the initial reform year, and increases rapidly to about \$108 by 2050, after which it gradually falls to about \$87 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 about 30% (relative to the 2019 price) by 2050. Interest rates fall until 2050, with wages declining until somewhat later, after which they both recover gradually. We assume that the government uses carbon tax revenues to reduce labor income taxes which implies that the average labor tax also falls until 2050 and recovers afterwards. Following a short-lived, small increase on impact, mostly due to a positive labor supply response, aggregate real output decreases by 0.5% until 2050 after which it slowly returns to its pre-reform long-run level.

⁵For each country, a pledge to the UNFCCC, known as a nationally determined contribution, specifies an emissions-reduction pathway and associated policy commitments. The emissions pathways set out in countries' NDCs form part of the broader Paris Agreement framework for limiting global warming. We use the emissions pathway specified in the 2024 U.S. NDC as a policy benchmark for our analysis.

The reform lowers average household welfare. The welfare cost follows a U-shaped pattern across cohorts: households in their fifties at the time of the reform bear the largest losses, while retirees, particularly those over 70, are affected the least. To understand this pattern, it is useful to consider how the direct and indirect effects of the reform vary across cohorts. The direct welfare cost declines monotonically with age at the time of the reform. This reflects the time path of carbon taxes under the net-zero reform: carbon taxes are initially low but rise rapidly over time. Older cohorts, having fewer remaining years, therefore have substantially lower lifetime exposure to carbon taxes. The indirect channel combines gains from lower labor taxes with losses from lower returns to investment. It benefits younger households, whose income relies more heavily on labor earnings, and hurts those approaching retirement, for whom asset income is relatively more important. Taken together, these forces imply that households in their fifties, who are also the wealthiest age group, bear the largest welfare losses: they face relatively high lifetime exposure to carbon taxes while also being particularly adversely affected by the decline in investment returns, given their relatively high reliance on asset income. Retirees are affected the least because of their lower lifetime exposure to carbon taxes and their general equilibrium effects.

The reform affects households relatively uniformly across the income distribution. Initially, this may appear counterintuitive, given that the share of energy expenditure in total spending declines with income, suggesting that poorer households should be disproportionately affected by higher energy prices. This mechanism is present in our framework; however, this static, short-run logic does not capture the full picture. Once the dynamic, lifetime effects of the reform are taken into account, retirees experience substantially smaller welfare losses. Because retirees are disproportionately represented in the lower income quintiles, they reduce the average direct burden of the reform on low-income households. In addition, the indirect welfare effects of the reform are more adverse for higher-income households. They benefit less from the reduction in labor income taxes and suffer larger losses from lower returns to investment, as capital income constitutes a larger share of their total income. Together these effects generate the relatively uniform

distribution of the welfare cost of the net-zero reform across the income distribution. This finding contrasts with the common characterization of carbon taxation as regressive in both the academic literature and policy debates (see, for example, Blanchard, Gollier and Tirole (2023)), a characterization that largely reflects a focus on the short-run incidence of higher energy prices rather than the full dynamic effects of the reform.

Motivated by concerns about the regressivity of climate reforms, some authors argue that such reforms should be accompanied by increased redistribution toward poorer households. Our findings, however, provide little rationale for redistribution based solely on income or age.⁶ This does not imply that redistributive policy has no role. Transfers conditioned jointly on age and income could instead be used to target households that bear the largest welfare losses from the reform.

A broader insight that emerges from our analysis is that the distributional consequences of climate reforms depend crucially on the time path of carbon taxation. To illustrate this point further, we consider an alternative reform that limits global warming by 2100 to the same extent as our baseline reform but implements a constant carbon tax. Under this policy, the distributional effects closely resemble the short-run incidence of the baseline reform and are distinctly regressive. Thus, policies that achieve similar climate outcomes can have markedly different distributional consequences depending on the timing of carbon taxation. This raises the question of whether the carbon tax path implied by the U.S. commitment under the Paris Agreement has any particular economic rationale. Perhaps surprisingly, we find that the time path of the optimal carbon tax reform that maximizes a weighted-average of household welfare closely resembles that required to implement the U.S. Paris Agreement emissions target, thereby yielding near-identical distributional impacts across cohorts and income groups.

Our choice of an overlapping-generations framework with no altruism toward future generations is consequential for our findings. At the opposite extreme, under perfect

⁶One might argue that these findings depend on the particular way in which carbon tax revenues are recycled in our analysis. More generally, the distributional implications of a carbon tax reform depend on how the government uses the revenues it generates. A common alternative is to redistribute these revenues to households through lump-sum rebates. Such a policy would, if anything, strengthen our finding that the net-zero reform is not regressive.

altruism—the assumption underlying infinitely-lived dynastic models—the time path of carbon taxes would matter less for the distribution of the burden of a climate reform. This is because all households alive at the time of the reform would face the same “lifetime” exposure to carbon taxes independent of their age, since they would effectively be exposed to the same future path of carbon taxes through their descendants. We believe that the overlapping-generations framework with non-altruistic agents is at least as well suited to the study of climate reforms as models with infinitely-lived agents. In particular, the assumption of limited altruism toward future generations appears consistent with the observed degree of inaction on climate change despite its potentially severe consequences beyond the lifetimes of currently living generations. Our finding that the time path of carbon taxes matters for distributional outcomes would continue to hold qualitatively in the, arguably most realistic, intermediate case of positive but imperfect altruism toward future generations.

Literature review Seminal contributions to the literature on the economics of climate change include Nordhaus (1994), Stern (2008) 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 overlapping generations models with a representative agent in each generation to explore the differential welfare impacts of climate policies across generations.

More recently, a new generation of papers that analyze climate policies in the presence of rich heterogeneity has emerged. Douenne et al. (2025), Belfiori, Carroll and Hur (2026) and Douenne, Hummel and Pedroni (2026) calculate optimal carbon taxes under income heterogeneity but using infinitely-lived agent models and without taking into account life-cycle considerations. In this sense, more closely related to the current paper are Fried, Novan and Peterman (2018) and Fried, Novan and Peterman (2024) who also analyze the consequences of carbon taxation using rich, quantitative overlapping generations models

with idiosyncratic productivity shocks and incomplete markets.⁷ These studies primarily examine alternative uses of carbon-tax revenue under reforms with constant tax rates. By contrast, we study both an actual net-zero proposal and an optimal reform, each featuring a time-varying carbon-tax path. We show that this time variation is critical for the distributional effects of the reform.⁸

Methodologically, we differ from all the aforementioned papers by employing a non-homothetic CES household demand system introduced by Comin, Lashkari and Mestieri (2021), rather than the more standard Stone-Geary formulation, which we estimate directly using household-level data on expenditures. An advantage of these preferences over Stone-Geary preferences is that they allow for a meaningful degree of non-linearity in Engel curves even at higher income/wealth levels, which is in line with observed energy shares. In this sense, our paper also relates to a growing literature that embeds non-homothetic CES preferences into macroeconomic models to address distributional questions.⁹ We contribute to this literature by incorporating these preferences into a heterogeneous-agent, overlapping generations framework.

On the empirical side, we build on the literature that combines input-output accounts with household expenditure surveys to measure the pollution content of consumption, see e.g., Levinson and O’Brien (2019) and Sager (2019).¹⁰ We confirm the finding in this literature that the carbon intensity of consumption baskets declines with income. Our substantive contribution to this literature is twofold. First, we document how the carbon intensity of expenditure varies over the life cycle, a dimension that has received less attention than income, and find the profile to be mildly hump-shaped and considerably flatter than the income gradient. Second, and more importantly for what follows, we show that essentially all of the variation in carbon intensity across both income and age

⁷See also Kuhn and Schlattmann (2024) and Dinlemez (2026) who use life-cycle models to analyze the equity-efficiency tradeoff that arises from subsidizing the adoption of carbon-neutral consumer durables.

⁸Benmir and Roman (2022) also evaluate the net-zero reform according to the Paris Agreement; however, unlike us, they do not model household energy consumption or the life cycle.

⁹See, in particular, Oni (2024), Labrousse and Perdureau (2026), and Olivi, Sterk and Xhani (2026) for applications that relate to energy price shocks and carbon taxation.

¹⁰These papers provide descriptive associations between household income and emissions intensity. For causal evidence, see, e.g., Kanzig (2023).

operates through the share of energy in total spending, rather than through differences in the composition of non-energy consumption.

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, the details of the calibration of the model, and a further decomposition of welfare gains.

II Empirical Findings

This section describes the data, presents the empirical findings, and explains how these findings inform our modeling choices regarding the household demand system in the quantitative model presented in Section III.

Data construction We follow Levinson and O’Brien (2019) and Sager (2019) among others in constructing consumption-based accounts of household 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, including the emissions produced at every stage of the production process. Assigning emissions to households in this way is generally considered to account for 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 3.10.1 for the years 1999–2019. EXIOBASE provides a multi-region environmentally extended input-output table, which means that it accounts for imported inputs in production processes which may have different carbon intensities from domestic inputs. We use the product-by-product tables, which distinguish 200 products, and invert the global technology matrix to obtain total CO₂ released per unit of final demand for each product.

Linking these intensities to household spending requires a concordance between EXIOBASE products and the roughly 850 expenditure codes (UCC) in the CEX. Rather than adopt an existing crosswalk, we construct one in three steps. We first map EXIOBASE

products to BEA detail commodity codes using the EXIOBASE-provided concordance to USEEIO tables. We then use the BEA Personal Consumption Expenditure (PCE) bridge table to map BEA commodities to PCE consumption categories. Finally, we map PCE categories to UCCs using the BLS CE–PCE concordance, which we extend by hand for the UCCs it does not cover. In the PCE bridge step, we convert the EXIOBASE coefficients, which are in terms of basic prices, into their purchaser price equivalent coefficients, accounting for retail and wholesale trade, as well as transport emissions. We also net taxes on products out of the intensities, since tax dollars generate no emissions. Appendix A gives the details.

Following Sager (2019), we augment these supply-chain estimates with Energy Information Administration (EIA) data on the direct combustion emissions of gasoline, natural gas and other heating fuels, which input-output accounts do not capture. For electricity, we instead replace the EXIOBASE intensity with the EIA estimate, as emissions from electricity are classified as Scope 2, whose generation emissions EXIOBASE already measures, and the EIA figure is more reliable for the United States specifically.

We combine these intensities with the Interview Files of the Consumer Expenditure Survey (CEX) over the same period. Households are interviewed quarterly about their expenditures, and report their income in the first and final waves. We aggregate these files to an annual frequency and apply the sample selection of Aguiar and Bils (2015): we retain households observed in all four interviews and reporting complete income, restrict attention to urban households, require food spending to be positive and below total expenditure, drop households reporting negative after-tax income or implausibly large spending in any single category, and trim the top and bottom five percent of the within-year distribution of real before-tax income. We depart from Aguiar and Bils (2015) in one respect: because our interest is in behavior over the life cycle, we retain households with heads aged 25 to 84 rather than 25 to 64. We also drop reference year 2015, for which the CEX sample redesign breaks the longitudinal link across interviews. These criteria leave a final sample of 58,979 households across the remaining 20 years.

Expenditures are under-reported in the CEX relative to the national accounts, and differentially so across spending categories. We therefore rescale each component of household spending to its PCE counterpart using the BLS-published CE-to-PCE ratios by category and year. Because these ratios differ across categories, the correction is not a uniform scaling of household budgets: it alters the composition of measured spending, and hence energy shares and emission intensities, by an amount that depends on a household’s basket.

As in previous work employing consumption-based estimates, we are unable to account for differences in emission intensity within products, for instance between more and less environmentally friendly versions of the same good, as all goods are assigned the average emission intensity in their product category. This also assumes that doubling expenditure on a good will double the emissions generated, regardless of whether the price or the quantity has doubled. Fortunately, this is unlikely to be a major issue for our main products of interest, namely gasoline, electricity and natural gas, as these goods are relatively homogeneous in price, conditional on geographical fixed effects. 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.

Estimation By combining estimated carbon intensities with consumption data, we are able to construct total annual CO₂ emissions per household. We then define the carbon intensity of consumption as 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 bins, state of residence and year, respectively, and $\mathbf{X}$ is a set of household-size dummies. 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 households in higher income deciles have lower carbon intensity of expenditure. The differences are quantitatively substantial, with households in the bottom income decile exhibiting 23% higher carbon intensity of consumption than those in the top decile. The emission intensity of expenditure has a hump shape across the life cycle, although the differences across age are much less pronounced than across income: the carbon intensity of expenditure for the highest intensity age group (55–64) is only about 8% greater than for the lowest intensity group (25–34).




Figure 1: CO₂/\$ of expenditure (2019\$)

Because within-product variation in emission intensity is not captured by our consumption-based accounting, 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 analyze these emissions patterns. We find that a simple division into energy and non-energy expenditure goes surprisingly far. 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 shows that most of the variation in emission intensity is driven by energy emissions, which also account for a substantial majority of total household emissions. There are two possible explanations for the patterns in energy emissions per dollar we see in Figure 2. One possibility is that the composition of energy expenditures varies




Figure 2: CO₂ intensity of energy and non-energy expenditure

across households; another is that energy expenditure as a share of total expenditure varies across households. Figure 3 shows that the latter effect dominates, as households' energy emissions per dollar almost exactly track their energy expenditure shares.




Figure 3: Energy-related CO₂ intensity and the energy expenditure share

These figures demonstrate the stylized facts that motivate our structural model. We find that the carbon intensity of expenditure falls with income whereas it exhibits a mild hump shape over the life cycle. These patterns are driven largely by differences in energy consumption, specifically the energy share of overall expenditure.¹¹ We also find that while more complex patterns are present in non-energy expenditure, their quantitative significance is limited (see Appendix B.1). We therefore omit this channel from the model that follows.

¹¹Regressing CO₂ per dollar of expenditure on energy expenditure share plus year fixed effects produces a within- R^2 of 0.85.

The next section presents a model that rationalizes the income gradient in energy expenditure. Following the literature, we capture this gradient through non-homothetic preferences over energy and non-energy consumption. Given that the hump shape in energy spending shares across the life cycle is muted, we do not complicate the model further by tailoring it to match this pattern.¹²

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 25 to 84. Life prior to labor market entry is not modeled. The entrants draw their initial assets from a distinct distribution that will be calibrated to match the empirical distribution of assets among the young. Between periods 1 and 40, households make consumption, saving and labor supply decisions. In period 41, 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 used to fund the initial asset distribution, with any excess received as revenues by the government, following Krueger and Wu (2025).

Preferences Households consume two distinct goods: e represents the energy good and c represents a composite non-energy good. Without loss of generality, the non-energy consumption good is assigned to be the numeraire good, with a price equal to 1 in

¹²In an earlier version of the paper, we targeted the observed energy pattern over the life cycle by allowing the parameters that control the non-homotheticity of preferences over energy and non-energy goods to be age-dependent. We find that such modeling does not affect our main quantitative findings despite complicating the model substantially.

each period. Preferences over the consumption of non-energy and energy goods and labor in a period are defined according to a utility function $u(c, e, l)$. People discount utility across periods by $\beta \in (0, 1)$.

Production There are two sectors. The representative firm in sector $s = 1$ uses capital, labor and the sector 1 intermediate good to produce the non-energy good. In any period t , taking the rental rate, the wage rate, and the price of the sector 1 intermediate good, r_t , w_t and $p_{1,t}^x$, as given, this firm solves:

$$(1) \quad \max_{Y_t, K_{1,t}, L_{1,t}, X_{1,t}} Y_t - r_t K_{1,t} - w_t L_{1,t} - p_{1,t}^x X_{1,t} \quad \text{s.t.} \quad Y_t \leq F_{1,t}(K_{1,t}, L_{1,t}, X_{1,t}),$$

where Y_t is the gross output in the non-energy sector and $(K_{1,t}, L_{1,t}, X_{1,t})$ are the amounts of capital, labor and intermediate goods used in its production.

The representative firm in sector $s = 2$ uses capital, labor and the sector 2 intermediate good to produce the energy good. In any period t , taking all prices as given, this firm solves:

$$(2) \quad \max_{E_t, K_{2,t}, L_{2,t}, X_{2,t}, \mu_t} (p_t - \tau_{e,t}(1 - \mu_t))E_t - r_t K_{2,t} - w_t L_{2,t} - p_{2,t}^x X_{2,t} - \Theta_t(\mu_t E_t) \\ \text{s.t.} \quad E_t \leq F_{2,t}(K_{2,t}, L_{2,t}, X_{2,t}),$$

where E_t is the gross output in the energy sector, μ_t is the share of abatement in energy production, $\tau_{e,t}$ is an excise carbon tax that applies to unabated energy production, and $\Theta_t(\mu_t E_t)$ is the physical cost of abating $\mu_t E_t$ amount of energy production in period t . Capital depreciates at rate δ in both sectors.

In each sector, there is an intermediate-input retailer that aggregates sectoral intermediate inputs $(Y_{s,t}, E_{s,t})$ into an intermediate-input bundle $X_{s,t}$ using a constant returns to scale technology and sells this to the final good firm in that sector at price $p_{s,t}^x$. The problem of this firm is:

$$(3) \quad \max_{X_{s,t}^f, Y_{s,t}, E_{s,t}} p_{s,t}^x X_{s,t}^f - Y_{s,t} - p_t E_{s,t} \quad \text{s.t.} \quad X_{s,t}^f \leq G_{s,t}(Y_{s,t}, E_{s,t}),$$

where $X_{s,t}^f$ is the intermediate good firm's supply in sector s 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 25) 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 addition to the carbon tax on unabated energy production, $\tau_{e,t}$, the government collects revenue via non-linear labor income taxes, $T_t(y)$, where y is labor income, a linear consumption tax $\tau_{c,t}$ and a linear tax τ_t applied to capital income net of depreciation. 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 $\{G_t\}_{t=0}^{\infty}$ and the pension system $\{P_t(\cdot)\}_{t=0}^{\infty}$, and to repay government debt $\{D_t\}_{t=0}^{\infty}$. Pensions are not subject to the labor tax, following Kindermann and Krueger (2022). The government also distributes accidental bequests from deceased agents, B_t , to newborn agents in the subsequent period, after accruing interest. The details of the distribution of these initial assets are discussed in Appendix C.2. Finally, the government collects any profits made by the energy sector, which will be positive when carbon taxes are positive, because the abatement cost function, Θ , is convex.

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 $\mathcal{A} = [\underline{a}, \infty)$ denotes the set of possible asset levels that working-age agents can hold. We force the retirees to have non-negative net worth. Our assumptions imply that, in every period, the total asset holdings of consumers (including the bequest fund) must equal the government debt 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:

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

Retiree's problem Following Krueger and Wu (2025), we assume households receive a pension according to their productivity state at retirement. An individual of age $h > 40$ with retirement period productivity shock and current asset level (z^{40}, a_t) solves:

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

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 of each firm in each sector, $(Y_t, K_{1,t}, L_{1,t}, X_{1,t})_{t=0}^\infty$, $(E_t, K_{2,t}, L_{2,t}, X_{2,t}, \mu_t)_{t=0}^\infty$, and $(X_{s,t}, Y_{s,t}, E_{s,t})_{t=0}^\infty$ for $s = 1, 2$, value and policy functions, $(v_{h,t}(z, a), c_{h,t}(z, a), e_{h,t}(z, a), l_{h,t}(z, a), a_{h+1,t+1}(z, a))_{t=0}^\infty$ for $h = 1, \dots, 60$, a price system $(p_t, w_t, R_t, p_{1,t}^x, p_{2,t}^x)_{t=0}^\infty$, a bequest fund $(B_t)_{t=0}^\infty$, 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, production choices of the firm in sectors $s = 1, 2$ solve the firms' problems given by (1), (2) and (3).
2. Given government policy and the price system, the policy functions solve individual problems given by (4) and (5), 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 $(z', a') \in \mathcal{Z} \times \mathcal{A}$:

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

for working age agents, and

$$\Lambda_{h+1,t+1}(z', a') = \int_{\{a: a_{h+1,t+1}(z, a) \leq a'\}} d\Lambda_{h,t}(z, a)$$

for retirees

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.

4. The bequest fund evolves according to

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

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

$$\begin{aligned} K_{1,t} + K_{2,t} + D_t &= \sum_{h=1}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} a \cdot d\Lambda_{h,t}(z, a) + B_t, \\ L_{1,t} + L_{2,t} &= \sum_{h=1}^{40} \pi_h \kappa_h \int_{\mathcal{Z} \times \mathcal{A}} l_{h,t}(z, a) z d\Lambda_{h,t}(z, a), \\ C_t + I_t + G_t + Y_{1,t} + Y_{2,t} &= Y_t - \Theta_t(\mu_t E_t), \\ E_{1,t} + E_{2,t} + E_{c,t} &= E_t, \\ X_{s,t}^f &= X_{s,t}, \quad \text{for } s = 1, 2, \end{aligned}$$

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

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

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

$$\begin{aligned} G_t + R_t D_t + P_t^{agg} - D_{t+1} + \pi_1 \int_{\mathcal{Z} \times \mathcal{A}} a d\Lambda_{1,t}(z, a) &= \tau_{e,t}(1 - \mu_t) E_t + \tau_{c,t}(C_t + p_t E_{c,t}) + \\ &\quad \tau_t(r_t - \delta) \sum_{s=1}^2 K_{s,t} + T_t^{agg} + R_t B_t + \Pi_t^E, \end{aligned}$$

where

$$\begin{aligned}
T_t^{agg} &= \sum_{h=1}^{40} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} T_t(l_{h,t}(z, a) w_t z) d\Lambda_{h,t}(z, a) \\
P_t^{agg} &= \sum_{h=41}^{60} \pi_h \int_{\mathcal{Z} \times \mathcal{A}} P_t(z^{40}) d\Lambda_{40,t-(h-40)}(z = z^{40}, a) \quad \text{and} \\
\Pi_t^E &= (p_t - \tau_{e,t}(1 - \mu_t))E_t - r_t K_{2,t} - w_t L_{2,t} - p_{2,t}^x X_{2,t} - \Theta_t(\mu_t E_t),
\end{aligned}$$

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

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

IV Parameterization

This section explains how we parameterize the model to align it 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 the estimation of 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. The internal calibration procedure is summarized in Table 2. Whenever data are not available up to 2019 for some variable, we use the most recent data. The details and definitions of the data and further details regarding calibration can be found in Appendix C.

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

We assume that the period utility function is separable between labor and the two other goods and isoelastic in labor; that is, $u(c, e, l) = \tilde{u}(c, e) - \Psi \frac{l^{1+1/\gamma}}{1+1/\gamma}$. Following Comin, Lashkari and Mestieri (2021), we assume that the utility level U that results from consumption bundle (c, e) under $\tilde{u}(\cdot)$, that is, $\tilde{u}(c, e) = U$, is implicitly given by

$$(6) \quad (U^{\varepsilon_c} b_c)^{\frac{1}{\bar{\sigma}}} c^{1-\frac{1}{\bar{\sigma}}} + (U^{\varepsilon_e} b_e)^{\frac{1}{\bar{\sigma}}} e^{1-\frac{1}{\bar{\sigma}}} = 1.$$

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 the Human Mortality Database for 2019. The discount factor β , the labor disutility parameter Ψ and the taste weights (b_c, b_e) are calibrated internally as discussed later in the text. The parameter that controls the elasticity of substitution between energy and non-energy consumption, $\bar{\sigma}$, and the expenditure elasticity parameters, $(\varepsilon_c, \varepsilon_e)$, are estimated as follows.

Estimating preferences for energy We adopt the estimation strategy and identifying assumptions of Comin, Lashkari and Mestieri (2021) to recover our nonhomothetic CES preference parameters. The demand system defined in equation (6) implies the following expression for the relative shares of expenditure for energy and non-energy goods

$$(7) \quad \ln \left(\frac{\omega_{et}}{\omega_{ct}} \right) = (1 - \bar{\sigma}) \left(\ln \left(\frac{p_{et}}{p_{ct}} \right) + \left(\frac{\varepsilon_e}{\varepsilon_c} - 1 \right) \left[\ln \left(\frac{X_t}{p_{ct}} \right) + \frac{\ln \omega_{ct}}{1 - \bar{\sigma}} \right] \right) + \ln \left(\frac{b_e}{b_c^{\varepsilon_e/\varepsilon_c}} \right),$$

where X_t refers to the household's total expenditure, and p_{ct} and ω_{ct} refer to the price index and expenditure share of the non-energy sector, which serves as the base good. Normalising b_c and ε_c to 1, without loss of generality, the empirical counterpart to equation (7) is

$$(8) \quad \ln \left(\frac{\omega_{et}^n}{\omega_{ct}^n} \right) = (1 - \bar{\sigma}) \left(\ln \left(\frac{p_{et}^n}{p_{ct}^n} \right) + (\varepsilon_e - 1) \ln \left(\frac{X_t^n}{p_{ct}^n} \right) \right) + (\varepsilon_e - 1) \ln \omega_{ct}^n + \ln (b_e^n) + \nu_{et}^n$$

Table 1: External Calibration

Parameter	Symbol	Value	Source
<i>Preferences and demographics</i>			
Coefficient of relative risk aversion	σ	2	Literature
Frisch elasticity of labor supply	γ	0.82	BPS
Elasticity of substitution, energy–non-energy	$\bar{\sigma}$	0.21	CEX (1999–2019)
Non-homotheticity of energy demand	ε_e	0.46	CEX (1999–2019)
Age-dependent survival probability	S_h	See Appendix C.1	HMD (2019)
Life-cycle profile of wages	κ_h	See Appendix C.1	HP
Newborn asset endowment distribution	a_1	See Appendix C.2	SCF, ages 23–27
Borrowing limit for working age households	$\underline{a}$	-\$96,214	SCF, ages 23–27
Persistence of productivity shocks	ρ	0.97	HSV
<i>Production</i>			
Capital share of gross output, non-energy	α_1	0.18	Barrage (2024)
Capital share of gross output, energy	α_2	0.17	Barrage (2024)
Energy intermediate share, non-energy	χ_1	0.03	BEA, 2012–2017
Energy intermediate share, energy	χ_2	0.40	BEA, 2012–2017
Non-energy intermediate share, non-energy	ϕ_1	0.40	BEA, 2012–2017
Non-energy intermediate share, energy	ϕ_2	0.18	BEA, 2012–2017
Cost of abatement	$\Theta(\cdot)$	See Appendix C.4	Barrage (2024)
<i>Government</i>			
Labor tax progressivity	τ_l	0.16	GRR
Linear tax on consumption	τ_c	0.05	MRT
Linear tax on capital income	τ	0.36	TU
Government consumption	G/Y	0.15	NIPA, 2010–2019
Government debt	B/Y	0.71	NIPA, 2010–2019
Linear tax on carbon	τ_e	0	Barrage (2019)
Pension function	$P(\cdot)$	See Appendix C.5	KW

This table reports the parameters that we take directly from the literature or estimate using the data. BPS, GRR, HMD, HP, HSV, KW, MRT, NIPA, and TU stand for Blundell, Pistaferri and Saporta-Eksten (2016), Grand, Ragot and Rodrigues (2025), Human Mortality Database (2019), Huggett and Parra (2010), Heathcote, Storesletten and Violante (2010), Krueger and Wu (2025), Mendoza, Razin and Tesar (1994), the National Income and Product Accounts, and Trabandt and Uhlig (2011), respectively.

where n refers to a given household. We follow Comin, Lashkari and Mestieri (2021) in parameterizing the relative taste weights as a linear function of household observables, including age groups, family size dummies and the number of earners in the household. This is intended to soak up household-specific taste shocks. Household-specific prices are constructed as Stone indices; a household n 's index for each good is an expenditure-weighted regional CPI index of the component categories. To account for the endogeneity of prices, we construct Hausman instruments from the prices of other census regions. We also instrument for total expenditure with real after-tax income and income quintile, following Aguiar and Bils (2015), as income is better measured in the CEX and the instrument should, therefore, avoid attenuation bias from measurement error in the regressor. Fi-

nally, we include region fixed effects to soak up region-specific shocks in the error term ν_{et}^n . We estimate the demand system using the generalized method of moments, which delivers estimated values $\bar{\sigma} = 0.21$ and $\varepsilon_e = 0.46$.

Asset market Following Krueger and Wu (2025), the initial distribution of assets of the entrants, a_1 , is set to a discrete version of the empirical net worth distribution of young households estimated using SCF data. Further details of the procedure we follow can be found in Appendix C.2. The lowest point of this grid also serves as the borrowing limit for working-age households, $\underline{a}$.

Production We specify production functions in both sectors in gross output terms in order to be consistent with the interpretation of our demand parameters. The functions that summarize production in sector s are given by

$$F_{s,t}(K_{s,t}, L_{s,t}, G_{s,t}(Y_{s,t}, E_{s,t})) = Z_{s,t} K_{s,t}^{\alpha_s} L_{s,t}^{1-\alpha_s-\chi_s-\phi_s} E_{s,t}^{\chi_s} Y_{s,t}^{\phi_s}.$$

As explained further in Appendix C.3, for each sector, the factor shares for the intermediate goods are calculated from the BEA Supply-Use Detail table. The split of the remaining factor share between capital and labor is taken from Barrage (2024). Total factor productivity terms in the two sectors, Z_1 and Z_2 , are internally calibrated for the initial steady state and assumed constant thereafter. The energy sector firm also optimizes over the share of emissions it chooses to abate, μ , with the abatement cost function $\Theta_t(\mu_t E_t)$. This function is a per-tonne approximation of the U.S. abatement cost function from RICE-2010, as parameterized by Barrage (2024). The details of this abatement function and its calibration are discussed in Appendix C.4.

Government The government spending-to-output ratio is set to 15%, which is roughly in line with the National Income and Product Accounts (NIPA) average for the period 2010–2019. This excludes government investment spending.¹³ In the status quo, the

¹³The corresponding capital-output and investment-output ratios used as targets for internal calibration also exclude government investment and capital stock for consistency.

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 $\tau = 36\%$. We follow Mendoza, Razin and Tesar (1994) and assume that the consumption tax $\tau_c = 5\%$. We take the value of labor tax progressivity, τ_l , equal to 0.16, from Grand, Ragot and Rodrigues (2025). They estimate this parameter for the U.S. from the Luxembourg Income Study, excluding capital taxation but including other transfers in the labor tax system, which is consistent with our fiscal framework. The function that represents the pension system, $P(\cdot)$, follows Krueger and Wu (2025) and is assumed to be a piecewise linear function of mean earnings. For further details, see Appendix C.5. We calibrate λ , which controls the average labor tax in the economy, to clear the government budget.

Productivity shocks Our parameterization of the productivity process closely follows that of Kaymak, Leung and Poschke (2026). Productivity is made up of a fixed component, f , a life-cycle profile, κ_h , and an idiosyncratic component. Following a large literature, notably including Castaneda, Díaz-Giménez and Ríos-Rull (2003), we include a superstar shock, in order to match the high degree of income and wealth concentration at the top which standard log-normal earnings processes fail to deliver. Specifically, productivity, z , can be either in an ordinary or a superstar state. The fixed component is equal to $-f$ or $+f$ for each half of the population, and is assigned upon labor market entry. As in Kaymak, Leung and Poschke (2026), we impose that the fixed component accounts for 60% of the total stationary variance of log productivity across the ordinary states, σ_z^2 , so that $\text{var}(f) = 0.6\sigma_z^2$, and the persistent component carries the remaining $0.4\sigma_z^2$. We set the annual persistence of the persistent component to $\rho = 0.97$, following Heathcote, Storesletten and Violante (2010), and discretize it using the Rouwenhorst method. The scale σ_z^2 is therefore the single free parameter governing ordinary productivity dispersion, and we calibrate it to the variance of log earnings among households aged 50–54 in the PSID. Our life-cycle labor efficiency profile is taken from Huggett and Parra (2010), and is normalized to 1 for labor market entrants.

When in the superstar state, the productivity shock is ι_1 higher than the highest normal state. Productivity transitions to a superstar state at any given time from any

ordinary state with probability ι_2 . The probability of remaining in the superstar state is ι_3 . These superstar state parameters are internally calibrated to match the working-age earnings Gini, the wealth Gini, and the wealth top-1% share. When agents return to an ordinary state, they draw a new labor market ability uniformly from the normal state productivity levels. Labor market entrants draw their initial productivity shock from a distribution which is parameterized by ζ as explained in Appendix C.6. This parameter is calibrated to match the variance of log earnings for 25–29 year-old households, estimated from the PSID. In total, the productivity process introduces six parameters, $(f, \sigma_z^2, \iota_1, \iota_2, \iota_3, \zeta)$, to be internally calibrated. The details of the computations of the data moments needed for the calibration of the distributional moments can be found in Appendix C.7.

Internal calibration In addition to the six parameters of the productivity process, we also calibrate six production and preference parameters to match aggregate moments. Although these twelve parameter values are chosen simultaneously to align the model economy with the U.S. economy, it is instructive to assign parameters to moments. Z_1 , Z_2 and b_e are calibrated to match output-per-household and household energy spending’s share in total output, and to scale the model units of energy such that initial steady-state gross energy output is equal, without loss of generality, to 1.32 (which equals the quantity of gigatonnes of carbon produced in the U.S. in 2019).¹⁴ As a result of this normalization, the price of a unit of energy in the resulting equilibrium, p , is equal to \$16,000. β and Ψ are calibrated to match the capital-output ratio and average hours. We calibrate the depreciation rate δ to hit the investment-output ratio. Aggregate calibration targets are detailed in Appendix C.8. Table 2 reports calibrated parameter values and the model’s ability to match data targets.

¹⁴The unit in which we count energy is not disciplined by the model. We choose the unit of energy good so that one unit of (unabated) energy production generates 1 gigatonne of carbon. See Appendix C.9 for further details. A different choice does not affect any of our results.

Table 2: Internal Calibration

Panel A: Moments	Data	Model	Source
<i>Distributional moments</i>			
Earnings Gini, working age	0.51	0.53	SCF, 2010–2019
Var. log earnings, 25–29	0.52	0.52	PSID, 2010–2018
Var. log earnings, 50–54	0.76	0.75	PSID, 2010–2018
Wealth Gini	0.85	0.83	SCF, 2010–2019
Wealth top-1% share	0.36	0.36	SCF, 2010–2019
Fixed effect share of variance	60%	60%	KLP
<i>Aggregate moments</i>			
GDP per household	167500	167469	NIPA, 2010–2019
Capital–output ratio	2.37	2.35	NIPA, 2010–2019
Investment–output ratio	0.17	0.17	NIPA, 2010–2019
Household energy spending / GDP	0.03	0.03	BEA, 2012–2017
Average hours worked	0.35	0.35	Literature
Gross energy output	1.32	1.32	Normalization
Panel B: Parameters	Symbol	Value	
Superstar productivity gap	ι_1	3.44	
Prob. of entering superstar state	ι_2	7.23×10^{-4}	
Prob. of remaining in superstar state	ι_3	0.83	
Initial productivity dispersion	ζ	0.15	
Variance of ordinary state productivity shocks	σ_z^2	0.89	
Fixed component of productivity	f	0.73	
<i>Aggregate parameters</i>			
TFP, non-energy sector	Z_1	288.04	
TFP, energy sector	Z_2	0.16	
Discount factor	β	0.97	
Depreciation rate of capital	δ	0.08	
Taste weight on energy	b_e	0.04	
Labor disutility weight	Ψ	1.11×10^{-3}	

This table reports the parameters calibrated internally and the corresponding targeted moments and the model’s ability to match them. Except for worker fixed effects, data moments are authors’ own calculations. BEA, KLP, NIPA, PSID, and SCF stand for Bureau of Economic Analysis, Kaymak, Leung and Poschke (2026), the U.S. Bureau of Economic Analysis (2021), Panel Study of Income Dynamics and Survey of Consumer Finances, respectively.

IV.1 Model Fit

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

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 while Panel B and

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

	1 st	2 nd	3 rd	4 th	5 th
A: Earnings quintile shares					
Data	3.5%	8.2%	13.2%	19.8%	55.2%
Model	4.2%	7.5%	12.1%	19.7%	56.5%
B: Income quintile shares					
Data	0.0%	4.2%	10.3%	18.6%	66.9%
Model	1.7%	5.6%	10.1%	19.0%	63.6%
C: Wealth quintile shares					
Data	-0.5%	0.8%	3.4%	9.4%	87.0%
Model	-2.8%	0.5%	5.7%	14.9%	81.7%

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 are derived by averaging the corresponding moments of the 2010, 2013, 2016 and 2019 waves of the SCF.

Table 4: Non-Targeted Moments: Joint Distribution of Income and Wealth

	1 st	2 nd	3 rd	4 th	5 th
A: Income shares by wealth quintile					
Data	6.8%	9.3%	12.6%	15.5%	55.9%
Model	13.0%	11.2%	12.0%	16.4%	47.4%
B: Wealth shares by income quintile					
Data	5.2%	6.4%	6.5%	10.1%	71.9%
Model	5.9%	8.3%	8.5%	9.7%	67.6%

This table reports the fit of the model with respect to non-targeted moments of the joint distribution of income and wealth. Income refers to capital, labor and pension income combined. All data moments are derived by averaging the corresponding moments of the 2010, 2013, 2016 and 2019 waves of the SCF.

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. Following the literature and policy debates, our primary measure of inequality and regressivity is (after-tax) income. As wealth is a more welfare-relevant metric, it is important that the model also reproduces the joint empirical distribution of income and wealth. Table 4 shows that this objective is largely achieved, despite not being directly targeted in calibration.

Expenditure moments As Figure 4a displays, the model does well in replicating the overall empirical profile of consumption expenditure 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 counterfactually steep consumption growth in early life. As our model includes no notion of family size, life-cycle consumption is converted to household equivalized units in the data, hence the much more muted hump shape than is visible in unconditional empirical consumption profiles.

A main channel through which the burden of carbon taxes falls heterogeneously across households is the expenditure channel. For this reason, it is crucial that the model captures well the share of energy spending in total household expenditure across income and age. Figure 4b shows that the model captures qualitatively the declining pattern of energy share by income although it overshoots for the lowest income quintile and somewhat undershoots for upper income quintiles. As shown in Appendix B.2, the model provides a near-perfect fit for energy share across expenditure quintiles. The excess steepness in the income-energy path is therefore a consequence of the pass-through of income into total expenditure being somewhat too high, a known feature of models that do not allow for large transitory income shocks. Figure 4c shows that the life-cycle profile of the energy spending share is close to flat, a pattern that is replicated by the model.¹⁵




Figure 4: Non-targeted moments: life-cycle consumption and energy expenditure shares

These figures report the fit of the model with respect to some non-targeted moments of expenditure. All data moments are derived using CEX data from 2010 to 2019.

Earnings, Income and Wealth over the Life Cycle The life-cycle path of wealth, income and earnings, shown in Figure 5, is also matched very closely, with the exception

¹⁵In the calculation of the data moments for both Figures 4b and 4c, we use the BLS provided CEX to Personal Consumption Expenditure conversion factors to scale up total household expenditure uniformly to account for the substantial share of consumption that is not picked up by the CEX, which only counts out-of-pocket expenditures plus imputed housing. We do so because our calibration also targets energy's share of total value added in data and total value added includes these items not captured by CEX.

of the excessively rapid decumulation of wealth in later life, and the corresponding over-accumulation in middle age, which is an artifact of our model lacking features such as bequest motives and health shocks, which induce higher savings among the retired.




Figure 5: Non-targeted moments: life-cycle wealth, income and earnings

These figures report the fit of the model with respect to earnings, income and wealth over the life cycle. All data moments are derived by averaging the corresponding moments of the 2010, 2013, 2016 and 2019 waves of the SCF.

Aggregate emissions This section evaluates the model’s fit to the 2019 U.S. economy in terms of matching aggregate moments related to emissions and energy use. Total emissions are targeted in our calibration. The remaining moments are untargeted. Table 5 shows that our model reproduces the relative importance of household emissions in total emissions (as approximated from EPA aggregates) very closely. We also match the share of household emissions that is attributable to energy expenditure and the average energy share of expenditure by households.

Table 5: Aggregate Emissions in 2019

Moment	Data	Model
Total emissions	1.32 GtC	1.32 GtC
Share of direct household emissions in total	41.5%	43.6%
Share of energy emissions in household emissions	65.5%	66.4%
Average energy expenditure share	6.0%	6.2%

This table reports model fit regarding some key aggregate emissions and energy share moments. All data moments correspond to carbon emissions in 2019 U.S. economy and are taken or approximated from U.S. Environmental Protection Agency (2024) or our combined EXIOBASE-CEX dataset. Energy expenditure shares are, as in the model fit figures, scaled down to account for components of PCE that are not picked up by the CEX.

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 do 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 U.S. administration in its latest NDC. We then consider alternative reforms.

V.1 Baseline Reform




Figure 6: (a) Annual emissions targets that are prescribed in the U.S. 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 U.S. in its 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 the 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 rate results in larger abatement, and therefore, a larger 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 since, following Nordhaus (2008), we assume that abatement is costless from this point onward.





Figure 7: Dynamics of the price of energy, the wage rate, the net return to capital and average labor income tax rate 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 perhaps notable is the magnitude of the response: the model predicts that the energy price will rise by 30% by 2050 in response to the 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 the equilibrium level of energy, raising its price and lowering its use in the non-energy sector, and lowering the demand for capital and labor in that sector as well. As a result, we observe a decline in wage and rental rates until 2050 after which they converge back to their initial levels in the long run. Figure 7d shows that the average labor earnings tax rate, defined as the ratio of total earnings tax revenue to total earnings, follows the opposite pattern to 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. 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 (Figure 8a). Figures 8b and 8c show that equilibrium levels of labor and capital both increase until 2050 after which they fall back. Interpreted together with the trajectory of the wage and interest rates in response to the reform, it is evident that both aggregate labor supply and savings curves shift out to the right, meaning households are willing to supply more labor and save more for a given level of wage and interest rate, until 2050 after which the curves shift back. The labor supply curve shifts out because (i) average labor tax rate falls with the reform until 2050 (substitution effect) and (ii) the rise in energy price induces people to work harder due to an income effect. The savings curve shifts out because (i) rising energy prices require households to save for higher costs of living in the future and (ii) rising after-tax labor income over time induces higher precautionary savings. The rise in labor and capital dominates the fall in energy as an input to production on impact and as a result real output (net of abatement costs) initially increases. However, as Figure 8d shows, this is quickly reversed, in part due to the rising cost of abatement (Figure 8e), and output remains below its long-run level until after 2250. Figure 8f displays that 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




(a) Energy production




(b) Labor




(c) Capital stock



(d) Real GDP



(e) Abatement spending as share of GDP




(f) Energy: intermediate vs. household use

Figure 8: Dynamics of aggregate quantities following the reform

assumptions we make regarding preferences and non-energy good production function, consumers have less room to substitute away from energy use relative to producers.

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 circles



(a) Welfare cost by age at the time of reform (b) Welfare cost by income at the time of reform

Figure 9: Welfare costs of hitting the U.S. NDC emissions reduction path

display the discounted lifetime welfare consequences of the net-zero reform in expenditure equivalent units. We measure welfare gains of the reform for an individual household by calculating the uniform expenditure change across dates and states that is needed to make the household indifferent between pre- and post-reform economies. We then aggregate the welfare gains for a cohort or income quintile by averaging the welfare gains of households in that group. 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 orange triangles isolate the direct welfare effect of the reform by comparing household welfare in the pre-reform economy to that in an auxiliary economy in which households face post-reform energy prices but wages, labor income taxes, interest rates and pensions are kept at the pre-reform levels. The green squares 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 has a U-shaped pattern across cohorts, with those in their fifties 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 combination of the direct and indirect effects gives rise to the U-shaped pattern. The direct welfare cost, displayed by orange triangles in Figure 9a, decreases monotonically

with age at the time of the reform. This may come as a surprise since, as depicted by Figure 4c, the share of energy expenditure in total spending is fairly flat across age groups, which seems to suggest that households of different ages should suffer uniformly from rising energy prices. This is, in fact, true in the short run. A different picture emerges, however, when we take into account lifetime effects. As depicted by Figure 6b, carbon taxes implied by the net-zero reform start low but rise rapidly with time, implying that average lifetime carbon tax exposure is lower for older cohorts. Regarding the indirect channel, the reform implies welfare gains through lower labor taxes and losses via lower returns to investment. The green squares in Figure 9a show that this channel benefits younger households more as labor income constitutes a larger share of their total income, and hurts those approaching retirement more, for whom capital income is relatively more important. Those who are already retired suffer smaller losses despite the high share of capital income in their total income, because their lower lifetime exposure to carbon taxes mitigates the adverse indirect effects of the reform.¹⁶

Along the income distribution The blue circles in Figure 9b show that the reform hurts households fairly uniformly across income. This may seem counterintuitive given the fact that, as depicted by Figure 4b, energy expenditure share declines with income, which implies that the poorest should be hurt more when energy prices increase. This mechanism exists in our framework; however, this static, short-term logic fails to capture the complete picture. Recall that, once we include the reform’s dynamic impacts, the retired experience substantially smaller losses. Given that this demographic group is overrepresented in low income quintiles, it lowers the average burden of the reform that falls on low income groups. As depicted by the orange triangles in Figure 9b, this implies a direct effect that is mostly flat along the income distribution. Furthermore, as the green squares in Figure 9b show, the indirect welfare effects of the reform favor higher income groups less 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

¹⁶A further breakdown of the indirect effects into those working through changes in after-tax wages, interest rates and pensions is relegated to Appendix D.

income). A combination of these effects implies the rather flat impact of the net-zero reform along the income distribution.

An alternative way of recycling carbon tax revenues would be to redistribute them to households through lump-sum rebates. Although we do not explicitly consider lump-sum rebates in this paper, incorporating them would, if anything, reinforce our finding that carbon tax reforms need not be regressive. Put differently, our findings indicate that carbon tax reforms do not have to be accompanied by progressive redistributive policies to prevent regressive outcomes.

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.¹⁷ Figure 10a shows that, relative to the baseline reform, the carbon tax starts much higher at just below \$80 per tonne of emissions but does not go as high in later years. The front-loading of carbon taxes in the flat reform is also reflected in the energy price, shown in Figure 10b, which is higher earlier and lower later on relative to the baseline reform.

With carbon taxes held constant over time, the direct welfare implications of the reform are quite uniform across cohorts as displayed by Figure 10c. This is in line with rather uniform energy spending shares across the life cycle depicted in Figure 4c. The indirect effect of the reform across cohorts is quite similar to the case of the baseline reform except that the cohort which is hurt the most is a bit older at the time of the reform and in general older cohorts are hurt more. This reflects the fact that the old, who have relatively higher capital income, are now also facing high carbon taxes and their general equilibrium implications. Overall, as the blue circles display, the welfare cost of a constant carbon tax reform rises with age until about retirement after which it stabilizes, reflecting the shape of the general equilibrium effects. Interestingly, Figure 10d

¹⁷Our model already implies a mapping from economic activity to emissions. We use the climate model in Barrage (2024) to map the model implied emissions to temperature changes. This gives us a mapping from carbon tax pathways to temperature changes. For details, see Appendix C.10.




Figure 10: Welfare costs of a constant carbon tax

shows that carbon taxes are regressive under the constant tax reform. This again reflects the spending patterns displayed in Figure 4b: 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. The fact that the distributional consequences of the reform follow energy spending shares is intuitive: when carbon taxes are constant over time, the short-term and the lifetime impacts of the reform coincide. Finally, 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.

Importance of the path of carbon taxes The constant tax reform highlights the broader message that the distributional consequences of a transitional climate policy depend critically on the evolution of carbon taxes over time, rather than only on the

overall magnitude of the intervention. This naturally raises the question of whether there is an economic rationale for the carbon tax path implied by the U.S. commitment under the Paris Agreement. This question partly motivates our analysis of the optimal carbon tax policy in the next section.

Optimal reform In this section, we consider the problem of a government that chooses carbon tax paths to maximize a social welfare function while keeping the temperature increase by 2100 at the level implied by the baseline reform. The welfare aggregator we consider is

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

This aggregator assumes that society weighs the welfare of each cohort alive at the time of the reform equally and that of each future generation with a geometric discount factor of $\hat{\beta} \in (0, 1)$. 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 enjoying that consumption. We restrict ourselves to linear carbon tax paths for computational tractability.

Figures 11a and 11b reveal that optimal carbon taxes and the implied change in energy prices are quite similar to those in the baseline reform albeit slightly more front-loaded. As a consequence, as in the baseline case, the welfare losses have a U-shape across cohorts and are uniform across the income distribution, as depicted by Figures 11c and 11d.

VI Conclusion

This paper studies the distributional consequences of net-zero transitions. We document that the carbon intensity of consumption declines with income, reflecting systematic differences in energy expenditure shares. We build a heterogeneous-agent overlapping-generations model with non-homothetic preferences that matches this pattern as well as




Figure 11: Welfare costs of the optimal reform

key aggregate, distributional and life-cycle moments and use it to evaluate the impact of the carbon tax path consistent with the U.S. net-zero emissions pathway targeted under the Paris Agreement.

We find that even though carbon tax incidence falls disproportionately on poorer households in the short run, welfare costs of the reform aggregated over a lifetime are uniform across the income distribution. Across cohorts, the lifetime costs of the reform follow a U-shaped profile, with middle-aged cohorts at the time of the reform being affected the most. We highlight a lesson that may be useful for policymakers considering climate reforms: the specific timing of carbon taxes is an important determinant of the distribution of the burden of climate reforms across households when one takes into account the whole impact over a lifetime and not just the short-run incidence. Finally, an optimal reform exercise shows that the optimal time path of carbon taxes, and hence, the distributional implications of an optimal reform coincide almost exactly with the U.S. policy under the Paris Agreement.

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.** 2026. “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.”
- Blanchard, Olivier, Christian Gollier, and Jean Tirole.** 2023. “The Portfolio of Economic Policies Needed to Fight Climate Change.” *Annual Review of Economics*, 15(1): 689–722.
- 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.
- Comin, Diego, Danial Lashkari, and Martí Mestieri.** 2021. “Structural Change With Long-Run Income and Price Effects.” *Econometrica*, 89(1): 311–374.
- Dinlemez, Sevim.** 2026. “The Incidence of Carbon Taxes in an Electrifying Economy.” Working paper.
- Douenne, Thomas, Albert Jan Hummel, and Marcelo Pedroni.** 2026. “Optimal Fiscal Policy in a Climate-Economy Model with Heterogeneous Households*.” *The Economic Journal*, ueag006.
- Douenne, Thomas, Sebastian Dyrda, Albert Jan Hummel, and Marcelo Pedroni.** 2025. “Optimal Climate Policy with Incomplete Markets.” Working paper.
- Federal Reserve Bank of St. Louis.** 2023. “Federal Reserve Economic Data.” <https://fred.stlouisfed.org/>. Last accessed: 2023-01-12. For details, see the online appendix and the readme.pdf file in the replication package.
- Fried, Stephie, 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.
- Grand, François Le, Xavier Ragot, and Diego Rodrigues.** 2025. “The Welfare of Nations: Social Preferences and the Macroeconomy.” , (DP19847).
- Heathcote, Jonathan, Fabrizio Perri, Giovanni L. Violante, and Lichen Zhang.** 2023. “More Unequal We Stand? Inequality Dynamics in the United States, 1967–2021.” *Review of Economic Dynamics*, 50: 235–266.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2010. “The Macroeconomic Implications of Rising Wage Inequality in the United States.” *Journal of Political Economy*, 118(4): 681–722.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante.** 2017. “Optimal Tax Progressivity: An Analytical Framework.” *The Quarterly Journal of Economics*, 132(4): 1693–1754.
- Huggett, Mark, and Juan Carlos Parra.** 2010. “How Well Does the U.S. Social Insurance System Provide Social Insurance?” *Journal of Political Economy*, 118(1): 76–112.
- Kanzig, Diego.** 2023. “The Unequal Economic Consequences of Carbon Pricing.”
- 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.
- 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.
- Krueger, Dirk, and Chunzan Wu.** 2025. “Taxes on Lifetime Income: A Good Idea?”
- Kuhn, Moritz, and Lennard Schlattmann.** 2024. “Distributional Consequences of Climate Policies.” C.E.P.R. Discussion Papers CEPR Discussion Papers 18893.
- Labrousse, Charles, and Yann Perdereau.** 2026. “Geography versus income: The heterogeneous effects of carbon taxation.” *Journal of Monetary Economics*, 161: 103962.
- 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.** 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. Yale University Press.
- Olivi, Alan, Vincent Sterk, and Dajana Xhani.** 2026. “Optimal Monetary Policy during a Cost-of-Living Crisis.” *American Economic Review*. Forthcoming.
- Oni, Mehedi Hasan.** 2024. “Commuting, Home Utilities, and Production: The Distributional Effects of Energy Price Shocks.” *Unpublished manuscript*.
- 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).
- Stadler, Konstantin, Richard Wood, Tatyana Bulavskaya, Carl-Johan Södersten, Moana Simas, Sarah Schmidt, Arkaitz Usubiaga, José Acosta-Fernández, Jeroen Kuenen, Martin Bruckner, Stefan Giljum, Stephan Lutter, Stefano Merciai, Jannick H. Schmidt, Michaela C. Theurl, Christoph Plutzar, Thomas Kastner, Nina Eisenmenger, Karl-Heinz Erb, Arjan de Koning, and Arnold Tukker.** 2018. “EXIOBASE 3: Developing a Time Series of Detailed Environmentally Extended Multi-Regional Input-Output Tables.” *Journal of Industrial Ecology*, 22(3): 502–515.
- Stern, Nicholas.** 2008. “The Economics of Climate Change.” *American Economic Review*, 98(2): 1–37.
- Stern, Nicholas, Joseph Stiglitz, and Charlotte Taylor.** 2022. “The economics of immense risk, urgent action and radical change: towards new approaches to the economics of climate change.” *Journal of Economic Methodology*, 29(3): 181–216.
- 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. Bureau of Economic Analysis.** 2026*a*. “Fixed Assets Table 1.1: Current-Cost Net Stock of Fixed Assets and Consumer Durable Goods.” <https://apps.bea.gov/iTable/?reqid=10&step=2&isuri=1>. Last accessed: 2026-06-24.
- U.S. Bureau of Economic Analysis.** 2026*b*. “NIPA Table 1.1.5: Gross Domestic Product.” <https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=survey>. Last accessed: 2026-06-24.
- U.S. Bureau of Economic Analysis.** 2026*c*. “NIPA Table 3.9.5: Government Consumption Expenditures and Gross Investment.” <https://apps.bea.gov/iTable/?reqid=19&step=2&isuri=1&categories=survey>. Last accessed: 2026-06-24.
- U.S. Bureau of Economic Analysis.** 2026*d*. “PCE Bridge Tables, Detail Level, 2012 and 2017 Benchmark Input-Output Accounts.” <https://www.bea.gov/industry/input-output-accounts-data>. Last accessed: 2026-05-17.
- U.S. Bureau of Economic Analysis.** 2026*e*. “Use Tables, After Redefinitions, Purchaser Value, Detail Level, 2012 and 2017 Benchmark Input-Output Accounts.” <https://www.bea.gov/industry/input-output-accounts-data>. Last accessed: 2026-05-20.
- U.S. Bureau of Economic Analysis and U.S. Bureau of Labor Statistics.** 2025. “Integrated Industry-Level Production Account (KLEMS), 1997–2023.” <https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems>. Last accessed: 2025-05-21.
- U.S. Environmental Protection Agency.** 2024. “Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990–2022.”

Online Appendix

A Data Construction

Emission intensities We compute basic-price CO₂ intensities from the EXIOBASE 3.10.1 product-by-product tables, taking the Leontief inverse of the global technology matrix so that each product’s intensity embodies the emissions produced at every stage of production, domestic and imported. We convert the intensities from euros to dollars at the annual average exchange rate.

From products to expenditure codes The concordance proceeds EXIOBASE → BEA → PCE → UCC. EXIOBASE products are mapped to BEA 2007 detail commodity codes using the USEEIO concordance provided by EXIOBASE. Where several EXIOBASE products map to one BEA code we take the expenditure-weighted average intensity across the categories. BEA commodities are then mapped to PCE consumption categories using the BEA PCE bridge, and PCE categories to UCCs using the BLS CE–PCE concordance. The BLS concordance covers individual item UCCs but omits three groups. Aggregate interview codes which the CE uses when households do not itemize are assigned the expenditure-weighted intensity of the corresponding PCE series. Home improvement and repair, which PCE classifies as residential investment rather than consumption, is retained at the intensity of the residential construction sector. Items outside the PCE consumption boundary, principally the costs of home ownership such as mortgage interest and property taxes, together with taxes, fees and transfers receive zero intensity. This is consistent with our expenditure measure which follows Aguiar and Bils (2015) in replacing owner costs with the owners’ estimated rental value of the home. This rental equivalence is mapped to owner-occupied housing services in PCE and carries its intensity, so that housing consumption is priced once as in the national accounts.

Basic to purchaser prices For PCE category c in sample year y , the purchaser-price intensity is

$$I(c, y) = (1 - \theta_c) \sum_j \frac{1}{P_c} \left[\text{prod}_{jc} I_j(y) + \text{trans}_{jc} I_{\text{trans}}(y) + \text{whol}_{jc} I_{\text{whol}}(y) + \text{ret}_{jc} I_{\text{ret}}(y) \right],$$

where prod_{jc} , trans_{jc} , whol_{jc} and ret_{jc} are the producer and margin components of commodity j within category c from the bridge table, P_c is the category’s total purchaser value, and I_{trans} , I_{whol} , I_{ret} are the intensities of the truck transport, wholesale trade and retail trade sectors. This is more precise than applying sector-average margin rates, because the bridge gives the margin decomposition commodity by commodity. The factor θ_c is the share of taxes on products in the category’s purchaser value, taken from the BEA Supply table; downweighting by $(1 - \theta_c)$ ensures that tax dollars, which generate no emissions, are not assigned an intensity. The correction is small in aggregate but material for heavily taxed goods, scaling the tobacco intensity by 0.71 and gasoline by 0.89.

Direct combustion We add EIA estimates of point-of-use combustion emissions for gasoline and natural gas, computed as the fuel’s emission coefficient and heat content divided by its retail price, so that the two components — upstream supply chain and

direct combustion — are distinct physical flows and can be summed. For electricity the EIA estimate replaces rather than supplements the EXIOBASE figure, because electricity is a scope 2 emitter. EIA intensities require no product-tax correction, since they divide emissions by the price households actually pay.

B Further Empirical Results

B.1 Non-energy expenditure and emissions

Figure B.1 shows that, unlike energy emissions, patterns in non-energy emissions across income and age are not well explained by the non-energy expenditure share of households along these dimensions. This is because the emission intensity of non-energy goods changes across income and age groups, whereas it remains essentially flat for energy goods. This reflects changing expenditure shares within the non-energy category towards less emission-intensive goods for richer and younger households. Note, however, that the actual quantitative significance of these patterns is extremely small, and we therefore omit this channel from our quantitative model.




Figure B.1: CO₂/\$ of expenditure from non-energy vs. non-energy share of expenditure

B.2 Energy share by expenditure quintile




Figure B.2: Energy expenditure share by expenditure quintile, model and data

C Calibration

C.1 Life-cycle parameters

We fit a quartic to the life-cycle profile in Huggett and Parra (2010), and take survival probabilities from the Human Mortality Database for the U.S. in 2019, with the exception of the final age, when households die with certainty by assumption. The resulting values are displayed in Table C.1.

Age	κ_h	S_h	Age	κ_h	S_h
25	1.0000	0.9989	55	2.0231	0.9939
26	1.0889	0.9989	56	2.0270	0.9935
27	1.1701	0.9988	57	2.0281	0.9929
28	1.2443	0.9987	58	2.0262	0.9924
29	1.3119	0.9987	59	2.0209	0.9917
30	1.3737	0.9986	60	2.0117	0.9910
31	1.4301	0.9986	61	1.9981	0.9904
32	1.4817	0.9986	62	1.9798	0.9896
33	1.5289	0.9985	63	1.9562	0.9890
34	1.5722	0.9985	64	1.9268	0.9883
35	1.6121	0.9984	65	0.0000	0.9875
36	1.6488	0.9983	66	0.0000	0.9867
37	1.6828	0.9982	67	0.0000	0.9857
38	1.7144	0.9981	68	0.0000	0.9849
39	1.7438	0.9981	69	0.0000	0.9838
40	1.7712	0.9979	70	0.0000	0.9820
41	1.7970	0.9979	71	0.0000	0.9805
42	1.8213	0.9978	72	0.0000	0.9788
43	1.8443	0.9977	73	0.0000	0.9767
44	1.8659	0.9976	74	0.0000	0.9748
45	1.8865	0.9973	75	0.0000	0.9712
46	1.9059	0.9972	76	0.0000	0.9687
47	1.9242	0.9969	77	0.0000	0.9656
48	1.9415	0.9967	78	0.0000	0.9625
49	1.9576	0.9964	79	0.0000	0.9584
50	1.9724	0.9961	80	0.0000	0.9532
51	1.9859	0.9958	81	0.0000	0.9485
52	1.9980	0.9954	82	0.0000	0.9425
53	2.0083	0.9951	83	0.0000	0.9355
54	2.0167	0.9945	84	0.0000	0.0000

Table C.1: Life-cycle labor efficiency and survival probabilities

C.2 Initial asset distribution

Following Krueger and Wu (2025), newborn households are endowed with assets drawn from a discretization of the empirical net worth distribution of young households. We take household heads aged 23 to 27 in the 2010, 2013, 2016 and 2019 waves of the SCF — a window around the model’s entry age of 25, chosen for cell size — and construct fifty mass points of equal probability by evaluating the weighted empirical CDF of net worth at the midpoint quantiles. The resulting grid runs from $-\$96,214$ to $\$479,010$. The lowest mass point also serves as the borrowing limit for working-age households.

Assets are drawn independently of initial productivity. Krueger and Wu (2025) report a correlation between net worth and log wages of approximately zero among the young; in our sample the Pearson correlation is 0.12. The initial asset distribution is funded by accidental bequests, with any surplus returned to the government.

C.3 Intermediate good production parameters

The production function parameters χ_s and ϕ_s are the shares of energy and non-energy intermediate inputs in the gross output of sector s . We measure them from the BEA Use tables at the detail level, after redefinitions and at purchaser prices, for the 2012 and 2017 benchmark years, and average the two for our parameter values. The energy sector comprises oil and gas extraction, petroleum and coal products, electric power generation and natural gas distribution, aligned with how this is defined in the consumption data. The non-energy sector is every other industry. For each sector, χ_s is the sum of purchases of energy commodities by its industries divided by their gross output, and ϕ_s the same for non-energy commodities. This gives $\chi_1 = 0.026$, $\phi_1 = 0.404$, $\chi_2 = 0.402$ and $\phi_2 = 0.184$. Three details are worth noting: First, purchaser prices place trade and transport margins inside the input cells, so the coefficients are on the same delivered-price basis as the demand parameters estimated from the CEX. Second, we do not net out imports. As the model economy is closed, its energy sector must supply total U.S. energy absorption, and χ_2 accordingly includes refiners' purchases of crude and other feedstocks. Third, the capital and labor split of the remaining share follows Barrage (2024).

C.4 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 tonne 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. When employed in our model, we multiply the function by 1.31 to convert it into 2019\$

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

$\bar{a}$	0.0662
P_t^{backstop}	$0.567(1 + \exp(-0.05 \cdot t))$
a_t	$49.8896 + 0.8551 \log(t)$
b_{0t}	$14.3338 - 6.4698 \log(t)$
b_{1t}	$15.1937 - 6.6864 \log(t)$
b_2	9.468×10^{-4}
b_x	2.6931

Table C.2: 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.

C.5 Pension system

Following Krueger and Wu (2025), pensions are a piecewise linear function of earnings, as with the true social security formula in the United States. Households receive a pension according to their productivity state at retirement. Writing $b_i = \beta_i \bar{y}_t$ for the bend points, where $\bar{y}_t$ is mean earnings among earners in year t , benefits are

$$P_t(z_{40}) = \min \left\{ \tilde{P}_t(z_{40}), \bar{P}_t \right\}, \quad \tilde{P}_t(z) = \begin{cases} r_1 y_t(z) & y_t(z) \leq b_1, \\ r_1 b_1 + r_2 (y_t(z) - b_1) & b_1 < y_t(z) \leq b_2, \\ r_1 b_1 + r_2 (b_2 - b_1) + r_3 (y_t(z) - b_2) & y_t(z) > b_2, \end{cases}$$

where $y_t(z)$ is mean earnings among households of productivity type z , $\{r_1, r_2, r_3\}$ are progressive replacement rates and the maximum benefit $\bar{P}_t = r_1 b_1 + r_2 (b_2 - b_1) + r_3 (b_3 - b_2)$ is the benefit accruing at the third bend point. We set $(\beta_1, \beta_2, \beta_3) = (0.21, 1.29, 2.42)$ and $(r_1, r_2, r_3) = (0.90, 0.32, 0.15)$ again following Krueger and Wu (2025).

C.6 Initial productivity distribution

The parameter ζ determines productivity dispersion at model age 1 in the following way. The six ordinary states are ordered as the two fixed-effect types crossed with the three values of the persistent component, and the superstar state z_7 sits ι_1 above the highest ordinary state in logs:

$$\mathbf{z} = [z_1 \quad z_2 \quad z_3 \quad z_4 \quad z_5 \quad z_6 \quad \exp(\log z_6 + \iota_1)]$$

Entrants draw from

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

so that within each fixed-effect type a fraction ζ of entrants begins at one of the two extreme values of the persistent component and $1 - \zeta$ at its middle value. Since $\zeta < 1/2$, entrants are less dispersed than the stationary distribution of the chain, and dispersion in the persistent component rises with age. We assume that no one begins life in the superstar state. This is the same way that Kaymak, Leung and Poschke (2026) handle the initial distribution.

C.7 Distributional calibration targets

Survey of Consumer Finances Our data construction and definitions follow Kaymak, Leung and Poschke (2026); our only adjustment is to calculate the relevant distributional moments across the 2010, 2013, 2016 and 2019 waves of the SCF, rather than the broader 2001–2019 sample they use. Wealth is SCF net worth, and households with negative net worth are retained. The wealth Gini (0.85) and the top-1% share (0.36, the share of net worth held above the weighted 99th percentile) are therefore computed over the full distribution, consistent with the model, which allows borrowing. Earnings are defined as wage income, with active business owners assigned imputed labor earnings as in the appendix to Kaymak, Leung and Poschke (2026), so that entrepreneurs remain in the sample. The working-age earnings Gini (0.51) is computed over heads aged 25 to 64 with positive earnings, matching the model’s working-age population.

Panel Study of Income Dynamics The variances of log earnings at ages 25–29 and 50–54 are computed from the PSID extract in the replication package of Heathcote et al. (2023), which covers the SRC and immigrant core samples and applies their sample selection criteria. We use household heads for income years 2010 to 2018 (from the survey waves 2011 to 2019) and keep observations with positive labor income and at least 260 annual hours, following Heathcote, Storesletten and Violante (2010). Earnings are the annual labor earnings of the household head. Log earnings are residualized around age cells within each survey year, and the residual variances are then pooled within each five-year window, giving 0.52 at ages 25–29 and 0.76 at ages 50–54. The model counterparts are the variances of log earnings among households in the same age windows in the stationary equilibrium.

C.8 Aggregate calibration targets

Aggregate ratios All aggregate ratios targeted in calibration are averages of annual data from 2010–2019 inclusive. Our model has no counterpart to government assets or investment. Consequently, the aggregate ratios we compute in the data are adjusted in the following ways: output is GDP (NIPA Table 1.1.5, line 1) less government gross investment (NIPA Table 3.9.5, line 3), capital stock is the current-cost net stock of private fixed assets (Fixed Assets Table 1.1, line 3), investment is gross private domestic investment (NIPA Table 1.1.5, line 7), and government spending is consumption expenditure only (NIPA Table 3.9.5, line 2). The resulting values are $K/Y = 2.37$, $I/Y = 0.175$ and $G/Y = 0.15$. The data counterpart to government debt is federal debt held by the public, equal to 71% of GDP over the same window (FRED series FYGFGDQ188S). The first two aggregate ratios are targeted in calibration, while the latter two are imposed as a requirement of equilibrium.

C.9 Model implied emissions

Total emissions in the model are $(1 - \mu_t)E_t M$, where M converts model units of energy into tonnes of carbon. Rather than estimate M , we choose the units in which energy is measured so that $M = 1$. This is a normalization: rescaling energy by a factor λ leaves every real allocation unchanged provided the energy taste weight and the two sectoral productivity terms are adjusted to $b_e \lambda^{1-\sigma}$, $Z_1 \lambda^{-\chi_1}$ and $Z_2 \lambda^{1-\chi_2}$. We choose λ so that gross energy production in the initial steady state equals U.S. carbon emissions in 2019, 1.32 GtC, which we take from U.S. Environmental Protection Agency (2024). Aggregate emissions are then read directly off the model as $(1 - \mu_t)E_t$, and the price of a unit of energy in the resulting equilibrium is approximately \$16,000.

C.10 Temperature change

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), contributes to two ‘temperature boxes’ which increase global temperature T_t according to:

$$FORC_t = \frac{F_{2x}}{\ln(2)} \ln \left(\frac{MAT_t}{MAT_{1765}} \right) + FORC_t^{EX}$$

$$Tbox_{j,t+1} = Tbox_{j,t} \cdot \exp \left(\frac{-\Delta}{d_j} \right) + teq_j \cdot FORC_{t+1} \left(1 - \exp \left(\frac{-\Delta}{d_j} \right) \right) \quad \text{for } j \in \{1, 2\}$$

$$T_t = Tbox_{1,t} + Tbox_{2,t}.$$

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:

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

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

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

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

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

All other parameters are exogenous. We take the initial conditions of the above dynamic system from Barrage (2024) and iterate forward with a given model implied path of emissions to calculate temperature change by 2100. To be precise, there is a negligible discrepancy in our implementation of Barrage (2024)’s initial conditions in that we apply them to our system which starts in 2019 instead of 2020. We leave emissions from the rest of the world, E_{ROW}^M , at their business-as-usual level as defined by RICE-2010.

D Appendix: Decomposition of Indirect Welfare Gains




Figure D.1: Decomposition of indirect welfare effects, Paris Agreement exercise




Figure D.2: Decomposition of indirect welfare effects, constant tax path exercise




Figure D.3: Decomposition of indirect welfare effects, optimal tax path exercise