

How Should Carbon Tax Revenue be Recycled?

Stephie Fried*, Kevin Novan†, William B. Peterman‡

August 25, 2020

Abstract

This paper explores how carbon-tax revenue should be recycled back to households to maximize welfare. Using a general equilibrium lifecycle model calibrated to reflect the heterogeneity in the U.S. economy, we find the optimal policy uses two thirds of carbon-tax revenue to reduce the distortionary tax on capital income while the remaining one third is used to increase the progressivity of the labor-income tax. The optimal policy attains higher welfare and more equality than the lump-sum rebate approach preferred by policymakers as well as the approach originally prescribed by economists – which called exclusively for reductions in distortionary taxes.

Keywords: Carbon tax; overlapping generations; revenue recycling
JEL codes: E62; H21; H23

*Arizona State University, W.P. Carey School of Business. Email: sdfried@asu.edu

†University of California, Davis, Department of Agricultural and Resource Economics. Email: knovan@ucdavis.edu

‡Federal Reserve Board of Governors. Email: william.b.peterman@frb.gov. Views expressed in this paper are those of the authors' and do not reflect the views of the Federal Reserve System or its staff.

1 Introduction

Policymakers face two fundamental questions when designing a carbon tax. First, at what level should the tax be set? Second, what should be done with the new stream of government revenue? While numerous studies shed light on the optimal level and trajectory for a carbon tax (e.g., Acemoglu et al. (2012), Golosov et al. (2014), Barrage (2018), Lemoine and Rudik (2017)), economists have yet to identify the welfare-maximizing way to return carbon-tax revenue to the public.

In this paper, we draw on an approach from the macro public finance literature (e.g., Conesa et al. (2009), Heathcote et al. (2017)) and solve for the welfare-maximizing way to recycle carbon-tax revenue.¹ To do so, we construct a general equilibrium model calibrated to reflect the heterogeneity in the U.S. economy. We find that the revenue-neutral carbon-tax policy that maximizes the expected steady-state welfare recycles revenue back to the public using two mechanisms. Two thirds of the carbon-tax revenue is used to reduce the distortionary tax on capital income while the remaining one third is used to increase the progressivity of the labor-income tax. While we find that the carbon tax itself is regressive, the optimal recycling approach more than unwinds this regressivity, resulting in a progressive overall change to the tax system.

The optimal recycling approach we identify differs from the method originally prescribed in the economics literature. Environmental and public economists have traditionally called for carbon-tax revenue to be returned exclusively through reductions in preexisting, distortionary taxes – the approach that maximizes economic surplus (Parry (1995), Goulder (1995), de Mooij and Bovenberg (1998), Bovenberg (1999)). However, this literature largely abstracts from heterogeneity and the welfare consequences arising from the distributional impacts of carbon-tax policies. Accounting for heterogeneity, we find it is welfare maximizing to instead use a substantial portion of the revenue to increase equality.²

¹Following much of the literature studying revenue-recycling options, we do not model the environmental benefits from carbon tax policies. Rather, we focus on the non-environmental welfare consequences.

²We focus exclusively on revenue-neutral carbon tax policies. Related research considers non-revenue-neutral approaches (Carbone et al. (2013)).

Our optimal recycling method also generates higher welfare and more equality than the approach advocated for by many policymakers. Motivated by distributional concerns, the carbon-tax policy proposals garnering the greatest support among policymakers call for carbon-tax revenue to be returned to individuals through equal, lump-sum rebates.³ However, our results demonstrate that providing lump-sum rebates is not the optimal way to obtain a progressive outcome. Policymakers can achieve an even more progressive outcome with far higher welfare by instead increasing the progressivity of the labor tax.

To explore whether our optimal rebate is unique to a carbon tax, we conduct a second experiment in which we search for the optimal recycling approach assuming that the new stream of revenue is provided exogenously instead of coming from a carbon tax. Like in the carbon tax experiment, we find that it is optimal to use one portion of the revenue to reduce the capital tax and the other portion to increase the progressivity of the labor tax. However, the fraction of revenue used to reduce the capital tax is smaller when the stream of revenue is exogenous instead of from a carbon tax. This is because the carbon tax itself depresses capital. Hence, when the revenue comes from a carbon tax, it is optimal to mitigate this effect by using relatively more of the revenue to reduce the capital tax.

The novel insights provided by our analysis stem from the combination of two modeling innovations. First, we construct a quantitative overlapping generations model that incorporates three types of heterogeneity that are crucial to quantifying the distributional impacts, and consequently the aggregate welfare impacts, of a carbon tax. To begin, we model heterogeneity over the lifecycle. In addition, we include idiosyncratic shocks to labor-income, which produce a full income distribution within each age group.⁴ Finally, our model uses Stone-Geary preferences to capture the fact that low-income agents use a higher fraction of

³The Carbon Dividend proposal, put forward by the Climate Leadership Council (CLC), calls for the U.S. federal government to institute a carbon tax and return the revenue “directly to U.S. citizens through equal lump-sum rebates.” See “Economists’ Statement on Carbon Dividends,” January 16th, 2019 *Wall Street Journal*. Similarly, Canada’s recently adopted climate policy returns revenues to households through lump-sum payments, which The Citizens’ Climate Lobby Canada (CCL) states will “equitably recycle the revenue obtained from carbon fees” (CCL (2018)).

⁴Related work examines the welfare impacts of carbon tax policies using lifecycle models with within-cohort heterogeneity (Chiroleu-Assouline and Fodha (2014), Williams et al. (2015), Fried et al. (2018)). However, these studies have all focused on a small subset of revenue-recycling options while the present analysis searches over a continuum of rebate options.

their expenditures for energy, implying that the carbon tax by itself is regressive (Metcalf (1999), Grainger and Kolstad (2010)).⁵

The second key modeling innovation centers around the set of revenue recycling options we consider. The existing literature studying revenue-neutral carbon taxes has primarily focused on a small set of blunt approaches for recycling carbon-tax revenues – i.e. returning revenue exclusively through lump-sum rebates, a reduction in the capital-income tax rate, or a reduction in the labor-income tax rate.⁶ In practice, however, policymakers have a much broader set of options at their disposal. To capture this fact, we model an entire continuum of potential rebate approaches. In particular, we consider convex combinations of the following four rebate options for the carbon-tax revenue: (i) reduce the capital-income tax, (ii) reduce the level of the labor-income tax, (iii) increase the progressivity of the labor-income tax, and (iv) provide direct rebate payments that may or may not vary with an agent’s income. Intuitively, the first two rebate mechanisms allow policymakers to unwind the distortions caused by the preexisting labor or capital-income taxes while the second two rebate mechanisms provide options for policymakers to achieve a more progressive outcome.

The rich set of rebate options we consider are crucial for uncovering the result that the welfare-maximizing policy results in a progressive change to the tax system. If, as in much of the previous literature, policymakers can only use lump-sum rebates to achieve a progressive outcome, we find that it is optimal to return all of the revenue through a reduction in the capital tax, using none to increase equality. A key reason why our optimal policy differs from this standard result is that the lump-sum rebate is an ineffective way to increase equality. By providing uniform payments to all agents, lump-sum rebates do a poor job of targeting revenue back towards low-income agents. Moreover, by providing payments to agents of all ages, including retirees, lump-sum rebates reduce the need to save for retirement, crowding out capital. In contrast, by increasing the progressivity of the labor tax, policymakers are able to target carbon-tax revenues towards low-income, working-age individuals, making it

⁵Related work by Cronin et al. (2019) finds that the direct regressivity of a carbon tax is mitigated once one accounts for the indexing of transfers, an approach we follow in our subsequent analysis.

⁶In one notable exception, Goulder et al. (2019) consider combinations of lump-sum rebates and reductions in federal taxes in an infinitely-lived agent model.

a far more effective mechanism for increasing equality.

Stepping back, the analysis presented in this paper highlights the value of bringing the modeling tools from the macroeconomic literature to bear on a question traditionally studied by environmental and public economists. The macro public finance literature has long used general equilibrium, lifecycle models with rich within-cohort heterogeneity to quantify the welfare and distributional effects of alternative tax policies. This literature has primarily focused on which taxes to use to achieve a given revenue target. Instead, we focus on which taxes to decrease, given a new stream of revenue from a carbon tax, to satisfy the same revenue target. By using the macro modeling tools to incorporate heterogeneity, we are able to provide a much more thorough understanding of the welfare and distributional consequences of potential carbon tax policies.

2 Model

2.1 Demographics

Our model incorporates overlapping generations of agents. Agents enter the model when they start working, which we approximate with a real-world age of 20. Each period, agents age one year and a continuum of new 20-year-olds enters the model. The size of the new-born cohort grows exogenously at rate n . Agents make labor-supply and savings decisions each period until they are forced to retire at a real-world age of 65. Retired agents finance consumption from Social Security payments and accumulated assets. Lifetime length is uncertain and mortality risk varies over the lifetime.⁷ Since individuals are uncertain how long they will live, they may die with positive asset holdings. We treat these assets as accidental bequests and redistribute them as lump-sum transfers T_t^a across individuals during period t .

⁷We impose a maximum age of 100.

2.2 Agents

Agents maximize the expected sum of discounted utility. We model agents as having time-separable preferences specified by:

$$U(\tilde{c}_{i,j,t}, h_{i,j,t}) = \frac{\tilde{c}_{i,j,t}^{1-\theta_1}}{1-\theta_1} - \chi \frac{h_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}}, \quad (1)$$

where $\tilde{c}_{i,j,t}$ represents the level of a composite good consumed by agent i , at age j , during period t and $h_{i,j,t}$ represents the hours worked. θ_1 is the coefficient of relative risk aversion and θ_2 is the Frisch elasticity of labor supply. χ determines the dis-utility of hours.

The composite good is comprised of a generic consumption good and carbon-emitting energy, capturing the fact that energy is not only used in production, but also directly by agents (e.g., gasoline). Importantly, previous work highlights that the share of expenditures that goes towards energy differs systematically across agents – with lower-income groups devoting a larger share of their budgets to energy (Metcalf (2007), Hassett et al. (2009)). Following Fried et al. (2018), we capture this negative relationship between income and energy budget shares by assuming that agents must consume a minimum amount of energy, \bar{e} , and that agents derive no utility from energy consumed up to this subsistence level. In particular, composite consumption is given by $\tilde{c}_{i,j,t} = c_{i,j,t}^\gamma (e_{i,j,t}^c - \bar{e})^{1-\gamma}$, where $c_{i,j,t}$ and $e_{i,j,t}^c$ denote the levels of the generic good and energy consumed, respectively.

Agents are endowed with one unit of time each period which they divide between labor and leisure. To generate a realistic distribution of income, we allow labor productivity to vary across agents and over time. In period t , at age j , agent i earns labor income $y_{i,j,t}^h \equiv w_t \cdot \mu_{i,j,t} \cdot h_{i,j,t}$, where w_t is the wage-rate, $h_{i,j,t}$ denotes hours worked, and $\mu_{i,j,t}$ is the agent's idiosyncratic productivity. Following Kaplan (2012), the log of an agent's idiosyncratic productivity consists of four additively separable components:

$$\log \mu_{i,j,t} = \epsilon_j + \xi_i + \nu_{i,j,t} + \pi_{i,j,t}. \quad (2)$$

ϵ_j governs age-specific human capital and evolves over the lifecycle in a predetermined manner. $\xi_i \sim NID(0, \sigma_\xi^2)$ is an agent-specific fixed effect observed when an agent enters the model. $\pi_{i,j,t} \sim NID(0, \sigma_\pi^2)$ is an idiosyncratic transitory productivity shock, and $\nu_{i,j,t}$ is an idiosyncratic persistent productivity shock which follows a first-order autoregressive process:

$$\nu_{i,j,t} = \rho\nu_{i,j-1,t-1} + \kappa_{i,j,t} \text{ with } \kappa_{i,j,t} \sim NID(0, \sigma_\kappa^2) \text{ and } \nu_{i,20,t} = 0. \quad (3)$$

To partially self-insure against productivity shocks and to finance consumption during retirement, agents can save by accumulating shares of physical capital, $a_{i,j,t+1}$, which they rent to firms at rate R_t . Capital accumulates according to the law of motion:

$$k_{t+1} = (1 - \delta)k_t + i_t,$$

where δ denotes the depreciation rate and variable i denotes new investment. We define $r_t \equiv R_t - \delta$ to be the agent's net rate of return. Working-age agents can borrow up to an exogenously-determined debt limit: $a_{i,j,t} \geq \underline{a}$.⁸

2.3 Firms

The final good, Y , is produced competitively from capital, K^y , efficiency labor, N^y , and carbon-emitting energy, E^y . Following Golosov et al. (2014), the production technology is Cobb-Douglas between the three inputs:

$$Y_t = A_t^y (K_t^y)^{\alpha_y} (N_t^y)^{1-\alpha_y-\zeta} (E_t^y)^\zeta. \quad (4)$$

A^y denotes total factor productivity. α_y and ζ denote capital share and energy share, respectively. The final good is the numeraire and can be used for consumption and investment.

Carbon-emitting energy is produced competitively from capital, K^e , and efficiency labor,

⁸Agents borrow at the rate of r_t divided by their probability of surviving period t .

N^e , according to the production technology:

$$E_t = A_t^e (K_t^e)^{\alpha_e} (N_t^e)^{1-\alpha_e}. \quad (5)$$

Parameter α_e denotes capital's share in the production of energy.

2.4 Government

The government runs a balanced-budget, pay-as-you-go Social Security system and raises revenue to finance an exogenous level of unproductive spending, G . The Social Security system is financed with a flat tax, τ^s , on labor income, up to a taxable maximum, $y^{h,max}$. In practice, the Social Security benefits provided to retired agents are a concave, piecewise linear function of each agents' average labor earnings over their highest 35 years of earnings. Instead of including an agent's whole history of labor earnings as an additional state variable, we follow Kindermann and Krueger (2018) and approximate lifetime labor earnings using agents' ability, ξ , and the value of the last realization of their persistent wage shocks, ν_{65} . Specifically, we compute $x(\xi, \nu_{65})$, the average lifetime labor earnings over the population, conditional on the ability and final persistent shock values. The social security benefit an agent of type (ξ, ν_{65}) receives during each period of retirement is determined using a piecewise-linear function of $x(\xi, \nu_{65})$ with marginal benefit rates, ϕ_i , $i \in \{1, 2, 3\}$, given by:

$$\begin{aligned} &\phi_1 \text{ for } 0 \leq x < b_1 \\ &\phi_2 \text{ for } b_1 \leq x < b_2 \\ &\phi_3 \text{ for } b_2 \leq x < b_3. \end{aligned} \quad (6)$$

To finance spending G , the government can tax capital income, labor income, and, once a climate policy is adopted, carbon emissions. The government taxes an agent's capital income, $y_{i,j,t}^k$, according to a constant marginal tax rate τ^k . An agent's capital income is the return on her assets plus the return on assets she receives as accidental bequests, $y_{i,j,t}^k \equiv r_t(a_{i,j,t} + T_t^a)$.

Labor income is taxed according to a progressive tax schedule. An agent's taxable la-

bor income is her labor income, $y_{i,j,t}^h$, net of her employer's contribution to Social Security which is not taxable. Thus, $\tilde{y}_{i,j,t}^h \equiv y_{i,j,t}^h - \tau^s \min(y_{i,j,t}^h, y^{h,max})/2$ is the agent's taxable labor income, where $\min(y_{i,j,t}^h, y^{h,max})/2$ is the employer's Social Security contribution. Following the quantitative public finance literature (Benabou (2002), Guner et al. (2014), Heathcote et al. (2017)), we use the following two-parameter function to model total labor income taxes for an agent with labor income $\tilde{y}_{i,j,t}^h$:

$$T^h(\tilde{y}_{i,j,t}^h) = \max \left[1 - \lambda_1 \left(\frac{\tilde{y}_{i,j,t}^h}{\bar{y}_t^h} \right)^{-\lambda_2}, 0 \right] \tilde{y}_{i,j,t}^h, \quad (7)$$

where \bar{y}_t^h is the mean value of taxable labor income in the economy. We bound the labor-tax function below at zero since we do not observe negative labor-income taxes in the U.S.

The function specified by equation (7) allows us to flexibly alter the labor tax following the introduction of a carbon tax. As long as the zero lower-bound does not bind, decreasing λ_1 decreases the after-tax labor-income of all individuals by the same percentage – leaving the distribution of after-tax income across agents unchanged. In contrast, changing λ_2 alters the distribution of after-tax labor income. Increasing λ_2 reduces the average tax rate for low-income households and increases the average tax rate for high-income households, reducing the inequality in the distribution of after-tax labor income.

With the introduction of a climate policy, the government can finance a portion of spending with a carbon tax, τ^c , levied on each unit of carbon-emitting energy consumed.⁹ Using our model, we compare steady-state outcomes across a range of revenue-neutral carbon tax policies. The stationary competitive equilibrium, in which factor prices and aggregate macroeconomic variables are constant, is defined in Appendix A.

⁹Given that fossil fuel combustion accounts for over 80 percent of GHG emissions, a carbon tax behaves much like a tax on energy. This abstracts from substitution between fossil fuel energy sources with varying carbon intensities that could occur with a carbon tax.

3 Calibration

We calibrate the model to match key features of the U.S. economy. We choose one set of parameters from the data and literature. The remaining parameters are set to ensure moments in the model match their values in the data. Appendix B describes the data sources.

3.1 Production

We normalize the total factor productivity in energy and final-good production to unity, $A^e = A^y = 1$. Following Barrage (2018), we set capital's share in energy production equal to 0.597. Following Golosov et al. (2014), we set capital's share in the production of output equal to 0.3 and fossil energy's share in the production output equal to 0.03. We choose the depreciation rate on capital equal to 0.079 to match the investment to output ratio of 23.3 percent.

3.2 Preferences

The discount rate $\beta = 0.995$ is chosen to match the U.S. capital-output ratio of 2.586. Disutility of labor $\chi = 73.3$ is chosen to ensure agents spend an average of one third of their time endowment working. Following Conesa et al. (2009), the coefficient of relative risk aversion, θ_1 , equals two and following Kaplan (2012), the Frisch elasticity of labor supply, θ_2 , equals 0.5. We choose the debt limit, $\underline{a} = -0.156$ to match the ratio of total debt (among individuals with debt) to total savings in the U.S. of 0.05. The conditional survival probabilities are based on estimates in Bell and Miller (2002).

Subsistence energy, \bar{e} , governs how an agent's energy budget share changes with income. Following Fried et al. (2018), we choose $\bar{e} = 0.0013$ to target the energy-share difference between the top and bottom halves of the expenditure distribution based on data from the CEX (see Appendix B). The expression $1 - \gamma$ represents fossil energy's share in the consumption-energy composite, \tilde{c} . All else constant, an increase in γ reduces energy's share

in the consumption-energy composite and thus decreases the agent’s demand for energy. We choose $\gamma = 0.9907$ to match the empirical ratio of energy consumed directly by households to total energy consumption, 0.183.

3.3 Idiosyncratic Labor Productivity

We take the parameters of the idiosyncratic labor productivity processes from Kaplan (2012): $\sigma_\xi^2 = 0.065$, $\sigma_\kappa^2 = 0.017$, $\sigma_\pi^2 = 0.081$ and $\rho = 0.958$.¹⁰ Importantly, the annual variation in labor income that Kaplan (2012) uses to estimate the shock processes includes heads of households who have worked as little as one-quarter of a full-time work-year. Thus, the estimated labor-income process includes variation in annual labor income from any unemployment spells that last less than 39 weeks for a full-time worker. This incorporates the vast majority of unemployed workers.¹¹ The age-specific human capital parameters, $\{\varepsilon_j\}_{j=20}^{100}$ are from Huggett and Parra (2010).¹²

3.4 Government Policy

Government expenditure, $G = 0.106$, is set to ensure it equals 15.7 percent of output. Following Kaplan (2012), the tax rate on capital income, τ^k , is set to 36 percent. We set the Social Security marginal benefit rates, $\phi_1 = 0.9$, $\phi_2 = 0.32$ and $\phi_3 = 0.15$, to match the piecewise-linear benefit function used in the U.S. Social Security system. To determine the benefit function’s knot points, $b_1 = 0.12$, $b_2 = 0.72$ and $b_3 = 1.36$, we set the ratio of the knot point to average labor earnings in the model equal to the corresponding ratio of the actual knot point and the average labor earnings in the data.¹³ We choose the social security

¹⁰We discretize the shocks using two states to represent the transitory and permanent shocks and five states for the persistent shock. To discretize the persistent shock, we use the Rouwenhorst method which is well-suited for discretizing highly persistent shocks with a small number of states (Kopecky and Suen 2010).

¹¹The average U.S. long-term unemployment rate (duration greater than 27 weeks) equals 1 percent, and accounts for less than one quarter of total unemployment. Data are from the BLS, we take the average over the five most recent years, July 2014-July 2019.

¹²The values are displayed in Table 3 of Huggett and Parra (2010). Following Peterman and Sommer (Forthcoming), we extend and smooth the age-specific human capital values to 65 years using a quadratic polynomial.

¹³The maximum taxable labor income is set equal to $y^{h,max} = 1.36$.

tax, $\tau^s = 0.096$, so that the social security budget balances each period. With each carbon tax policy we simulate, we adjust the Social Security benefits so that the purchasing power is unchanged from the pre-carbon-tax baseline steady state (Goulder et al. 2019). Following Guner et al. (2014), we set the curvature parameter of the labor-tax function, λ_2 , equal to 0.031. The parameter determining the level of the labor tax, λ_1 , is set equal to 0.827 to clear the government budget constraint. These parameters imply that an agent with the mean labor income faces an average labor-tax rate of 17.4 percent and a marginal labor-tax rate of 20.0 percent.

To focus exclusively on the welfare consequences of alternative approaches for rebating the resulting carbon-tax revenue, we set the tax on carbon emissions at a fixed level of \$40 dollars per ton of CO₂ – the initial value proposed by the Climate Leadership Council (CLC, 2019). To calibrate the size of the tax in the model, we calculate the empirical value of the tax as a fraction of the price of a fossil energy composite of coal, oil, and natural gas. We calculate the price of this energy composite averaging over the price of each type of energy in each year, and weighting by the relative consumption in each year. Similarly, we calculate the carbon emitted from the energy composite by averaging over the carbon intensity of each type of energy in each year, and weighting by the relative consumption in each year. This process implies that a \$40 per ton carbon tax equals 49 percent of our composite fossil energy price in the baseline steady state, yielding $\tau^c = 0.26$.

4 Computational Experiments

We use the model to study the long-run welfare effects of policies that combine the carbon tax with one or more rebate instruments to return the revenue back to agents.

4.1 Rebate Policies

We allow the policymaker to return the carbon-tax revenue through direct rebate payments and by decreasing existing federal labor and capital tax rates. Since we are focused on

ways to return the carbon-tax revenue, not raise additional revenue, we do not allow the policymaker to increase income taxes for any individual agent. Additionally, following the macro-public finance literature, we do not permit age-dependent taxes and transfers. Based on these criteria, we analyze combinations of the following four rebate instruments: (i) a reduction in the capital tax, (ii) a reduction in the level of the labor tax, (iii) an increase in the progressivity of the labor tax, and (iv) a direct rebate payment that is either uniform across all agents (i.e. a lump-sum rebate) or varies with the agent's income.

The increase in the progressivity of the labor tax is designed to mimic a change in the tax code in which the government reduces the average labor-income tax rate for the lower-income agents but does not change the average labor-income tax rate for higher-income agents. While increasing the curvature parameter, λ_2 , in the labor-tax function (equation (7)) lowers the average labor tax rate for low-income agents, it increases the average labor tax-rate for high-income agents. This change would not constitute a pure rebate because the tax rate increases for a fraction of the population. Therefore, we augment equation (7) to ensure the average tax rate does not increase for any level of labor income. Specifically, the labor tax rate for an individual with taxable labor income, $\tilde{y}_{i,j,t}^h$, is:

$$\max \left[\min \left[1 - \lambda_1 \left(\frac{\tilde{y}_{i,j,t}^h}{\tilde{y}_t^h} \right)^{-\lambda_2}, 1 - \lambda_1' \left(\frac{\tilde{y}_{i,j,t}^h}{\tilde{y}_t^h} \right)^{-\lambda_2'} \right], 0 \right],$$

where parameters λ_1 and λ_2 are the baseline values of the level and curvature parameters and λ_1' and λ_2' are the corresponding values in the counterfactual simulation.¹⁴

We also allow the government to recycle carbon-tax revenue through direct rebate payments that can vary linearly with an agent's total income, y_{ij} , according to the equation:

$$T_{ij}^c = \max [\Upsilon_1 + \Upsilon_2 y_{ij}, 0]. \quad (8)$$

Again, we bound the rebate function below by zero to avoid raising taxes on any agent.

¹⁴To calculate labor taxes in each counterfactual simulation, we keep the value of average taxable labor income, \tilde{y}^h fixed at its value in the baseline.

4.2 Welfare and Distributional Metrics

Following the tradition in the macro literature, we use the consumption equivalent variation (CEV) to quantify the aggregate welfare impacts of each carbon tax policy. This welfare measure is ex-ante in that it depends on the agent’s expected lifetime consumption before information about the agent is revealed. Specifically, the CEV measures the uniform percentage change in an agent’s expected non-energy consumption that is required to make her indifferent – prior to observing her idiosyncratic ability, productivity, and mortality shocks – between the baseline steady state and the steady state under the carbon tax.

To quantify the distributional impacts, we follow Fried et al. (2018) and compute the Gini coefficient for lifetime welfare under each policy. The Gini coefficient ranges between zero (perfect equality) and one (perfect inequality).


5 Quantitative Results

To find the welfare-maximizing rebate, we calculate the steady state with a carbon tax over a grid of different rebate policies. The policies include all combinations of the capital tax, τ^k , the level and progressivity of the labor tax, determined by λ_1 and λ_2 , and the slope, Υ_2 , and intercept, Υ_1 , of the rebate-payment function that clear the government budget constraint and do not increase the capital or labor tax above the baseline levels.

5.1 Optimal Policy

The welfare-maximizing policy uses 62 percent of the revenue to reduce the capital tax by 5 percentage points to 31 percent. The remaining 38 percent of the revenue is used to increase equality, specifically by lowering the labor tax for agents earning low labor income. In particular, Figure 1 highlights that, under the optimal policy, agents with labor-income earnings below 48 percent of the mean see their average labor tax rates fall, with agents earning below 24 percent of the mean paying zero labor taxes.

Figure 1: Rebate From the Increase in Labor Tax Progressivity Under the Optimal Policy



Note: The figure displays the average labor income tax rate paid by an agent under the optimal policy and in the baseline steady state. The average tax rate is displayed as a function of an agent’s labor income relative to the mean level of labor income.

Ultimately, we find that the optimal approach for rebating the revenue eliminates almost all of the ex-ante non-environmental welfare loss from the carbon tax, with the CEV falling by only 0.11 percentage points (Table 1). For comparison, Table 1 also reports the CEV under the rebate approaches typically considered in the literature: exclusively providing uniform lump-sum rebates, reducing the capital tax, or reducing the level of the labor tax (for all agents). Recall, these welfare changes do not incorporate benefits stemming from improved environmental quality. However, Table 1 highlights that the change in energy use, and thus the environmental benefits, are stable across the policies. Therefore, abstracting from the environmental benefits will not impact the relative ranking of the policy options.

Table 1: Welfare and Distribution Effects

	Lump sum rebate	Capital tax rebate	Labor tax rebate	Optimal rebate
CEV	-0.64	-0.27	-0.54	-0.11
Percent change in the welfare Gini	-1.09	-0.00	0.15	-2.35
Percent change in capital	-3.37	1.78	-1.83	-0.55
Percent change in the production energy	-33.02	-31.18	-32.38	-32.13
Percent change in the consumption energy	-27.55	-26.47	-27.10	-27.14

Note: The table displays the CEV and the percentage change the aggregate variables and the Gini coefficient for lifetime welfare relative to their values in the baseline.

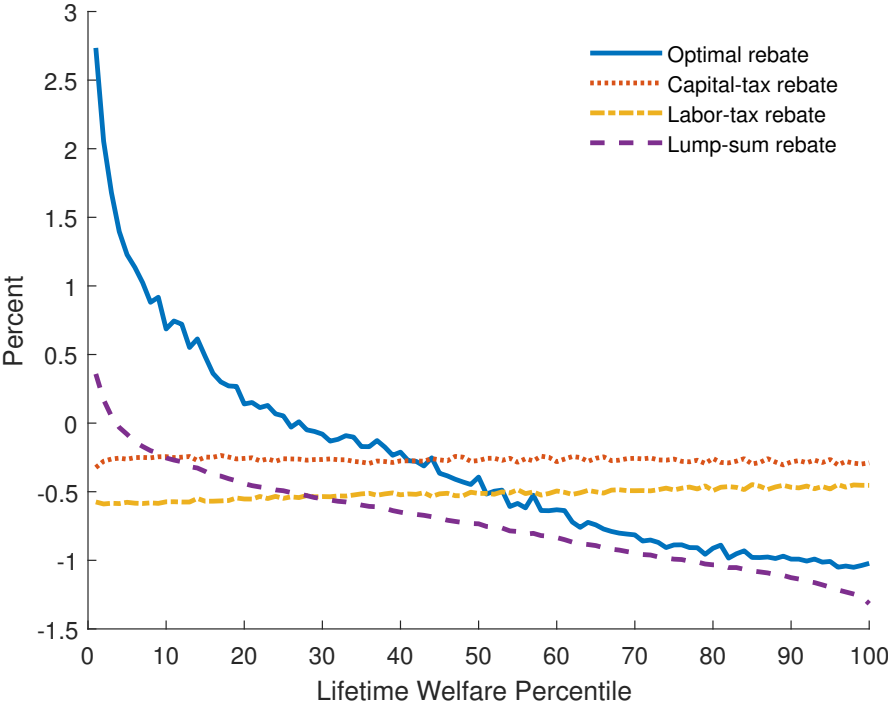
To understand how the optimal policy achieves the highest expected welfare, it is important to note that a recycling approach can boost expected welfare not only by reducing distortionary taxes and increasing economic surplus, but also by redistributing resources away from agents with high levels of lifetime welfare and low marginal utilities of consumption. The optimal policy does both. As Table 1 highlights, the optimal policy achieves a substantial amount of redistribution. The Gini coefficient of lifetime welfare falls by 2.35 percent under the optimal rebate, more than double the decrease under the lump-sum rebate, the most progressive of the three standard instruments. Importantly, the optimal policy achieves this increase in equality using only a fraction of the carbon tax revenue, allowing the remaining revenue to be used to reduce the distortionary capital tax.

To further highlight how the distributional and the expected welfare impacts differ across policies, we examine how each policy affects agents across the entire distribution of lifetime welfare. To do so, we calculate the percentage change in each agent’s baseline consumption that would be required to make her indifferent – after observing her idiosyncratic ability, productivity, and mortality shocks – between the baseline steady state and the steady state under each carbon-tax policy. In contrast to the CEV, which measures the expected ex-ante change in lifetime welfare, this exercise measures the realized ex-post change in lifetime welfare for each agent.

Figure 2 displays how the ex-post welfare impacts vary across agents based on their realized lifetime welfare in the baseline steady state, with the 1st percentile representing

agents with the lowest lifetime welfare. The optimal policy and the lump-sum rebates result in progressive changes to the tax system; the solid blue and dashed purple lines are downward sloping, implying that agents with highest lifetime welfare in the baseline experience the largest percentage declines in ex-post welfare under both policies. In contrast, the capital and labor tax rebate policies have relatively neutral distributional impacts; the orange and purple lines are approximately flat, implying that all agents experience roughly the same percentage decrease in ex-post welfare.

Figure 2: Heterogeneity in Ex-Post Welfare Changes



Note: The vertical axis represents the percentage change in baseline lifetime consumption required to make an agent indifferent between living in the baseline steady state and the steady state under a given climate policy. Agents are separated by their lifetime welfare in the baseline steady state, with the 1st percentile representing the agent with the lowest lifetime welfare.

While the slope of a line reveals the degree of progressivity or regressivity of the policy, the overall average height reflects how the policy impacts an agent’s expected welfare prior to being born. The blue line in Figure 2 is always above the purple line, implying the optimal

policy achieves not only higher expected welfare than the lump-sum rebate policy, but also higher ex-post welfare regardless of an agent's baseline lifetime welfare. The results are mixed, however, comparing the optimal policy to the capital and labor tax rebates. Agents with baseline welfare above the median experience larger welfare losses under the optimal policy compared to the capital or labor tax rebates (the right half of the blue line is below the orange and yellow lines). In contrast, agents with lower levels of lifetime welfare fare far better under the optimal policy (the left half of the blue line is above the orange and yellow lines). Notably, agents below the 25th percentile of lifetime welfare experience welfare gains under the optimal policy. Ultimately, the sizable ex-post welfare gains experienced by agents with the lowest lifetime welfare pull the average height of the blue line up, highlighting the importance of redistribution in arising the expected welfare achieved by the optimal policy.

5.2 Increasing Labor-tax Progressivity vs. Lump-sum Rebates

Our finding that it is optimal to use a portion of the carbon-tax revenue to increase equality ultimately hinges on the ability of the policymaker to rebate the carbon-tax revenue by increasing the progressivity of the labor tax. In this section, we highlight that this result stems from the fact that, by increasing the progressivity of the labor tax, policymakers can target the revenue along two important margins: towards agents that (i) have low income and (ii) are not retired.

To understand why it is important for the recycling mechanism to target low-income agents, note that using any revenue to increase equality reduces the policymaker's ability to unwind the distortions caused by pre-existing taxes. Therefore, for it to be welfare-maximizing to use any revenue to increase equality, the mechanism chosen must effectively redistribute from high to low-welfare agents. Recall from the Gini coefficients in Table 1, far more redistribution is achieved using just 38 percent of the revenue to increase in the progressivity of the labor tax as opposed to returning all of the revenue through lump-sum rebates. This is due to the fact that lump-sum rebates fail to directly target low-income agents while the optimal policy exclusively lowers the labor tax for the lowest earners.

Importantly, by itself, the ability to target low-income agents does not make it optimal to use carbon-tax revenues to increase equality. To demonstrate this point, we search for a ‘restricted’ optimal policy in which the policymaker cannot change the progressivity of the labor tax, but she can still increase equality by providing the direct rebate payments which can vary with an agent’s total income (equation (8)). In this setting, we find that it is optimal to exclusively return revenue through a reduction in the capital tax; none is used for equality-increasing rebate payments. The reason that it is only optimal to increase equality when the policymaker can increase the progressivity of the labor tax stems from differences in how the increase in the labor-tax progressivity and the rebate payments affect savings behavior. Importantly, agents receive the rebate payments in every year of life, including after retirement. As a result, the rebate payment reduces an agent’s need to save for retirement, crowding out capital. In contrast, agents only receive the rebate from the increase in the labor-tax progressivity during their working years. Consequently, it does not crowd out as much capital and, thus, is less costly.

Similarly, by itself, the ability to target rebates towards working-age individuals does not make it optimal to use carbon-tax revenues to increase equality. To highlight this point, we search for a new restricted optimal policy in which the policymaker cannot increase the progressivity of the labor tax, but she can provide uniform lump-sum rebates to working-age agents only. Again, the optimal policy in this case rebates all revenue through a reduction in the capital tax.

It is only when we allow the policymaker to (i) target rebate payments to working-age agents, and (ii) vary the rebate payment with the agent’s income that we find it is optimal to combine reductions in the capital tax with equality-increasing rebate payments. In our baseline scenario in which explicit age-dependent rebates are not an option, the ability to increase the progressivity of the labor tax provides policymakers with a simple way in which to target carbon-tax revenues towards low-income, working-age individuals.

5.3 The Effect of a Carbon Tax on Optimal Policy

At this point, it is worth stepping back and asking whether there is something unique about a carbon tax that makes it optimal to recycle the revenue through a combination of a decrease in the labor tax for low labor income earners and a reduction in the capital tax? Qualitatively, the answer is no but quantitatively, the answer is yes. To highlight these points, we conduct a different experiment. Instead of imposing a carbon tax, we assume that the government receives an exogenous stream of revenue that exactly equals the amount that would be raised by the carbon tax under the optimal policy. We search for the optimal way to recycle this new stream of revenue back to agents.

Again, we find that expected welfare is maximized by recycling revenue through a combination of a reduction in the capital tax and a reduction in the average labor tax for low labor income earners. Compared to the carbon tax experiment, we find that relatively more of the exogenous stream of revenue should be used to decrease the labor tax for low labor-income earners and less should be used to reduce the capital tax. In particular, the optimal rebate under the exogenous revenue stream only reduces the capital tax rate by three percentage points, instead of by five percentage points, while the average labor tax falls for agents earning below 70 percent of the mean labor income, instead of 48 percent.

The quantitative differences between the optimal rebates of carbon tax revenue and the exogenous revenue stream stem from the fact that the carbon tax itself depresses capital.¹⁵ Intuitively, the carbon tax reduces energy use, which, all else constant, decreases the marginal product of capital, leading to lower aggregate savings. To mitigate the drop in capital in the carbon-tax steady state, it is optimal to devote a larger share of revenue to reducing the capital tax. Even so, a sizable portion of the carbon tax revenue is still used to increase equality.

¹⁵If the revenue from the carbon tax is not recycled, and instead, say, thrown into the ocean, we find that, compared to the baseline steady state, capital is reduced by 2.62 percent.

6 Conclusion

The environmental and public economics literature has long studied how to optimally design a revenue-neutral carbon tax. However, this literature has typically focused on a small set of blunt options for rebating carbon-tax revenues (e.g., reducing the labor or capital tax rate versus providing lump-sum rebates). Moreover, much of the literature has abstracted from heterogeneity and has focused exclusively on maximizing economic surplus.

In this paper, we use a lifecycle model with rich heterogeneity to search over a continuum of rebate options to find the welfare maximizing way to recycle carbon-tax revenue. In contrast to the early recommendations from the double-dividend literature calling for carbon-tax revenues to be returned exclusively through reductions in pre-existing distortionary taxes, we find that it is optimal to use a sizable portion of the revenue to increase equality. Importantly, however, the welfare maximizing way to achieve a more progressive outcome is not through the use of lump-sum rebates – the approach that is garnering the greatest support among policymakers. Instead, we find that a more progressive distributional outcome can be achieved with far lower welfare costs by rebating carbon-tax revenues by increasing the progressivity of the labor tax.

References

- Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous,** “The environment and directed technical change,” *American Economic Review*, 2012, 102 (1), 131–66.
- Barrage, Lint,** “Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy,” *Review of Economic Studies*, 2018.
- Bell, Felicitie and Michael Miller,** “Life Tables for the United States Social Security Area 1900-2100,” *Actuarial Study 120, Office of the Chief Actuary, Social Security Administration*, 2002.

- Benabou, Roland**, “Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?,” *Econometrica*, 2002, 70 (2), 481–517.
- Bovenberg, Lans**, “Green Tax Reforms and the Double Dividend: An Updated Reader’s Guide,” *International Tax and Public Finance*, 1999, 6(3), 421–443.
- Carbone, Jared, Richard D. Morgenstern, Robertson C. Williams III, and Dallas Burtraw**, “Deficit Reduction and Carbon Taxes: Budgetary, Economic, and Distributional Impacts,” *Resources for the Future Report*, August 2013.
- CCL**, “A Solution: The Carbon Fee-and-Dividend Proposal,” *Citizens’ Climate Lobby Canada*, 2018.
- Chiroleu-Assouline, Mireille and Mouez Fodha**, “From regressive pollution taxes to progressive environmental tax reforms,” *European Economic Review*, 2014, 69, 126–142.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger**, “Taxing Capital? Not a Bad Idea After All!,” *American Economic Review*, 2009, 99(1), 25–38.
- Cronin, Julie Anne, Don Fullerton, and Steven Sexton**, “Vertical and horizontal redistributions from a carbon tax and rebate,” *Journal of the Association of Environmental and Resource Economists*, 2019, 6 (S1), S169–S208.
- de Mooij, Ruud A. and A. Lans Bovenberg**, “Environmental Taxes, International Capital Mobility and Inefficient Tax Systems: Tax Burden versus Tax Shifting,” *International Tax and Public Finance*, 1998, 5(1), 7–39.
- Fried, Stephie, Kevin Novan, and William B. Peterman**, “The distributional effects of a carbon tax on current and future generations,” *Review of Economic Dynamics*, 2018, 30, 30–46.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski**, “Optimal Taxes on Fossil Fuel in General Equilibrium,” *Econometrica*, 2014, 82 (1), 41–88.

- Goulder, Lawrence**, “Environmental Taxation and the Double Dividend: A Reader’s Guide,” *International Tax and Public Finance*, 1995, 2(2).
- Goulder, Lawrence H, Marc A.C. Hafstead, GyuRim Kim, and Xianling Long**, “Impacts of a carbon tax across US household income groups: What are the equity-efficiency trade-offs?,” *Journal of Public Economics*, 2019, 175, 44–64.
- Grainger, Corbett A. and Charles D. Kolstad**, “Who Pays a Price on Carbon?,” *Environmental and Resource Economics*, 2010, 46 (3), 359–376.
- Guner, Nezh, Remzi Kaygusuz, and Gustavo Ventura**, “Income taxation of US households: Facts and parametric estimates,” *Review of Economic Dynamics*, 2014, 17 (4), 559–581.
- Hassett, Kevin, Aparna Mathur, and Gilbert E. Metcalf**, “The Incidence of a U.S. Carbon Tax: A Lifetime and Regional Analysis,” *Energy Journal*, 2009, 30(2), 157–180.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante**, “Optimal tax progressivity: An analytical framework,” *The Quarterly Journal of Economics*, 2017, 132 (4), 1693–1754.
- Huggett, Mark and Juan Carlos Parra**, “How well does the US social insurance system provide social insurance?,” *Journal of Political Economy*, 2010, 118 (1), 76–112.
- Kaplan, Greg**, “Inequality and the Lifecycle,” *Quantitative Economics*, 2012, 3(3), 471–525.
- Kindermann, Fabian and Dirk Krueger**, “High Marginal Tax Rates on the Top 1%? Lessons from a Life Cycle Model with Idiosyncratic Income Risk,” Technical Report, Working Paper 2018.
- Kopecky, Karen A. and Richard M.H. Suen**, “Finite state Markov-chain approximations to highly persistent processes,” *Review of Economic Dynamics*, 2010, 13 (3), 701–714.

- Lemoine, Derek and Ivan Rudik**, “Steering the climate system: using inertia to lower the cost of policy,” *American Economic Review*, 2017, 107 (10), 2947–57.
- Metcalf, Gilbert E.**, “A Distributional Analysis of Green Tax Reforms,” *National Tax Journal*, 1999, pp. 655–681.
- , “A Proposal for a US Carbon Tax Swap: An Equitable Tax Reform to Address Global Climate Change,” *Hamilton Project, Brookings Institute*, 2007.
- Parry, Ian W.H.**, “Pollution Taxes and Revenue Recycling,” *Journal of Environmental Economics and Management*, 1995, 29 (3), S64–S77.
- Peterman, William B. and Kamila Sommer**, “How Well Did Social Security Mitigate the Effects of the Great Recession?,” *International Economic Review*, Forthcoming.
- Williams, Roberton III, Hal Gordon, Dallas Burtraw, and Jared Carbone**, “The Initial Incidence of a Carbon Tax Across Income Groups,” *National Tax Journal*, 2015, 68(1), 195–214.

Appendix

A Definition of an equilibrium

Let $z_{i,j,t} = (j, a_{i,j,t}, \nu_{i,j,t}, \xi_i)$ denote the vector of household state variables and let Z denote the corresponding state space. We define a sequence-of-markets equilibrium for this economy as a sequence of prices, $\{w_t, r_t, p_t^e\}_{t=0}^\infty$, allocations for each household i age j , $\{c_{i,j,t}, e_{i,j,t}^c, a_{i,j,t+1}, h_{i,j,t}\}_{t=0}^\infty$, allocations for firms, $\{E_t^y, K_t^y, N_t^y, K_t^e, N_t^e\}_{t=0}^\infty$, a social security tax, $\{\tau_t^s\}_{t=0}^\infty$, a carbon tax, τ^c , transfers, $\{T_t^a, T_t^c\}_{t=0}^\infty$, and the distribution of individuals over the state space, Φ_t , such that the following holds:

1. Given prices, household allocations maximize:

$$\frac{\tilde{c}_{i,j,t}^{1-\theta_1}}{1-\theta_1} - \chi \frac{h_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} + \mathbb{E} \left\{ \sum_{k=j+1}^J \beta^{k-j} \prod_{q=j}^{k-1} \psi_q \left(\frac{\tilde{c}_{i,j,t}^{1-\theta_1}}{1-\theta_1} - \chi \frac{h_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} \right) \right\},$$

subject to the budget constraint:

$$c_{i,j,t} + (p_t^e + \tau_t^c) e_{i,j,t}^c + a_{i,j,t+1} = \mu_{i,j,t} h_{i,j,t} w_t - T_{i,j,t}^s + (1 + r_t(1 - \tau^k))(a_{i,j,t} + T_t^a) - T_t^h (\mu_{i,j,t} h_{i,j,t} w_t - 0.5 T_{i,j,t}^s) + T_t^c \quad \text{for } j < j^r$$

$$c_{i,j,t} + (p_t^e + \tau_t^c) e_{i,j,t}^c + a_{i,j,t+1} = b^{ss}(x_{i,j,t}) + (1 + r(1 - \tau^k))(a_{i,j,t} + T_t^a) + T_t^c \quad \text{for } j \geq j^r$$

the evolution of labor productivity (equations (2) and (3)) and the non-negativity constraints, $c_t \geq 0$, $a_t \geq 0$, $h_t \geq 0$, and $e_t^c \geq 0$.

2. Given prices, final-good producer allocations solve the profit maximization problem for the representative final good firm:

$$\max_{K_t^y, N_t^y, E_t^y} A_t^y (K_t^y)^{\alpha_y} (N_t^y)^{1-\alpha_y-\zeta} (E_t^y)^\zeta - w_t N_t^y - (r_t + \delta) K_t^y - (p_t^e + \tau^c) E_t^y$$

3. Given prices, energy producer allocations solve the profit maximization problem for the representative energy firm:

$$\max_{K_t^e, N_t^e} p_t^e A_t^e (K_t^e)^{\alpha_e} (N_t^e)^{1-\alpha_e} - w_t N_t^e - (r_t + \delta) K_t^e.$$

4. The markets for capital, labor, and energy clear:

$$\begin{aligned}(1+n)(K_t^y + K_t^e) &= \int a_{i,j,t} d\Phi_t \\ N_t^y + N_t^e &= \int \mu_{i,j,t} h_{i,j,t} d\Phi_t \\ E_t &= E_t^y + \int e_{i,j,t}^c d\Phi_t.\end{aligned}$$

5. The government budget balances:

$$G_t = \int [\tau^k r_t (a_{i,j,t} + T_t^a) + T_t^h (\mu_{i,j,t} h_{i,j,t} w_t - 0.5\tau^s \min(y_{i,j,t}^h, y^{h,max})) + \tau_t^c e_{i,j,t}^c] d\Phi_t + \tau_t^c E_t^y - T_t^c.$$

6. Transfers from accidental bequests satisfy:

$$(1+n)T_{t+1}^a = \int (1 - \psi_j) a_{i,j,t+1} d\Phi_t.$$

7. The Social Security budget clears:

$$\tau^s = \frac{\int T^s(x_{i,j,t}) \Phi_{Z|j \geq j^r}}{\int [\min(y_{ij,t}^h, y^{h,max}) \partial \Phi_{Z|j < j^r}]}.$$

A *stationary competitive equilibrium* consists of prices, $\{w, r, p^e\}$, allocations for firms, $\{E^y, K^y, N^y, K^e, N^e\}$, a social security tax, τ^s , a carbon tax, τ^c , and transfers, $\{T^a, T^c\}$, that are constant over time and satisfy the conditions 2-7. Allocations for households, $\{c_{i,j,t}, e_{i,j,t}^c, a_{i,j,t+1}, h_{i,j,t}\}$, satisfy condition 1. The distribution of individuals over the state space, Φ , is stationary.

B Calibration

We use a five year average from 2013-2017 for all parameter values and targets that we calculate directly from the data. Data on investment, output, and capital are from NIPA Tables 1.1, 1.1.5, and 1.5. We define investment as the sum of investment in private fixed assets and consumer durables and we define capital as the sum of private fixed assets and consumer durables. Data on government budget outlays comes from the CBO.¹⁶ Since our

¹⁶See <https://www.cbo.gov/about/products/budget-economic-data>.

model includes Social Security separate from government spending, we calculate government spending as the difference between total government outlays and Social Security outlays. Data on the carbon intensity, energy prices, and energy consumption are from the EIA.

Using data from the Consumer Expenditures Survey (CEX) spanning 2013 through 2017, we find that the share of expenditures going towards energy is 33.84 percent lower in households in the top half of the total expenditure distribution compared to the bottom half of the total expenditure distribution. However, rather than setting \bar{e} to directly match this difference in the energy expenditure shares, we first must account for the fact that the variance in total expenditures in the CEX is larger than in our model.¹⁷ In particular, the percent difference in total expenditures between the top and bottom half of the expenditure distribution is 288.8 percent in the CEX and 66.7 percent in our model. Following Fried et al. (2018), we deflate the energy expenditure share difference observed in the CEX by $\frac{66.7}{288.8} = 0.231$. To target an energy expenditure share difference between the top and bottom halves of the expenditure distribution of 7.81 percent, we choose $\bar{e} = 0.0013$.

We choose energy-share parameter γ to target the ratio of energy consumed directly by households relative to total energy consumed in the US economy. We calculate the empirical value of E^c/E from data on total primary energy consumption from the Energy Information Administration (EIA). Total fossil energy consumption, E , equals total primary energy consumption of coal, oil, and natural gas reported in EIA Table 1.1. Total fossil energy consumption by individuals, E^c , equals total primary consumption of coal, oil, natural gas by the residential sector (see EIA Table 2.2).¹⁸ The average empirical value of E^c/E over the most recent five years of data, 2013-2017, equals 0.183.

C Computational Experiments

In the simulations, the carbon tax raises the price of the energy-good which reduces the relative price of the numeraire. Since Social Security benefits are denominated in terms of the numeraire, the purchasing power of the Social Security benefits falls from its value in

¹⁷The key reason for the smaller differential in total expenditures in our model is that the productivity shocks are assumed to be log normal. This distributional assumption, while standard in the literature, results in our model failing to capture the extreme top tail of the income distribution. We normalize the CEX data by the square root of family size in all of the calculations.

¹⁸The EIA data report residential energy consumption of coal, oil, natural gas and electricity. To convert residential electricity consumption to primary energy consumption of coal, oil, and natural gas, we calculate household electricity use relative to total electricity use (see EIA Table 7.6). We multiply this fraction the total amounts of coal, oil, and natural gas used in the electricity sector (see EIA Table 2.6).

the baseline. In practice, the U.S. government adjusts Social Security payments each year to ensure that the purchasing power remains constant. Consistent with this policy, we adjust the Social Security payment in each simulation to ensure that the retiree can buy the same bundle of energy and non-energy goods as she could in the baseline steady state. Specifically, Social Security payments in each simulation equal Social Security payments in the baseline times $\frac{c^e(p^e+\tau^c)}{c^e p^e+c}$ where c^e and c are the baseline values of energy and non-energy consumption, respectively. We adjust the Social Security tax to ensure that the Social Security budget balances.