Understanding the Inequality and Welfare Impacts of Carbon Tax Policies

Stephie Fried (ASU and SF Fed)

Kevin Novan (UC Davis)

William B. Peterman (Federal Reserve Board of Governors)

April 11, 2022

Abstract

This paper explores how to recycle carbon tax revenue back to households to maximize welfare. Using a general equilibrium lifecycle model calibrated to reflect the heterogeneity in the U.S. economy, we find that the welfare maximizing rebate uses two thirds of carbon-tax revenue to reduce the distortionary tax on capital income while using the remaining one third to increase the progressivity of the labor-income tax. This recycling approach 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. Importantly, we find that recycling a majority of the revenue by lowering the capital tax is unique to raising revenue through a carbon tax.

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

Email: sdfried@asu.edu, knovan@ucdavis.edu, william.b.peterman@frb.gov. The views expressed in this paper are our own and do not reflect the views of the Federal Reserve System or its staff. For helpful comments, we thank seminar and conference participants at AERE, Fordham University, UPenn and the CESifo Conference on Environmental Economics. All potential errors are our own. Declarations of interest: none. This paper was previously circulated under the title “Recycling Carbon Tax Revenue to Maximize Welfare.”

1 Introduction

From 2019 through 2021, fourteen different carbon tax bills were introduced in the U.S. Congress. On the surface, these bills are quite similar. Nearly all propose a tax starting in the range of \$25 to \$50 per ton of CO_2. Moreover, most call for the majority of the revenue to be recycled back to households. However, in terms of how to recycle the revenue back to households, there is no clear consensus. Many proposals call for the revenue to be returned through tax swaps or periodic, lump-sum payments (i.e. carbon dividends). Others instead propose a more targeted approach intended to achieve a more progressive outcome – returning the revenue through income-tested payments. The objective of this paper is to provide guidance surrounding which recycling options to use to maximize welfare.

Drawing on an approach from the macro public finance literature (e.g., Conesa et al. (2009), Heathcote et al. (2017)), we solve for the welfare-maximizing way to recycle carbon-tax revenue in a general equilibrium model of the U.S. economy. We calibrate the model to not only reflect heterogeneity across agents and their lifecycles, but also to match several important features of the U.S. economy, including the current tax system and the importance of energy in the production of the final good and also in utility. Consistent with the wide array of recycling approaches included in the existing policy proposals, we consider all convex combinations of tax swaps (i.e. uniform reductions in the labor or capital income tax rates), lump-sum payments, and targeted payments that vary with income. Importantly, we allow these targeted payments to vary not only with total household income, as many current proposals call for, but also as a function of labor income specifically. In particular, we allow the policymaker to direct revenue back to agents with low labor income by increasing the progressivity of the labor income tax. Measuring social welfare behind the veil of ignorance, we solve for the revenue recycling approach, or combination of approaches, that maximizes the expected lifetime welfare of an agent born into the future steady-state.[1]

We find welfare is maximized by recycling revenue back to households using two distinct methods. Two thirds of the revenue is used to reduce the marginal tax rate on capital income. The importance of this recycling approach is not a new insight but is instead consistent with a large literature – dating back to the orignal 'double-dividend' studies (Parry (1995), Goulder (1995), de Mooij and Bovenberg (1998), Bovenberg (1999)) – highlighting that efficiency

[1] Following 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. In addition, we focus exclusively on modeling income heterogeneity, abstracting from heterogeneity across other dimensions (e.g., spatial heterogeneity). Recent work by Cronin et al. (2019) explores the potential redistributional impacts of a carbon tax policy within income groups.

gains are achieved by using revenue from Pigouvian taxes to reduce distortionary taxes. The key new insight from our analysis is that welfare is maximized by using the remaining revenue to increase the progressivity of the overall policy, instead of uniformly lowering an existing tax. In particular, we find that the way to do so is by increasing the progressivity of the labor-income tax. We find that these optimal rebate methods do not depend on the sub-optimal mix of income taxes we begin with and they are robust across a wide range of carbon tax levels and specifications for the utility function.

Returning to the recent surge of carbon tax proposals, there is clearly growing support among policymakers for recycling the revenue back to households through direct payments that are either uniform across individuals or varying inversely with a household's total income. One reason that these direct payments are more politically palatable is because they are progressive, with relatively larger benefits accruing to lower income households. Consistent with these policy proposals, we find that it is welfare-maximizing to use a sizable portion of the revenue to increase the progressivity of the policy. However, we find that varying the payments with a household's total income is not an effective way to increase equity. We find that using the revenue to increase the progressivity of the labor income tax achieves far higher equity and welfare than payments that are uniform or vary with a household's total income.

Intuitively, returning carbon tax revenue through an increase in the progressivity of the labor income tax directs revenue back towards working-age agents with low labor income. In contrast, recycling the revenue through payments that vary inversely with total household income will direct revenue back towards any agents, working or retired, with low total income in that period. From an equity perspective, the ideal revenue recycling method would return the revenue to the agents with the highest marginal utility. Ultimately, labor income is more closely correlated with marginal utility than total income because retirees have low total income but, since they are consuming their savings, this does not imply that they have high marginal utility. The key advantage of returning the revenue by increasing the progressivity of the labor tax is that it better targets the revenue to the highest marginal utility agents. These could be agents who just experienced an adverse shock or agents who have low labor income over the entire course of their working lifetime.

The novel insights provided by our analysis stem from the combination of two modeling innovations. First, we don't restrict the recycling options to a small set of blunt approaches – e.g. returning revenue exclusively through lump-sum rebates, exclusively through a reduction in the capital income tax rate, or exclusively through a reduction in the labor income tax rate. Instead, we systematically consider a very rich set of recycling options. Second, the quantitative overlapping generations model we construct incorporates heterogeneity on several dimensions that are crucial for quantifying the distributional and overall welfare consequences of a carbon tax. First, building on recent studies demonstrating that the welfare impacts of a carbon tax can vary across the lifetime and between individuals (Chiroleu-Assouline and Fodha (2014), Williams et al. (2015), Fried et al. (2018)), we model agents' entire lifecycle, generating heterogeneity over age, and include idiosyncratic shocks to labor productivity, generating heterogeneity over income within each age group. Additionally, we use Stone-Geary preferences to capture the fact that low-income agents use a higher fraction of their expenditures for energy (Metcalf (1999), Grainger and Kolstad (2010)).[2]

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. 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.

2.2 Agents

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

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

[2] Recent work by Aubert and Chiroleu-Assouline (2019) and Jacobs and van der Ploeg (2019) also examine the welfare consequences of a carbon tax in models with income heterogeneity and homothetic preferences.

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). 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 $\tilde{c}_{i,j,t} = c^{\gamma}_{i,j,t}(e^c_{i,j,t} - \bar{e})^{1-\gamma}$, where $c_{i,j,t}$ and $e^c_{i,j,t}$ denote the levels of the generic good and energy consumed, respectively.[3]

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 $y^h_{i,j,t} \equiv w_t \cdot \mu_{i,j,t} \cdot h_{i,j,t}$, 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:

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

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

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

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

[3]As an alternative to the Stone-Geary specification, one could assume that agents have heterogeneous preferences over energy consumption and this heterogeneity is correlated with the agent-specific fixed effect in the labor-productivity process in such a way as to generate declining energy budget shares with income.

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: $a_{i,j,t} \geq \underline{a}$.[4]

2.3 Firms

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

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

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. The specification in equation (4) implies that the economy can reduce fossil energy consumption by either reducing total production, or by substituting capital and labor for fossil-energy. Implicitly, this substituted capital and labor corresponds to non-carbon emitting energy or improvements in energy efficiency.

Carbon-emitting energy is produced competitively from capital, K^e, and efficiency labor, N^e, according to the production technology:

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

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

[4] Agents borrow at the rate of r_t divided by their probability of surviving period t.

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, $\phi_i,\ i \in \{1, 2, 3\}$, given by:

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

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, $y^k_{i,j,t}$, 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, $y^k_{i,j,t} \equiv r_t(a_{i,j,t} + T^a_t)$.

Labor income is taxed according to a progressive tax schedule. An agent's taxable labor income is her labor income, $y^h_{i,j,t}$, net of her employer's contribution to Social Security which is not taxable. Thus, $\tilde{y}^h_{i,j,t} \equiv y^h_{i,j,t} - \tau^s \min(y^h_{i,j,t}, y^{h,max})/2$ is the agent's taxable labor income, where $\min(y^h_{i,j,t}, y^{h,max})/2$ 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 $\tilde{y}^h_{i,j,t}$:

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

where $\bar{\tilde{y}}^h_t$ 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.[5] 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.

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 discusses the data sources and summarizes the calibrated parameter values in greater detail.

3.1 Production

We normalize the total factor productivity in energy and final-good production to unity, $A^e = A^y = 1$. Following Barrage (2019), 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 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 follow Bell and Miller (2002) and we impose a maximum age of 100.

[5]Given 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.

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). We also explore the sensitivity of the results across higher and lower values for $\bar{e}$. 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 $\rho = 0958$.[6] 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.[7] The age-specific human capital parameters, $\{\varepsilon_j\}_{j=20}^{100}$ are from Huggett and Parra (2010).[8]

3.4 Government Policy

Government expenditure, $G = 0.106$, is set to ensure it equals 15.7 percent of output. 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

[6]We 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).

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

[8]The 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.

actual knot point and the average labor earnings in the data.[9] We choose the social security 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).

In our subsequent computational experiments, we consider two different pre-carbon tax baseline scenarios: one in which the labor and capital tax parameters are chosen to match the current U.S. labor and capital tax rates and another in which the labor and capital tax rates are set to maximize expected welfare without a carbon tax. To match the current tax rates, we follow Guner et al. (2014) and 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. Following Kaplan (2012), the tax rate on capital income, τ^k, is set to 36 percent.

To focus exclusively on the welfare consequences of alternative approaches for recycling the resulting carbon-tax revenue back to agents, we set the tax on carbon emissions at a fixed level of \$40 dollars per ton of CO_2 – the initial value proposed by the Climate Leadership Council (CLC, 2019). We also explore the sensitivity of the results across different carbon tax levels. 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 48 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 recycling approaches to return the revenue back to agents.

[9]The maximum taxable labor income for Social Security corresponds to the top bend point: $y^{h,max} = 1.36$.

4.1 Recycling Options

We allow the policymaker to return the carbon-tax revenue through direct payments and by altering 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. Therefore, the optimal recycling approach we identify should be viewed as a constrained optimal approach. Based on these criteria, we analyze combinations of the following five recycling instruments: (i) a uniform reduction in the capital income tax, (ii) a uniform reduction in the labor income tax, (iii) lump-sum payments that are uniform across agents, (iv) income-tested payments that vary with an agent's total income, (v) an increase in the progressivity of the labor income tax

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, $\tilde{y}^h_{i,j,t}$, is:

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

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.[10]

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

$$T^c_{ij} = \max\left[\Upsilon_1 + \Upsilon_2 y_{ij},\ 0\right]. \tag{8}$$

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

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

4.2 Welfare and Distributional Metrics

To compare social welfare under alternative policies, we must impose a social welfare function. Following the standard of the macro literature, we measure welfare behind the veil of ignorance. That is, we identify the carbon tax rebate that maximizes the expected welfare of a newborn in the future steady state prior to the realization of any idiosyncratic shocks.

We quantify the change in social welfare caused by a carbon tax policy using the consumption equivalent variation (CEV). Again, 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. Formally, we define the CEV as the value of Ω that solves the equality below:

$$\begin{aligned}
&\mathbb{E}\left\{\sum_{k=1}^{J}\beta^{k-j}\prod_{q=j}^{k-1}\psi_q\left(\frac{[((1+\Omega)\dot{c}_{i,j,t})^{\gamma}(\dot{e}^{c}_{i,j,t}-\bar{e})^{1-\gamma}]^{1-\theta_1}}{1-\theta_1}-\chi\frac{\dot{h}_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}}\right)\right\} \qquad (9)\\
&=\mathbb{E}\left\{\sum_{k=1}^{J}\beta^{k-j}\prod_{q=j}^{k-1}\psi_q\left(\frac{[\hat{c}^{\gamma}_{i,j,t}(\hat{e}^{c}_{i,j,t}-\bar{e})^{1-\gamma}]^{1-\theta_1}}{1-\theta_1}-\chi\frac{\hat{h}_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}}\right)\right\},
\end{aligned}$$

where the 'dots' denote values in the baseline economy without a carbon tax and 'hats' denote values in the counterfactual economy with the carbon tax in place. The expectation is taken over the lifetime draws of the labor-productivity shock. Note, when $\Omega = 0$, the left-hand-side of equation (9) equals the ex-ante expected lifetime welfare for an agent born into the baseline steady state and the right-hand-side of equation (9) equals the ex-ante expected lifetime welfare for an agent born into the counterfactual economy with the carbon tax.

Since our welfare measure is the CEV between two steady states, it captures the long-run welfare consequences of the carbon tax policy. It does not capture the near-term welfare effects of the policy as the economy transitions to the new steady state with the carbon tax in place.[11] Therefore, while our results provide insights surrounding the optimal way to

[11] Fried et al. (2018) highlight that the non-environmental welfare impacts of revenue-neutral carbon tax policies can differ meaningfully in the short vs. long run. In particular, agents nearing or post retirement at the time the carbon tax is adopted can be affected very differently than agents born in the future, long-run steady state. However, the welfare changes experienced by agents who are relatively young when the policy is adopted are similar to the welfare effects for agents born into the future steady state.

rebate carbon tax revenue in the long run, they do not illustrate how to transition to the optimal rebate.

To quantify the distributional impacts of alternative revenue-recycling mechanisms, we follow Fried et al. (2018) and compute the Gini coefficient for lifetime welfare under the carbon tax combined with each rebate. We define the Gini coefficient, $\mathcal{G}$, as:

$$\mathcal{G} = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N}|x_i - x_j|}{2N^2\bar{x}}, \tag{10}$$

where x_i represents lifetime welfare of agent i, $\bar{x}$ is the mean of lifetime welfare, and N is the total number of agents in the economy. The Gini coefficient ranges between zero (perfect equality) and one (perfect inequality). It is of course important to again stress that the cross-sectional heterogeneity in our model arises from differences in agents' productivity and lifetime earnings. Therefore, while the Gini coefficient effectively captures the distribution of the resulting welfare effects across different income groups, it abstracts from the horizontal distributional effects within income groups that would arise due to other dimensions of heterogeneity, such as geography or occupation, that are not included in our analysis.

5 Quantitative Results

To find the welfare-maximizing recycling approach, we calculate the steady state with a carbon tax over a grid of different recycling 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 Recycling Policy

Starting from a baseline representing our current labor and capital taxes, we find the welfare-maximizing policy uses 64 percent of the revenue generated by a carbon tax to reduce the capital tax by 5 percentage points to 31 percent. The remaining 36 percent of the carbon tax revenue is used to increase equality, specifically by lowering the labor tax for agents earning low labor income. Figure 1 highlights that, under the optimal recycling approach, 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

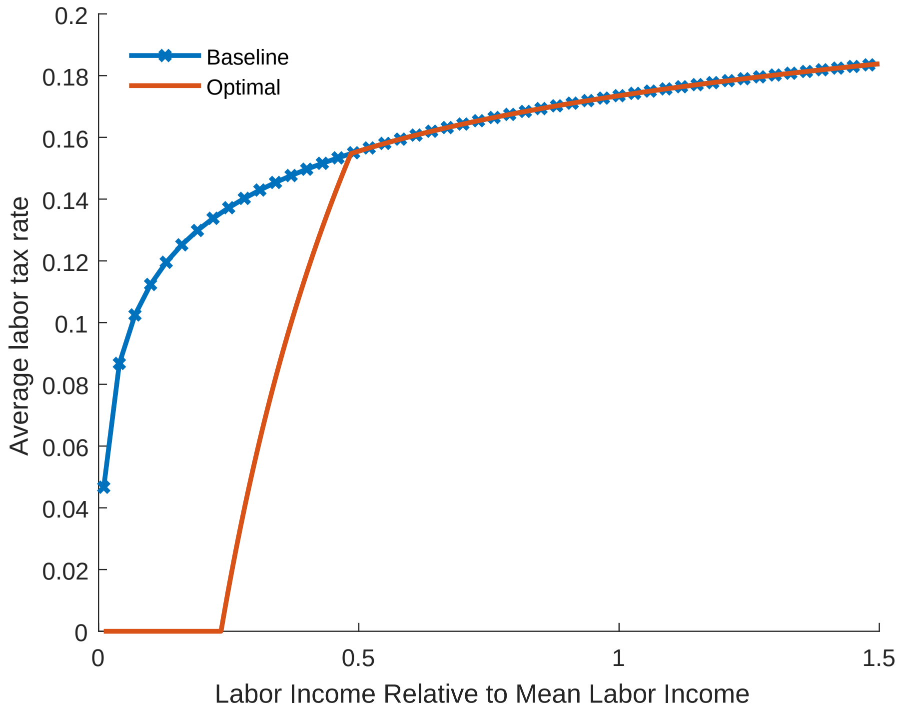

Note: The figure displays the average labor income tax rate paid by an agent under the optimal rebate and in the baseline steady state. The average tax rate is displayed as a function of an agent's labor income relative to the mean level of labor income.

We find that the optimal approach for recycling 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 welfare effects of using each of the five recycling approaches in isolation.[12] The first two columns display the effects of recycling all of the revenue through a uniform reduction in the capital or labor income tax rates while the third column displays the impacts of providing uniform, lump-sum payments. The fourth column displays the impacts of recycling the revenue through payments that are allowed to vary with an agent's total income. Importantly, when searching for the welfare maximizing way to vary these payments with total income, we find that it is optimal to *not* vary the payments with income. That is, the welfare maximizing income-

[12]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.

dependent payments are simply uniform, lump-sum rebates. Therefore, column four instead displays the impact of the unique income-dependent payment approach that achieves the same level of progressivity (i.e. the same percentage change in the Gini of lifetime welfare) as the optimal recycling approach while imposing the smallest possible decrease in expected welfare. Column five finally displays the effect if all of the revenue had to be used solely to increase the progressivity of the labor income tax.

Table 1: Distribution and Welfare Effects

	Capital tax	Labor tax	Lump sum	Income progressivity	Labor progressivity	Optimal
CEV	-0.27	-0.54	-0.64	-0.92	-0.13	-0.11
%Δ welfare Gini	-0.0	0.1	-1.1	-2.4	-3.7	-2.4
%Δ capital	1.8	-1.8	-3.4	-5.7	-3.8	-0.5
%Δ energy	-30.3	-31.4	-32.0	-32.9	-32.4	-31.2
%Δ output	0.6	-0.4	-1.1	-2.0	-1.6	-0.4

To understand how the optimal rebate 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, to agents with lower lifetime welfare, and higher marginal utilities of consumption. Ultimately, the optimal policy combines the two most effective recycling approaches to achieve both of these objectives. First, in terms of using the carbon tax revenue to uniformly reduce an existing tax, the first two columns of Table 1 highlight that reducing the capital income tax – as the optimal policy does – achieves greater expected welfare than uniformly reducing the labor income tax. However, the first two columns of Table 1 also reveal that, in isolation, uniformly reducing the capital or labor income tax would have neutral distributional impacts, with the Gini coefficients of lifetime welfare effectively remaining unchanged. Therefore, to increase welfare through the second channel – i.e. redistributing resources towards agents with lower lifetime welfare and higher marginal utility – at least a portion of the carbon tax revenue will need to be recycled in a progressive way.

Our analysis considers three broad recycling approaches that can achieve a progressive redistribution of the carbon tax revenue. First, policymakers can provide lump-sum payments. While these payments do not vary across households, they will nonetheless provide a relatively larger increase in utility for lower lifetime income households, and consequently, as column three of Table 1 highlights, reduce the Gini coefficient of lifetime welfare. However, far more redistribution can be achieved by targeting the revenue back to lower income agents. In particular, we consider two ways to do so. Consistent with the income-tested approach in several recent policy proposals, the payments can vary inversely with an agent's total income. Alternatively, by increasing the progressivity of the labor income tax, the carbon tax revenue received by agents can instead effectively vary with their labor income. Columns three, four, and five highlight that increasing the progressivity of the labor income tax is superior to providing lump-sum payments or payments that vary with an agent's total income. In particular, across the three progressive recycling approaches, returning the carbon tax revenue through an increase in the progressivity of the labor income tax achieves the greatest increase in equality – the Gini coefficient of lifetime welfare falls by 3.7 percent – while also achieving the highest expected welfare.

To understand why increasing the progressivity of the labor income tax is superior to the other progressive recycling options, it is helpful to consider what an ideal recycling approach would do to achieve an increase in progressivity. In particular, the ideal policy would return the revenue back to agents with the highest marginal utility. Figure 2 highlights that, among the three progressive recycling options, increasing the labor tax progressivity is the most effective at targeting the revenue back to agents with the highest marginal utility. By increasing the progressivity of the labor tax, 85 percent of the revenue is recycled back to the 50 percent of agents with the highest marginal utility. In contrast, by returning the revenue through payments that vary with agents' total income, only 64 percent of the revenue is returned to agents in the top half of the marginal utility distribution.

Intuitively, agents with the highest marginal utility will be those that have low labor income over the entire course of their working lifetime as well as those that just experienced large, adverse shocks to their labor productivity. Consequently, increasing the progressivity of the labor income tax, and reducing the tax liability of working-age agents with low labor earnings, targets the revenue back towards these agents with low labor productivity and high marginal utility. In contrast, recycling the revenue through payments that vary inversely with an agent's total income does not precisely target the carbon tax revenue back to agents with the highest marginal utility. This is due to the fact that payments will be provided to any agents with low total income in that period, working or retired. After retirement, agents will have low total income. However, because they are consuming savings, the low total income does not imply that they have high marginal utility. As a result, targeting the payments back to any agents, working or retired, with low total income does not achieve as large of an

Figure 2: Cumulative Share of Revenue Returned by Marginal Utility

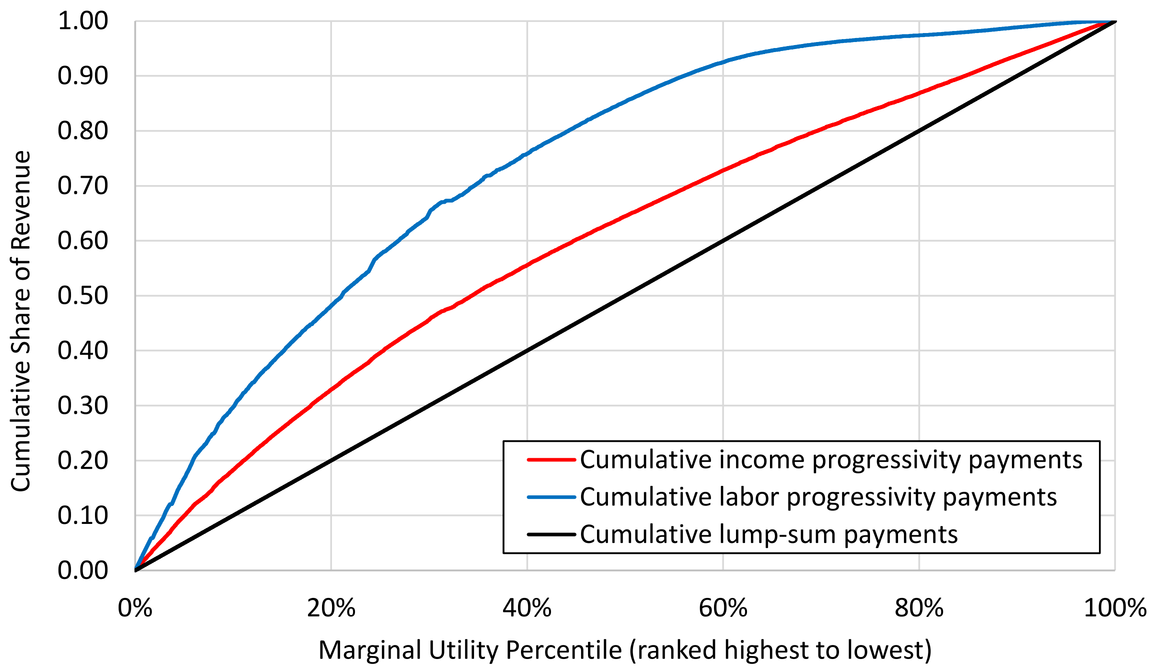

Note: The figure displays the cumulative share of the revenue recycled by agents' marginal utility under the three progressive recyling approaches. The agents are ranked from highest to lowest marginal utility.

increase in equality as would be achieved by increasing the progressivity of the labor tax.

Table 2, which displays the correlation between the amount of revenue an agent receives under the three progressive recycling options and their marginal utility, further highlights how the revenue targeting varies across the policies. Focusing exclusively on working age agents, recycling the revenue through payments that vary inversely with an agent's total income is the most effective at targeting the revenue back to the agents with the highest marginal utility. However, under this recycling approach, 40 percent of the revenue will be returned to retired age agents. Expanding the focus to the full population, the correlation between the payment an individual receives and their marginal utility is the highest when the revenue is recycled through an increase in the progressivity of the labor tax.

Table 2: Relationship between Payments and Marginal Utility

	Lump sum payments	Income progressivity payments	Labor progressivity payments
Correlation with marginal utility			
Working agents	0	0.71	0.59
Retired agents	0	0.65	·
Full population	0	0.41	0.57
Percent of revenue to retired agents	18	40	0

Providing uniform lump-sum payments are of course even less targeted. Indeed, the only reason they achieve an increase in progressivity is due to the fact that, relative to their level of consumption, the uniform payments are relatively larger for lower lifetime income agents. Consequently, as Table 1 highlights, the lump-sum payments achieve a smaller change in the Gini coefficient of lifetime welfare. However, compared to the targeted payments that vary inversely with an agent's total income, the lump-sum payments direct less of the carbon tax revenue back to agents later in their lifecycle, post-retirement. Consequently, as Table 1 highlights, the lump-sum payments don't crowd out as much savings – capital falls by 3.4 percent using the lump-sum payments compared to 5.7 percent when targeting the payments as a function of an agent's total income. As a result, the lump-sum payment approach ends up achieving even higher expected welfare than the targeted approach of providing payments that vary with a household's total income.

To further emphasize the importance of targeting the revenue back towards working-age agents with low labor income, we conduct an additional set of experiments. We first search for a 'restricted' optimal recycling policy in which the policymaker cannot change the progressivity of the labor tax, but she can still increase equality by providing targeted 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 payments.

Similarly, by itself, the ability to target payments 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 recycling policy in which the policymaker cannot increase the progressivity of the labor tax, but she can provide uniform lump-sum payments to working-age agents only. Again, the optimal approach in this case recycles 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 total income that we find it is optimal to combine reductions in the capital tax with equality-increasing, targeted payments. In a setting 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.

5.2 Distributional Impacts of Recycling Approaches

Figure 3: Heterogeneity in Ex-Post Welfare Changes

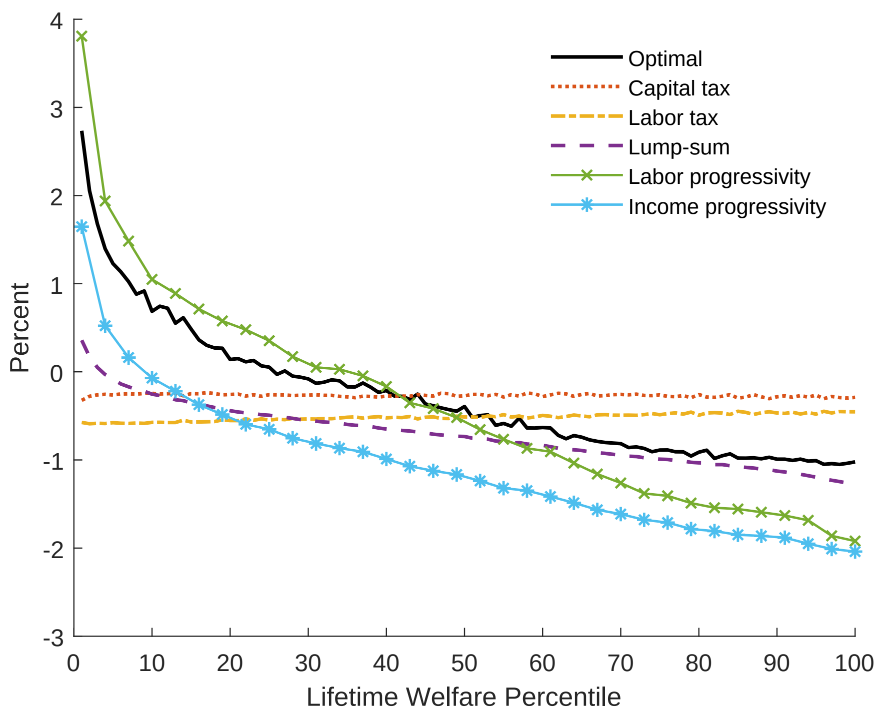

Note: The vertical axis represents the percentage change in baseline lifetime consumption required to make an agent indifferent between living in the baseline steady state and the steady state under a given climate policy. Agents are separated by their lifetime welfare in the baseline steady state, with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

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.

Table 3: Share of Agents Better Off

	Capital tax	Labor tax	Lump sum	Income progressivity	Labor progressivity	Optimal
% better off	1.9	0.0	3.1	9.2	32.6	29.4
% prefer optimal	40.4	50.7	83.8	99.9	60.1	-

Figure 3 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 average impact on agents with the lowest lifetime welfare. Recycling the revenue through lump-sum payments, payments that vary inversely with an agent's total income, or through an increase in the progressivity of the labor tax all result in progressive changes to the tax system; the corresponding lines are downward sloping if Figure 3, implying that agents with the highest lifetime welfare in the baseline experience the largest percentage declines is ex-post welfare under both policies. In contrast, the capital and labor tax rebates 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.

While the slope of the line reveals the degree of progressivity or regressivity of the policy, the height of the line reflects how the policy impacts an agent's welfare, given the value of lifetime welfare in the baseline on the horizontal axis. Focusing on the three progressive recycling options, we can see that recycling the revenue through an increase in the labor tax progressivity effectively dominates the approach of recycling the revenue through payments varying with an agent's total income. Further highlighting the undesirability of the payments varying with total income, all but the lowest lifetime welfare agents fare better off under the lump-sum rebates compared to varying the payments with total income.

Figure 3 reveals that the results are mixed, however, comparing the optimal recycling approach to the capital tax reduction. Agents with baseline welfare above 40th percentiles experience larger welfare losses under the optimal rebate compared to the capital tax reduction (Table 3). In contrast, agents with lower levels of lifetime welfare fare far better under the optimal rebate. Notably, agents in the bottom 30 percent of the lifetime welfare distribution experience welfare gains under the optimal rebate (Table 3). Ultimately, the sizable ex-post welfare gains experienced by agents with the lowest lifetime welfare cause the optimal rebate policy to achieve the highest expected lifetime welfare.

5.3 Sensitivity

To explore the sensitivity of the results, we examine how the optimal rebate differs as we vary the level and regressivity of the carbon tax. To vary the regressivity of the carbon tax, we consider different values of subsistence energy, $\bar{e}$. Table 4 highlights that, across the different carbon tax levels we consider, and for different levels of inherent regressivity, it remains optimal to use 60 to 66 percent of the revenue to reduce the capital tax while the remaining revenue is used to increase the progressivity of the labor tax. Regardless of the specification, the carbon tax paired with this combined rebate reduces the Gini coefficient on lifetime welfare from its value in the baseline, raising equality.

While the optimal rebate is very stable across the alternative specifications, Table 4 reveals small quantitative differences that illustrate how the level and regressivity of the carbon tax affect the relative importance of using the revenue to unwind existing distortions versus increase equality. When the carbon tax is more regressive (i.e. $\bar{e}$ is larger), the relative importance of using the revenue to increase equality increases. In contrast, as the level of carbon tax grows, the relative importance of using the revenue to reduce the distortions caused by the capital tax grows.

We also explore whether the optimal rebate approach is sensitive to the baseline composition of income taxes. Rather than starting from a baseline matching current U.S. income taxes, we start from a baseline using the optimal mix of income taxes. To identify the optimal pre-carbon tax mix of income taxes, we choose the capital tax and the level and progressivity of the labor tax to maximize expected welfare prior to the adoption of a carbon tax, subject to the same government revenue requirement. In the optimal pre-carbon-tax baseline, the capital tax falls from 36 percent to 11.6 percent while the labor tax becomes slightly more progressive with λ_1 falling to 0.793 and λ_2 increasing to 0.121.

Starting from this optimal pre-carbon-tax baseline, we again find that it is optimal to rebate the revenue by both reducing the capital tax rate and by increasing the progressivity of the labor-tax. While the optimal rebate mechanisms are unchanged, the share of carbon tax revenue going towards each mechanism does differ. Because the capital tax rate is far lower in the optimal pre-carbon-tax baseline, the importance of reducing the distortions caused by the capital tax is reduced relative to the importance of increasing equality. Consequently, when starting from the optimal baseline, the share of carbon tax revenue used to reduce the capital tax falls from 64 percent to 40 percent.

Table 4: Sensitivity

	Fraction of revenue used to reduce the capital tax	Percent change in the welfare Gini
Subsistence energy: $\bar{e}$		
$\bar{e} = 0$	0.65	-2.36
$\bar{e} = 0.0013$	0.64	-2.35
$\bar{e} = 0.0026$	0.62	-2.33
Carbon tax: τ^c		
\$30/ton CO_2	0.60	-2.21
\$40/ton CO_2	0.64	-2.35
\$50/ton CO_2	0.66	-2.38

Note: Column 1 displays the fraction of the carbon tax revenue used to reduce the capital tax for different values of subsistence energy, $\bar{e}$, and the carbon tax, τ^c. The remaining revenue is used to increase the progressivity of the labor tax. Column 2 displays the corresponding percent change in the Gini coefficient on lifetime welfare from its value in the baseline. The middle values of $\bar{e}$ and τ^c equal the values from the benchmark calibration.

5.4 Uniqueness of Carbon Tax Revenue

To explore whether there is something unique about a carbon tax that drives the optimal rebate approach identified in the preceding results, we conduct an additional 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 welfare maximizing rebate policy. Starting again from a pre-carbon-tax baseline reflecting our current, sub-optimal mix of income taxes, we search for the optimal way to recycle this new exogenous stream of revenue back to agents.

We find the optimal rebate uses the same two instruments to recycle the revenue. Quantitatively, however, the optimal rebate approach is very different. While we previously found the optimal rebate uses 64 percent of the revenue from a carbon tax to reduce the capital tax, it is optimal to use only 37 percent of the exogenous revenue stream to reduce the capital tax. This result highlights that the way a new stream of revenue is generated can alter the primary mechanism through which the revenue should be rebated.

The quantitative differences between the optimal rebates of carbon tax revenue and the exogenous revenue stem from the fact that the carbon tax itself depresses capital.[13]

[13] If 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.

Intuitively, the carbon tax reduces energy use, which, all else constant, decreases the marginal product of capital, leading to lower aggregate savings. This amplifies the relative importance of reducing the distortions caused by the capital tax when the revenue is generated from a carbon tax. Hence, when the revenue comes from a carbon tax, it is optimal to mitigate the resulting decrease in capital by using substantially more of the revenue to reduce the capital tax.

6 Conclusion

The mechanism through which carbon tax revenue is recycled back to households is an important piece of any revenue-neutral carbon tax policy. This paper solves for the welfare maximizing way to return carbon tax revenue in a general equilibrium lifecycle model of the U.S. economy. 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 many involved in the policy-making process. 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.

Key to our results is the broad set of rebate options we consider and the inclusion of household heterogeneity over age, income and energy expenditure shares. However, there are other sources of heterogeneity that our model does not include, such as heterogeneity over geography or occupation, that could also impact the welfare consequences of different carbon tax rebates. Additionally, our analysis abstracts from the near-term welfare consequences of the carbon tax rebate, focusing instead on comparing outcomes in the long-run steady state. Studying the welfare-maximizing rebate in a setting that accounts for additional sources of heterogeneity and for the near term welfare consequences over the transition is an interesting avenue for future work.

References

Aubert, Diane and Mireille Chiroleu-Assouline, "Environmental tax reform and income distribution with imperfect heterogeneous labour markets," *European Economic Re-*

view, 2019, *116*, 60–82.

Barrage, Lint, "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy," *Review of Economic Studies*, 2019.

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.

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, Rud 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, Nezih, 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.

Jacobs, Bas and Frederick van der Ploeg, "Redistribution and pollution taxes with non-linear Engel curves," *Journal of Environmental Economics and Management*, 2019, *95*, 198–226.

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.

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.

Online 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:

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

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

   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:

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

5. The government budget balances:

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

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:

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

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.[14] Since our

[14]See 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.[15] 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).[16] The average empirical value of E^c/E over the most recent five years of data, 2013-2017, equals 0.183.

[15]The 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.

[16]The 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).

Table A1: Parameter Values

Parameter	**Value**
Persistence: ρ	0.958
Persistent shock variance: σ_κ^2	0.017
i.i.d shock variance: σ_π^2	0.081
Fixed effect variance: σ_ξ^2	0.065
Final Good Capital Share: α_y	0.3
Energy Share: ζ	0.03
Energy Capital Share: α_e	0.597
Depreciation: δ	0.079
Risk Aversion: θ_1	2
Frisch Elasticity: θ_2	0.5
Conditional Discount: β	0.995
Disutility of Labor: χ	73.3
Subsistence Energy: $\bar{e}$	0.0013
Consumption Energy Share: $1-\gamma$	0.0093
Debt Limit: $\underline{a}$	-0.156
Labor Tax Function: λ_1	0.827
Labor Tax Function: λ_2	0.031
Capital Tax Rate: τ^k	0.36
SS Payroll Tax: τ^s	0.096
Government Spending: G	0.106
SS max income: $y^{h,max}$	1.358
SS function bend point: b_1	0.118
SS function bend point: b_2	0.724
SS function bend point: b_3	1.358
SS function marginal benefit: ϕ_1	0.9
SS function marginal benefit: ϕ_2	0.32
SS function marginal benefit: ϕ_3	0.15

Note: This table reports the calibrated parameter values.

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 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.