Optimally Rebating Carbon Tax Revenue (Preliminary)

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

October 20, 2019

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 that the welfare-maximizing revenue-neutral carbon tax policy returns revenue through a reduction in the distortionary capital tax combined with an increase in the progressivity of the labor tax. The optimal policy we identify attains higher welfare and more equality than the lump-sum rebate approach preferred by many policymakers as well as the approach originally prescribed by the economics literature – 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?[1] While numerous studies in the economics literature have 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 and Krueger (2006), Conesa et al. (2009), Heathcote et al. (2017)) and solve for the welfare-maximizing way to recycle carbon tax revenue.[2] 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 distinct mechanisms. First, approximately one third of the carbon tax revenue is used to reduce the distortionary tax on capital income. Second, the remaining two-thirds of the revenue is used to increase the progressivity of the labor income tax. Ultimately, the combination of these two mechanisms results in a progressive change to the tax system.

The optimal revenue recycling approach we identify attains higher welfare and more equality than 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 – an approach which would maximize economic surplus (Parry (1995), Goulder (1995), de Mooij and Bovenberg (1998),

[1] A report from the U.S. Department of the Treasury (Horowitz et al. (2017)) estimates that a carbon tax starting at \$49 per ton of CO_2 in 2019, and rising to \$70 by 2028, would generate \$2.2 trillion over ten years. Similarly, estimates from the U.S. Congressional Budget Office suggest that setting a modest CO_2 price of \$20/ton would raise \$1.2 trillion in revenue during the first decade the policy is in place (CBO (2011)).

[2] Following much of the literature studying alternative approaches for recycling carbon tax revenues, we do not model the environmental benefits from reduced carbon emissions, but rather, we focus specifically on the long-run, non-environmental welfare consequences of revenue-neutral carbon tax policies. While focusing on the long-run impacts abstracts from early transitional dynamics, Fried et al. (2018) highlight that the non-environmental welfare consequences of adopting a carbon tax are similar for young living agents (e.g., 20 years old) and those born into the future steady state.

Bovenberg (1999)). However, we demonstrate that using all of the revenue to reduce the distortionary capital or labor tax results in a regressive outcome, which raises the average welfare cost of the policy. We find that it would be welfare-maximizing to instead use a substantial portion of the revenue to increase equality.[3]

Our optimal revenue 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 the carbon tax revenue to be returned to individuals through equal, lump-sum payments.[4] However, our results demonstrate that providing a lump-sum rebates is not the optimal way to use the carbon tax revenue to obtain a progressive outcome. Policymakers can achieve an even more progressive outcome with far higher welfare by instead using a portion of the revenue to increase the progressivity of the labor 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 an agent's entire life cycle, leading to heterogeneity across age cohorts. In addition, we include idiosyncratic shocks to labor-income, which produce a full income distribution within each age cohort.[5] Finally, our model uses Stone-Geary preferences to capture the fact that low-income agents use a higher fraction of their expenditures for energy, implying that the

[3]Consistent with these previous studies, we focus exclusively on revenue-neutral carbon tax policies. Related research has also considered policies that are not revenue-neutral – e.g., using carbon tax revenue to reduce the federal deficit (Carbone et al. (2013)).

[4]The 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)).

[5]Related work has also examined the welfare impacts of carbon tax policies using life cycle models and 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 instead searches over a continuum of potential rebate options to identify the welfare-maximizing policy.

carbon tax by itself is regressive.[6]

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, exclusively through a reduction in the capital tax rate, or exclusively through a reduction in the labor tax rate.[7] 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 rate, (ii) reduce the level of the labor income tax rate, (iii) increase the progressivity of the labor income tax, and (iv) provide lump-sum rebates 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.

Ultimately, both of these modeling contributions – the rich heterogeneity as well as the rich set of policy instruments – are crucial to identifying the welfare-maximizing rebate approach. To highlight the importance of incorporating heterogeneity, we decompose the resulting welfare changes from each potential policy into two components: (i) welfare changes stemming solely from the policy's impacts on the average level of consumption and hours worked, and (ii) welfare impacts stemming from changes in the distributions of consumption and hours across different age and income groups. If we only focus on the welfare impacts stemming from changes in the average levels of consumption and hours worked, then we

[6]Previous partial equilibrium studies highlight that, because low income households typically devote a larger share of their total expenditures on energy, the direct burden of a carbon tax alone is likely to be regressive (e.g., Metcalf (1999), Grainger and Kolstad (2010)). Recent work by Cronin et al. (2019) find 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.

[7]In one notable exception, Goulder et al. (2019) consider combinations of lump-sum rebates and reductions in federal taxes in an infinitely lived agent model. Ultimately, however, our analysis demonstrates that (a) including heterogeneity over the life-cycle and (b) allowing for changes in the progressivity of the labor tax are vital for uncovering the optimal policy.

find that returning carbon tax revenue exclusively through a reduction in the capital tax would maximize welfare. However, when we consider the welfare changes stemming from both components, we find that it is instead welfare-maximizing to divert a sizable portion of revenue away from reducing the capital tax to increase the progressivity of the labor tax.

The rich set of rebate options we consider is crucial to our result that the welfare-maximizing policy is also progressive. If, as in the previous literature, the policymaker can only to use lump-sum rebates to mitigate the regressivity of the carbon tax, then we find that it is optimal settle for a regressive outcome and return all of the revenue through a reduction in the capital tax rate. A key reason why our optimal policy differs from this standard result is that the lump-sum rebate is a very costly way to increase equality. By providing lump-sum rebates to agents of all ages, the rebate reduces an agent's need to save for retirement, crowding out capital. While the outcome is more progressive, expected welfare for agents of all income levels is lower when the carbon-tax revenue is used to provide lump-sum rebates as opposed to exclusively reducing the capital tax.

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.[8] The macro public finance literature has long utilized general equilibrium, lifecycle models with rich within-cohort heterogeneity to quantify the welfare and distributional effects of alternative tax policies. While the macro literature has focused on which taxes to increase to achieve a given revenue target, we instead focus on which taxes to decrease, given a new stream of revenue from a carbon tax, in order 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.

The remainder of the paper proceeds as follows. Section 2 presents the model and Section 3 discusses how the model is calibrated to reflect the heterogeneity in the U.S. economy.

[8]Beginning with the original 'Double Dividend' studies (Terkla (1984), Pearce (1991), and Repetto et al. (1992)), the environmental and public economics literature has consistently compared the efficiency and welfare consequences of alternative approaches for recycling revenue from corrective environmental taxes in an economy with distortionary taxes.

Section 4 describes the set of revenue recycling options we study as well as the welfare measures we use to evaluate the policies. Section 5 summarizes the steady state welfare and distributional impacts of the continuum of revenue-neutral carbon tax policies. We first identify the welfare maximizing policy and then, through a welfare decomposition, highlight why the optimal policy is preferred to a range of alternatives. Finally, Section 6 concludes.

2 Model

2.1 Demographics

Our model incorporates many 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 at exogenous rate, n. Agents make labor-supply and savings decisions each period until they are forced to retire at a real-world age of 65. After retirement, agents finance consumption from Social Security payments and accumulated assets. Lifetime length is uncertain and mortality risk varies over the lifetime.[9] Since individuals are not certain how long they will live, they may die with positive asset holdings. In this case, we treat the assets as accidental bequests and redistribute them lump-sum across all living individuals during period t in the form of transfers, T_t^a.

2.2 Agents

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

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

[9]We impose a maximum age limit of 100. That is, conditional on surviving to 100, an agent exits the model after the period.

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. Parameter θ_1 is the coefficient of relative risk aversion and parameter θ_2 is the Frisch elasticity of labor supply. Parameter χ determines the dis-utility of hours.

The composite good is comprised of two goods – a generic consumption good as well as carbon-emitting energy. Including these two consumption categories allows us to capture the fact that energy is not only used in the production of consumption goods, but it is also consumed directly by agents (e.g., electricity and gasoline). Importantly, previous work highlights that the share of expenditures that goes towards purchasing energy differs systematically across agents – with lower income groups devoting a larger share of their budgets to energy (Metcalf (2007), Mathur and Metcalf (2009)). To ensure that our model captures this negative relationship between income and energy budget shares, we assume that all agents must consume a minimum amount of energy, $\bar{e}$, and that agents derive no utility from the energy consumed up to this subsistence level. In particular, the composite consumption good 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 consumption good and energy consumed by agent i at age j in period t.

Agents are endowed with one unit of time each period which they divide between labor and leisure. To generate a realistic distribution of income within and across age cohorts, we allow labor productivity to vary across agents. 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 market wage-rate during period t, $h_{i,j,t}$ 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,

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

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

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

The different ability types, as well as the initial realization of the transitory shock, generate an initial income distribution within the cohort of 20-year-old entrants to the model. The different realizations of the persistent shock over the lifetime cause the within-cohort variation to grow with age. Finally, the age-specific human capital generates variation in the average labor productivity of agents of different ages.

Agents can save by accumulating shares of physical capital, $a_{i,j,t+1}$, which they rent to firms in either sector at rate R_t. Physical capital accumulates according to the standard law of motion,

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

where parameter 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. Importantly, both precautionary and lifecycle factors motivate agents to save. Saving allows agents to partially self-insure against idiosyncratic productivity shocks and to finance consumption during retirement.[10] Agents are not allowed to borrow – we require that asset holdings are always weakly positive, $a_{i,j,t} \geq 0$.

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

[10] We assume that agents cannot insure against idiosyncratic productivity shocks by trading explicit insurance contracts and their are no annuity markets to insure against mortality risk.

Cobb-Douglas between the three inputs,

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

Parameters α_y and ζ denote capital share and energy share, respectively. Parameter A^y denotes total factor productivity. 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, N^e, according to the Cobb-Douglas 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 performs two tasks – it runs a balanced-budget, pay-as-you-go Social Security system and it generates revenue in order to finance an exogenous level of unproductive government spending, G. The Social Security system is financed with a flat tax on labor income. We model the Social Security benefits to mimic the U.S. Social Security system. In practice, the Social Security benefits provided to retired agents are a convex, piecewise linear function of each agents' average labor earnings over their highest 35 years of earnings. Rather than including an agent's past labor earnings as an additional state variable, we instead 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 $(\xi, \nu_6 5)$ receives during each period of retirement is determined using a convex, piecewise-linear function of $x(\xi, \nu_{65})$ with marginal benefit rates given by:

$$\begin{aligned} &\tau_1 \text{ for } 0 \leq x < b_1 \\ &\tau_2 \text{ for } b_1 \leq x < b_2 \\ &\tau_3 \text{ for } b_2 \leq x < b_3. \end{aligned} \tag{6}$$

To finance the required spending level, G, the government can use three different instruments: a tax on capital income, a tax on labor income, and once a climate policy is in place, a tax on carbon emissions.[11] The government taxes each 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 any assets she receives as accidental bequests, $y^k_{i,j,t} \equiv r_t(a_{i,j,t} + T^a_t)$.

The government taxes labor income according to a progressive tax schedule. A working 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 under U.S. tax law. Thus, $\tilde{y}^h_{i,j,t} \equiv y^h_{i,j,t} - T^s(y^h_{i,j,t})/2$ is the agent's taxable labor income, where $T^s(y^h_{i,j,t})/2$ is the employer's Social Security contribution. Following the quantitative public finance literature (Benabou (2002), Guner et al. (2014), and 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.[12]

The labor tax function specified by equation (7) allows us to alter the labor tax in a variety of ways following the introduction of a tax on carbon. In particular, parameter λ_1 determines the level of the labor tax and parameter λ_2 governs the curvature the tax function. A decrease in λ_1 decreases the after-tax labor-income of all individuals by the

[11] Consistent with the fact that the U.S. Federal Government does not levy a consumption tax, we do not include a tax on the final consumption good in the set of available tax instruments.

[12] This function form allows for the possibility of a negative average income tax rates. Since we do not observe negative income taxes in the US economy, we bound the tax function below by zero. This constraint rarely binds in our computational experiments.

same percentage – it does not impact on the distribution of after-tax income across agents.[13] In contrast, changing λ_2 does affect the distribution of after-tax labor income. Increases in λ_2 reduce the average tax rate for low-income households and increase the average tax rate for high-income households, reducing the inequality in the distribution of after-tax labor income.

Finally, with the introduction of a climate policy, the government not only finances the exogenous level of spending, G, using the capital and labor taxes, but also with the new carbon tax. The carbon tax, τ^c, is levied on each unit of carbon-emitting energy consumed.[14] Using our model, we compare long-run, steady-state outcomes across a range of different revenue-neutral carbon tax policies. The stationary competitive equilibrium of the model, in which factor prices and aggregate macroeconomic variables are constant, is defined in Appendix A.1.The following section describes how we calibrate the model.

3 Calibration

We calibrate the model to reflect the U.S. economy using two steps. First, we choose one set of parameters directly from the data and existing literature. Second, given these fixed parameters, we choose the remaining parameters so that a set of moments in the model match their empirical values in the data. Table 1 reports the main parameter values and their source. We use a five year average from 2013-2017 for all parameter values and targets that we calculate directly from the data.

[13]Specifically, the Gini coefficient calculated from after-tax labor income is independent of the value of λ_1, as long as the lower bound of zero does not bind for a substantial portion of the income distribution.

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

Table 1: Parameter Values

Parameter	Value	Source
Demographics		
Surv. Prob: Ψ_j	Bell and Miller (2002)	Data
Labor productivity		
Persistence: ρ	0.858	Kaplan (2012)
Persistent shock variance: σ_κ^2	0.017	Kaplan (2012)
i.i.d shock variance: σ_π^2	0.081	Kaplan (2012)
Firm Parameters		
Final Good Capital Share: α_y	0.3	Golosov et al. (2014)
Energy Share: ζ	0.04	Golosov et al. (2014)
Energy Capital Share: α_e	0.597	Barrage (2018)
Final Good Productivity: A^y	1	Normalization
Energy Productivity: A^e	1	Normalization
Depreciation: δ	0.083	Method of moments
Preference Parameters		
Risk Aversion: θ_1	2	Conesa et al. (2009)
Frisch Elasticity: θ_2	0.5	Kaplan (2012)
Conditional Discount: β	1.0002	Method of moments
Disutility of Labor: χ	74.9	Method of moments
Subsistence Energy: $\bar{e}$	0.00474	Method of moments
Consumption Energy Share: $1-\gamma$	0.069	Method of moments
Government Parameters		
Labor Tax Function: λ_1	0.813	Clears market
Labor Tax Function: λ_2	0.031	Guner et al. (2014)
Capital Tax Rate: τ^k	0.36	Trabandt and Uhlig (2011)
Replacement rate: η	0.5	Conesa and Krueger (2006)
SS Payroll Tax: τ^s	0.1096	Clears market
Government Spending: G	0.12	Method of moments
SS max income: $y^{h,max}$	0.96	Method of moments

3.1 Production

We normalize the total factor productivity in both 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 0.04. We choose the depreciation rate on capital, δ, to match the investment to output ratio of 23.3 percent.[15]

3.2 Preferences

We choose the discount rate $\beta = 1.0002$ to match the U.S. capital-output ratio of 2.6.[16] We choose the disutility of labor, $\chi = 74.9$, such that individuals spend an average of one third of their time endowments working. Following Conesa et al. (2009), we set the coefficient of relative risk aversion, θ_1, equal to 2 and, consistent with Kaplan (2012), we set the Frisch elasticity, θ_2, equal to 0.5. We choose the conditional survival probabilities based on the estimates in Bell and Miller (2002).

The energy subsistence parameter, $\bar{e}$, governs how an agent's energy budget share changes with income. Following Fried et al. (2018), we choose $\bar{e} = 0.00474$ to target an energy-share difference between the top and bottom halves of the expenditure distribution of 7.44 percent, based on data from the CEX (see Appendix A). 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 $\gamma = 0.931$ to match the empirical ratio of energy consumed directly by households relative to total energy consumption, 0.183.

3.3 Idiosyncratic Labor Productivity

We take the parameters of the idiosyncratic labor productivity processes, $\{\sigma^2_\xi, \sigma^2_\kappa, \sigma^2_\pi\}$, from the estimates in Kaplan (2012).[17] The values for the age-specific human capital parameters are taken from Kaplan (2012) and Huggett and Parra (2010).[18] Importantly, the annual

[15]Data are from Tables 1.1 and 1.1.5. We define investment as investment in private fixed assets plus investment in consumer durables.

[16]Data are from Tables 1.1 and 1.5. We define capital as the sum of private fixed assets and consumer durables.

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

[18]The parameter 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

variation in labor income that Kaplan (2012) uses to estimate the shock processes includes heads of households who have worked as little as 520 hours – or equivalently, 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 in the US.[19]

3.4 Government Policy

Government expenditure, G, is set to ensure that it equals 15.7 percent of output.[20] We set the tax rate on capital income, τ^k, to 36 percent based on estimates in Kaplan (2012), Nakajima (2010) and Trabandt and Uhlig (2011).

We set the social security marginal benefit rates, τ_1, τ_2, τ_3, to match the progressive, piecewise-linear benefit function used in the actual U.S. Social Security system. To determine the knot points, b_1, b_2, b_3, in the progressive benefit function, we set the ratio of the knot point to the model average labor earnings equal to the ratio of the corresponding ratio of the actual knot point and the average labor earnings in the data.

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 0.813 in order to clear the government budget constraint. These parameters imply that an agent with the mean labor income faces an average labor-tax rate of approximately 18.7 percent and a marginal labor-tax rate of 21.2 percent.

In contrast to previous work focusing on the optimal level for a carbon tax, we set the tax on carbon emissions at a fixed level. This allows us to focus exclusively on the welfare consequences of alternative approaches for rebating the resulting carbon tax revenue. We

quadratic polynomial.

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

[20] We use data on government budget outlays from the CBO: https://www.cbo.gov/about/products/budget-economic-data. Since our model includes Social Security separate from government spending, we calculate government spending as the difference between total government outlays and Social Security outlays.

analyze a carbon tax set at \$40 dollars per ton of CO_2 – 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.[21] This process implies that a \$40 per ton carbon tax equals 49 percent of our composite fossil energy price.

Table 2 reports the value of the targets in the model and in the data.

Table 2: Model Fit

Moment	Model	Target
Capital-output ratio	2.6	2.6
Investment-output ratio	0.233	0.233
Government spending-output ratio	0.157	0.157
Household energy consumption relative to total	0.183	0.183
Percent difference in energy budget share	7.44	7.44

4 Computational Experiments

Using the model, we study the long-run welfare effects of policies that combine the carbon tax with one or more rebate instruments to return all of the tax revenue back to agents. Our approach is to compare the welfare and distributional outcomes in a counterfactual steady state with a carbon tax policy in place to the outcome in the baseline steady state with no carbon tax policy. In each counterfactual simulation, we adjust the Social Security benefits so that the purchasing power is unchanged from the baseline (Goulder et al. 2019).

[21] Data on the carbon intensity, energy prices, and energy consumption are from the EIA.

4.1 Rebate Policies

We allow the policymaker to rebate the carbon tax revenue through direct lump-sum transfers to agents and by decreasing existing federal labor and capital tax rates. Since we are focused on ways to return the carbon-tax revenue to agents, and not raise additional revenue, we do not allow the policymaker to increase 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 rate, τ^k, (ii) a reduction in the level of the labor-tax rate, λ_1, (iii) an increase in the progressivity of the labor-tax rate, λ_2, and (iv) an income-dependent lump-sum transfer.

The increase in the progressivity of the labor-tax schedule is designed to mimic a change in the U.S. tax code in which the government reduces in the average labor-income tax rate for the lower-income agents while not changing 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 simultaneously increases the average labor tax-rate for high-income agents. This change alone would not constitute a pure rebate because the tax rate increases for a fraction of the population. Therefore, we require that the average labor tax rate cannot increase from the baseline 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_2' is the value of the curvature parameter in the counterfactual simulation.

In contrast to much of the previous literature, we do not require lump-sum rebates to be constant across all agents. Instead, we allow the lump-sum payments to vary linearly with an agent's total income (capital plus labor income). Specifically, the size of the lump-sum rebate for agent i, age j, with income y_{ij} is

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

where parameter Υ_1 is the intercept for the lump-sum rebate function and parameter Υ_2 is the slope. If $\Upsilon_2 = 0$, then equation (8) generates an equal lump-sum rebate across all agents. We are particularly interested in regions of the parameter space in which $\Upsilon_2 < 0$, so that the size of the rebate decreases with an agent's income. Again, we bound the lump-sum rebate function below by zero to avoid raising taxes on any agent.

4.2 Welfare and Distributional Metrics

Following the tradition in the macro literature, we use the consumption equivalent variation (CEV) to measure the aggregate welfare impacts of each carbon tax policy. Importantly, this welfare measure is ex-ante in that it depends on the agent's expected lifetime consumption before any information about the agent is revealed. Specifically, the CEV measures the uniform percentage change in an agent's expected 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 policy.

Following Peterman and Sager (2018), we decompose the aggregate welfare impact of each policy into three components: the level, age, and distribution components. The level component measures changes in aggregate welfare that result from changes in the aggregate, economy-wide values of hours worked and consumption of the generic and energy good. For example, suppose the carbon-tax policy increases the size of the economy, leading to more economy-wide consumption. The effect of the increase in consumption on aggregate welfare is measured by the level component.

The age component refers to changes in aggregate welfare that result from changes in the shapes of the lifecycle profiles of average consumption and hours worked, holding the aggregate level of consumption and hours constant. For example, suppose that the carbon-tax policy flattens the consumption profile across ages, implying that an individual's consumption is allocated more evenly over her lifetime. The welfare benefit from the flatter consumption-profile, holding the aggregate level of consumption constant, is captured by the age component.

The distribution component refers to changes in aggregate welfare that result from changes in the distribution of consumption and hours across agents. For example, suppose that the carbon tax policy increases the variance in consumption. The welfare cost of this increase in variance is the distribution component. Appendix A.4 reports the formula for each CEV component.

Finally, to further explore the distributional impacts of each policy option, we compute the following additional metrics. First, we calculate the CEV conditional on agents being a specific income quintile. In addition, following Fried et al. (2018), we compute the Gini coefficient for lifetime welfare to measure the level of equality under each policy. We define the Gini coefficient, $\mathcal{G}$, as

$$\mathcal{G} = \frac{\sum_{i=1}^{N}\sum_{j=1}^{N}|x_i - x_j|}{2N^2\bar{x}}, \tag{9}$$

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 and one, with zero implying perfect equality and one implying perfect inequality.

5 Quantitative Results

5.1 Optimal Policy

To search for the optimal policy, we calculated steady states over a grid of different tax policies. Specifically, for each combination of the capital tax rate (τ^k), the level of the labor tax rate (λ_1), the progressivity of the labor tax (λ_2), the intercept of the lump-sum rebate function (Υ_1), and the slope of the lump-sum rebate function (Υ_2), we solved for the resulting steady state outcome. In each policy simulation, we ensured the government budget constraint cleared by adjusting either the level of the labor or capital tax. Finally, we eliminated all steady states in which the capital or labor tax was higher than the pre-tax steady state.

Searching over parameter combinations, we find that the rebate policy that maximizes the expected welfare of an agent in the steady state uses approximately 37 percent of the revenue to reduce pre-existing distortions in the economy, through a reduction in the capital tax rate. The remaining 63 percent of the revenue is used to increase equality, through an increase in the progressivity of the labor-tax rate. Under the optimal rebate, the capital-tax rate equals 32 percent, 4 percentage points below its baseline value of 36 percent. The curvature parameter on the labor-tax function increases to 0.107, from its' baseline value of 0.031.

To interpret the increase in the labor-tax progressivity under the optimal rebate, Figure 1 plots the percent of labor-income taxes returned to an agent as a function of her labor income. Agents with labor income greater than the mean do not receive a rebate. If an agent has very low income, then she receives a rebate equal to 100 percent of her tax payments; her effective labor-income tax rate is zero. The ability of the increase in progressivity to transfer resources to low-income agents is limited by the level of labor-income taxes that they would have paid in the baseline. This constraint binds for those with very low-paying jobs as well as for people who were unemployed for part of the year, and as a result, have low annual labor income.

Figure 1: Rebate From Increase in Labor Tax Progressivity at the Optimum

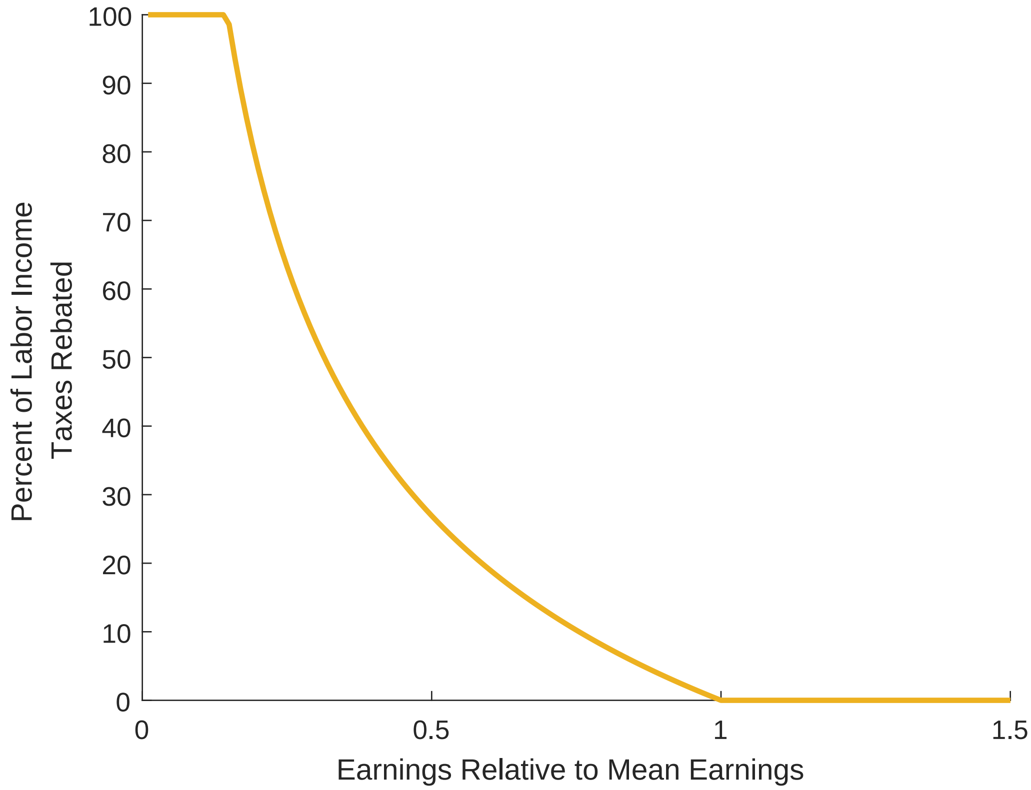

For the sake of comparison, we also quantify the welfare and distributional impacts of the three standard rebate approaches typically considered in the literature: exclusively providing uniform lump-sum rebates, exclusively reducing the capital tax rate, or exclusively reducing the level of the labor tax. Table 3 reports the CEV and percentage change in the Gini coefficient, relative the baseline steady state, under each rebate approach.

Table 3: Distribution and Welfare Effects

	Lump sum rebate	Capital tax rebate	Labor tax rebate	Optimal rebate
Change in Gini (percent)	-1.39	0.18	0.36	-2.93
CEV (percent)	-1.04	-0.41	-0.68	-0.29

Focusing first on change in expected welfare, we see that the CEV is always negative, implying that all four rebates reduce non-environmental welfare relative to the baseline. It is important to stress that these welfare changes do not incorporate any benefits that accrue through improved environmental quality – which ultimately is the justification for implementing the carbon tax in the first place. The change in energy use, and thus the resulting environmental benefits, are stable across each of the policy simulations. Thus, the relative differences between the non-environmental CEV represents the welfare differences across the alternative rebate options.

The optimal policy differs from the results in the previous literature on two key dimensions. First, the double-dividend literature consistently finds that it is optimal to use the carbon-tax revenue to exclusively reduce pre-existing distortionary taxes. Indeed, as Table 3 highlights, this result holds if the policymaker is restricted to the three simple policies. Among these three policies, welfare is highest when the carbon-tax revenue is used to reduce the capital-tax rate. However, when the policymaker can use the full continuum of rebate instruments, then it is only optimal to use approximately 37 percent of the revenue to reduce the capital-tax rate. The remaining 63 percent of the revenue is used to increase the progressivity of the labor tax. Importantly, the increase in the labor-tax progressivity actually increases the distortion on hours. The higher curvature parameter implies that the labor-tax rate rises more steeply with labor income, directly affecting individuals' labor-supply decisions.[22]

Second, the optimal policy is progressive, resulting in greater equality relative to the baseline steady state. In contrast, if policymakers are restricted to the three simple instruments, then the welfare-maximizing instrument, the capital-tax rebate, is regressive, instead of progressive. The more costly lump-sum rebate is the only instrument that increases equality. As a result, policymakers are forced to choose between a welfare-maximizing policy (capital-tax rebate) and a progressive policy (lump-sum rebate). However, when policymakers are not restricted to the three simple instruments, then this artificial trade-off between welfare-maximizing and progressive policies disappears; the optimal policy maximizes welfare and

[22]The steeper labor-tax function distorts the individual's intratemporal decision between hours and leisure and the intertemporal decision between hours today versus hours tomorrow.

increases equality.

Driven by concerns over equity, many policies and policy proposals use a lump-sum rebate to return the carbon-tax revenue to individuals. However, more equality does not necessarily imply higher welfare. Indeed, the welfare cost of the uniform lump-sum rebate is so high that even the lowest income agents are better off under the capital-tax rebate. Table 4 reports the CEV conditional on the agent's income quintile; quintile one corresponds to the poorest agents. The welfare cost for the poorest agents under the capital-tax rebate rebate is only -0.45, compared to -0.51 under the lump-sum rebate. Thus, even the low-income agents would prefer the regressive, but less costly, capital-tax rebate to the progressive, but more costly, lump-sum rebate.

Table 4: CEV by Quintile (percent)

	Lump sum rebate	Capital tax rebate	Labor tax rebate	Optimal rebate
Quintile 1	-0.76	-0.45	-0.76	0.34
Quintile 2	-0.79	-0.45	-0.74	0.24
Quintile 3	-1.05	-0.41	-0.68	-0.30
Quintile 4	-1.42	-0.37	-0.57	-1.09
Quintile 5	-1.44	-0.35	-0.58	-1.22

In sum, the optimal rebate stands in stark contrast to the results from the double dividend literature and to many policy proposals. These differences are driven by two separate channels. First, the policy's effect on equality is intimately linked to its' welfare cost. Much of the double dividend literature abstracts from income heterogenity, and thus does not model the welfare consequences of changes in equality. Second, the lifecycle features of our model demonstrate that it is much less costly to offset the regressivity of the carbon tax by increasing the progressivity of the labor tax, as oposed to using lump-sum rebates. Sections 5.2 and 5.3 explore each of these channels in turn.

5.2 Relationship between Welfare and Equality

To highlight the relationship between welfare and equality, Figure 2 plots the welfare-equality frontier. Specifically, the blue line plots the CEV for the carbon-tax policy that maximizes welfare and achieves the level of equality on the horizontal axis.[23] The optimal policy is the policy with the highest welfare, denoted with the gold star in the top middle of Figure 2. For reference, the purple open circle, green diamond, and red triangle plot the pure capital-tax, labor-tax, and lump-sum rebate policies, respectively.

Figure 2: Welfare-Equality Frontier

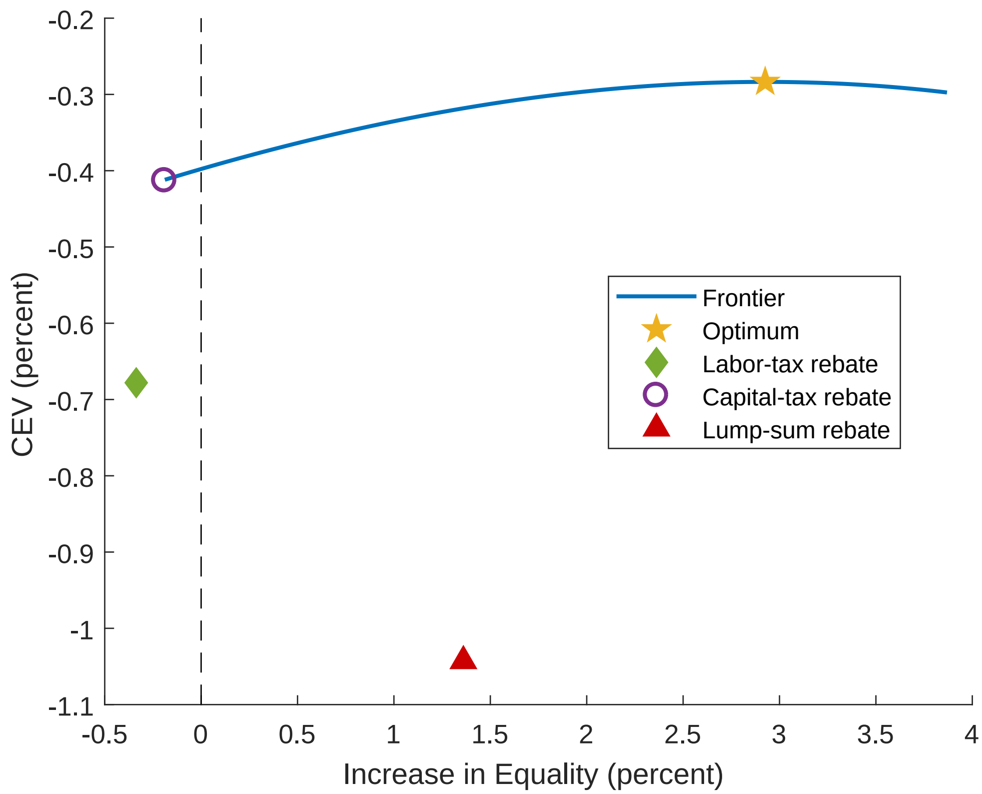

The two panels of Figure 3 plot the combination of rebate instruments that achieve

[23]The grid search we conduct identifies discrete policy combinations that reside along the welfare-equality frontier. The frontiers displayed throughout the remainder of the paper fit a smooth polynomial to the discrete, simulated data points in order to make a continuous frontier. To highlight the fit of the smoothed polynomial, Figure A1 displays the discrete simulated points and the corresponding smoothed polynomial frontier.

the frontier policies. Importantly, none of the frontier policies use carbon tax revenue to reduce the level of the labor tax or to provide lump-sum rebates, uniform or otherwise. The left-most point on the frontier represent the pure capital tax rebate policy. Moving from left to right along the frontier, the policymaker decreases the amount of revenue she returns to agents through reductions in the capital tax (the capital-tax rate rises from 0.21 to 0.36) and increases the amount of revenue returned through steadily larger increases in the progressivity of the labor tax (the curvature parameter rises from 0.031 to 0.014).

Figure 3: Tax Policies Along the Frontier

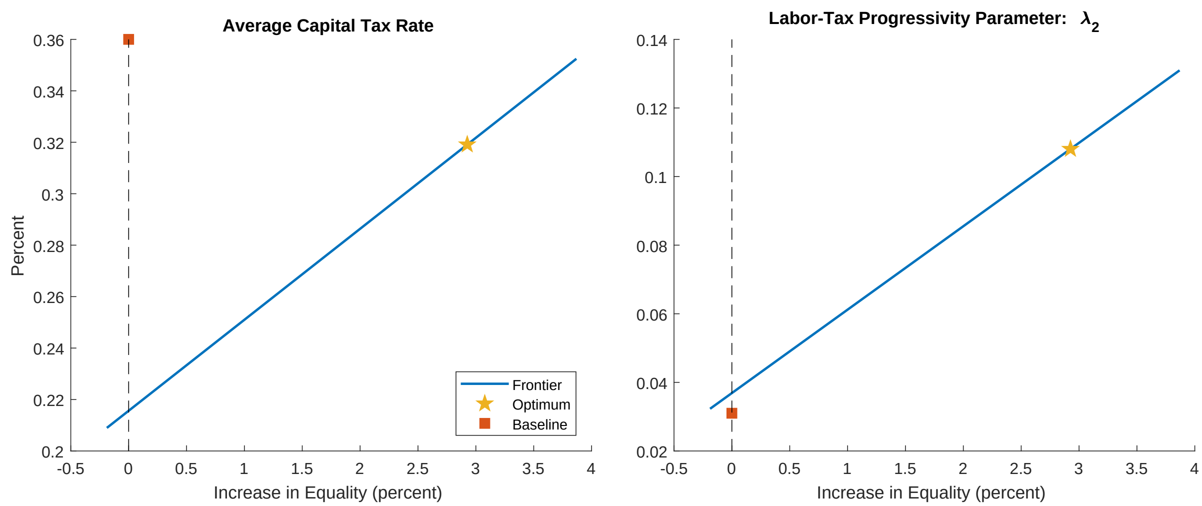

The welfare-equality frontier displayed in Figure 2 is hump-shaped. Increases in equality are accompanied by increases in welfare at low levels of equality and decreases in welfare at high levels of equality. This hump-shape results from the competing effects that changes in average consumption and changes in equality have on welfare. To quantify these competing effects, we decompose the CEV into the level, age, and distribution components. The level component captures the welfare consequences of changes in average consumption and hours. The age and distribution components measure the welfare consequences of changes in equality across age and income groups. Figure 4 plots the level (purple dashed line) age (red dashed and dotted line) and distribution (green dotted line) components of the total CEV for each policy along the frontier. For comparison, the blue solid line plots the welfare-equality frontier which is equal to the total CEV. The gold stars denote the value of each CEV component at the optimum.

Figure 4: CEV Decomposition

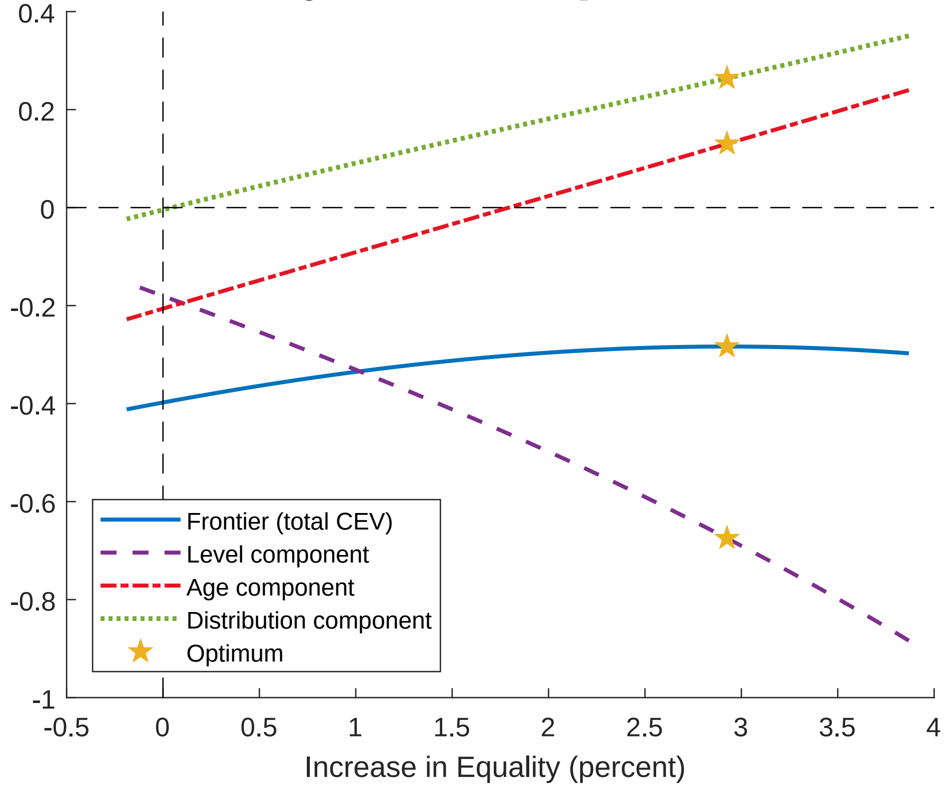

The level component is downward-sloping along the frontier, implying that the welfare cost of changes in average consumption and hours increase as equality increases. Figure 5 plots the percent change in capital, output, consumption and hours at each point on the frontier. Capital falls as one moves from left to right along the frontier because the capital-tax rate increases. Consistent with the decline in capital, hours, output, and consumption all decrease over the frontier. The welfare cost of the fall in consumption dominates the welfare gain from the fall in hours, causing the level component to slope downwards.

The level component measures the welfare costs of the policy if the policymaker does not account for the welfare consequences of changes in equality across age or income groups. This metric is the closest analog in our model to changes in economic efficiency, the metric used by much of the double dividend literature to compare the three simple rebate options. The downward-sloping level component reveals that if one ignores the effects of distributional changes on welfare, then the optimal policy is the simple capital-tax rebate. Thus, if the policymaker only focuses on the level component, then it is optimal to use all of the carbon-tax revenue to reduce pre-existing distortionary taxes, consistent with the findings from the double-dividend literature. Furthermore, increases in equality are accompanied by a decrease in the level component, restoring the trade-off between welfare and equality.

Figure 5: Capital and Output Along the Frontier

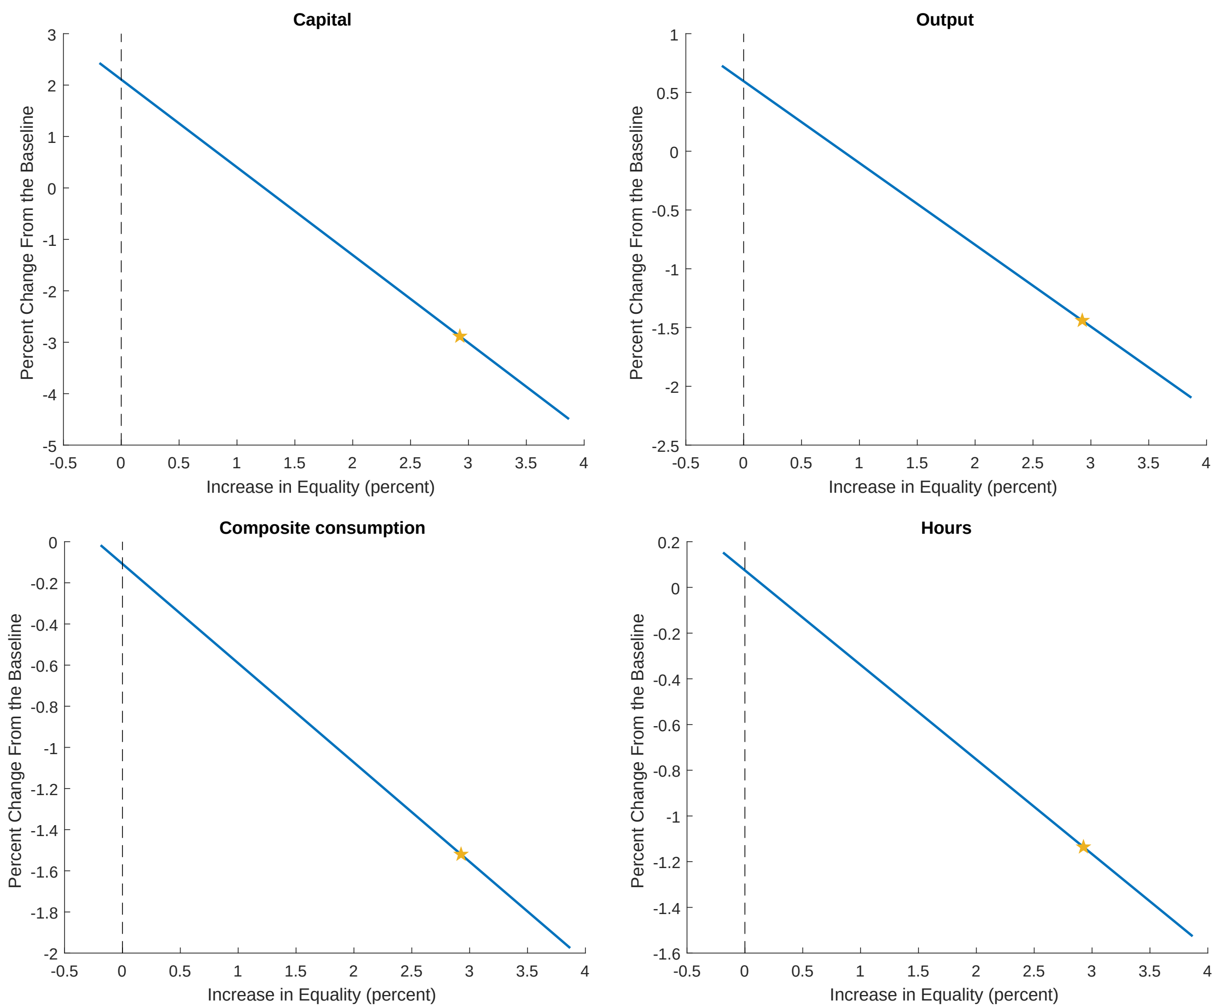

While the simple capital-tax rebate maximizes the level component, the welfare cost of the capital tax rebate from its effects on age and income heterogeneity are substantial. Along the welfare-equality frontier, both the age and distribution components are lowest under the capital tax rebate. The age and distribution components increase as one moves from left to right along the frontier. The distribution component rises over the frontier because an increase in equality reallocates consumption from high-income agents, with low marginal utilities of consumption, to low-income agents, with high marginal utilities of consumption, raising ex-ante welfare. The age component rises because the after-tax interest rate falls as the capital tax rate increases along the frontier. The fall in the after-tax interest rate leads agents to shift consumption to earlier in life, effectively flattening their consumption profiles. A flatter consumption profile brings the individual closer to perfectly equal consumption over the lifecycle, raising welfare.

The capital tax rebate is inferior to the optimal policy at the peak of the frontier because the welfare costs from low distribution and age components dominate the welfare gains from the high level component. These distributional costs imply that it is optimal for the policymaker to use some of the carbon-tax revenue to increase equality, at the expense of a lower level component. Importantly, it is not optimal for the policymaker to use all of the revenue to increase equality. To the right of the optimum, the decrease in the level component dominates the increases in the age and distribution components, causing total welfare to fall. Thus, it is critical to allow the policymaker to use a combination of rebate instruments to attain the welfare-maximizing outcome.

5.3 Increase in Labor Tax Progressivity vs. Lump-Sum Rebate

The characteristics of the optimal policy hinge on the ability of the policymaker to rebate the carbon-tax revenue through an increase in labor-tax progressivity. To demonstrate the importance of this rebate option, we calculate a restricted welfare-equality frontier in which the policymaker cannot change the progressivity of the labor tax. Along this restricted frontier, the lump-sum rebate function is the policymaker's only tool to increase equality. The dashed green line in Figure 6 plots the restricted welfare-equality frontier when the policymaker cannot change the progressivity of the labor tax. For comparison, the solid blue line plots the unrestricted frontier from Figure 2 when the policymaker can use the full set of rebate instruments. For computational tractability, we restrict the range of the restricted frontiers to go from the capital-tax rebate to the optimal policy from Figure 2.

Figure 6: Restricted and Unrestricted Welfare-Equality Frontiers

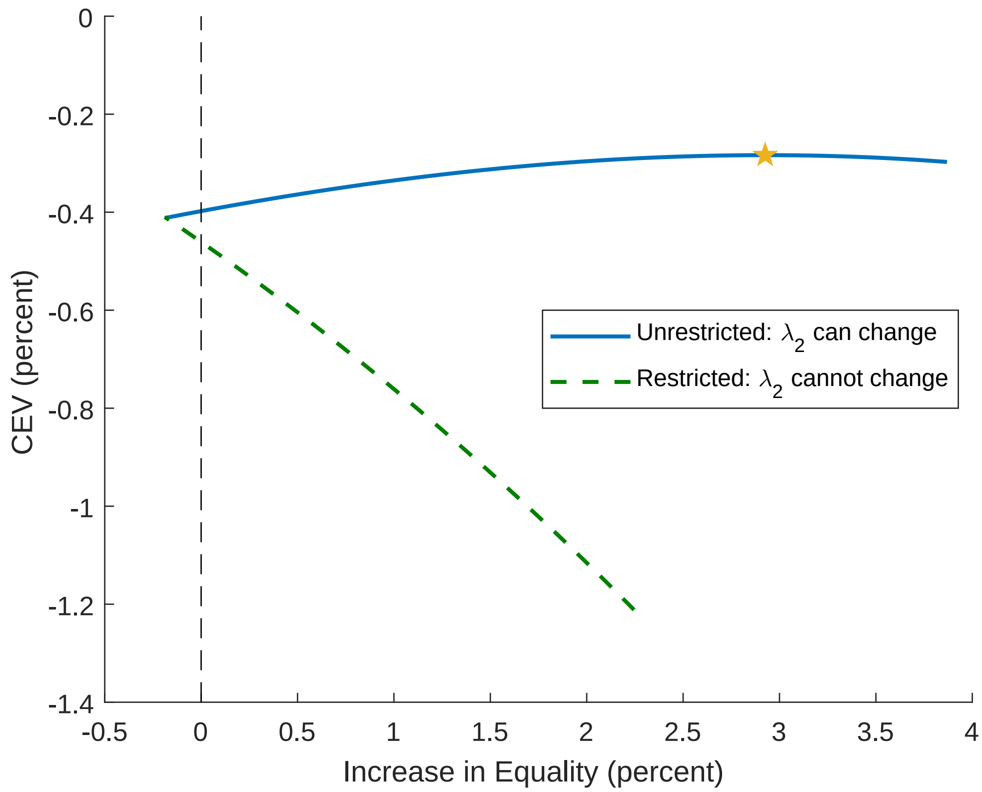

Like the unrestricted frontier, the set of policies on the restricted frontier never return any carbon tax revenue through reductions in the level of the labor tax rate. Instead of changing the progressivity of the labor-tax function to achieve a given level of equality, the policymaker on the restricted frontier increases the level of the lump-sum rebate function and decreases its slope (see Appendix B).

The restricted frontier is downward-sloping, and with the exception of the left-most point (the pure capital tax rebate policy), lies everywhere below the unrestricted frontier. Consistent with the results from the previous literature, the optimal policy along the restricted frontier rebates all of the carbon tax revenue through a reduction in the distortionary capital tax. This restricted optimum increases inequality from its value in the baseline, restoring the trade-off between welfare-maximizing and progressive policies.

The fact that the restricted frontier falls below the unrestricted frontier highlights that, in order to increase equality, providing lump-sum rebates is a much more costly approach as opposed to increasing the labor tax progressivity. The welfare cost difference stems from rebates' different impacts on agents' savings behavior over the lifecycle. Importantly, agents receive the lump-sum rebate in every year of life, including retirement. As a result, the lump-sum rebate directly reduces an agent's need to save for retirement, crowding out capital. In contrast, since agents only receive the labor-tax rebate during their working years, it does not crowd out as much capital and, thus, is less costly. Figure 7 plots the percent change in capital from the baseline at each point on the restricted and unrestricted frontiers. Capital is always lower on the restricted frontier.

Figure 7: Percent Change in Capital Along the Restricted and Unrestricted Frontiers

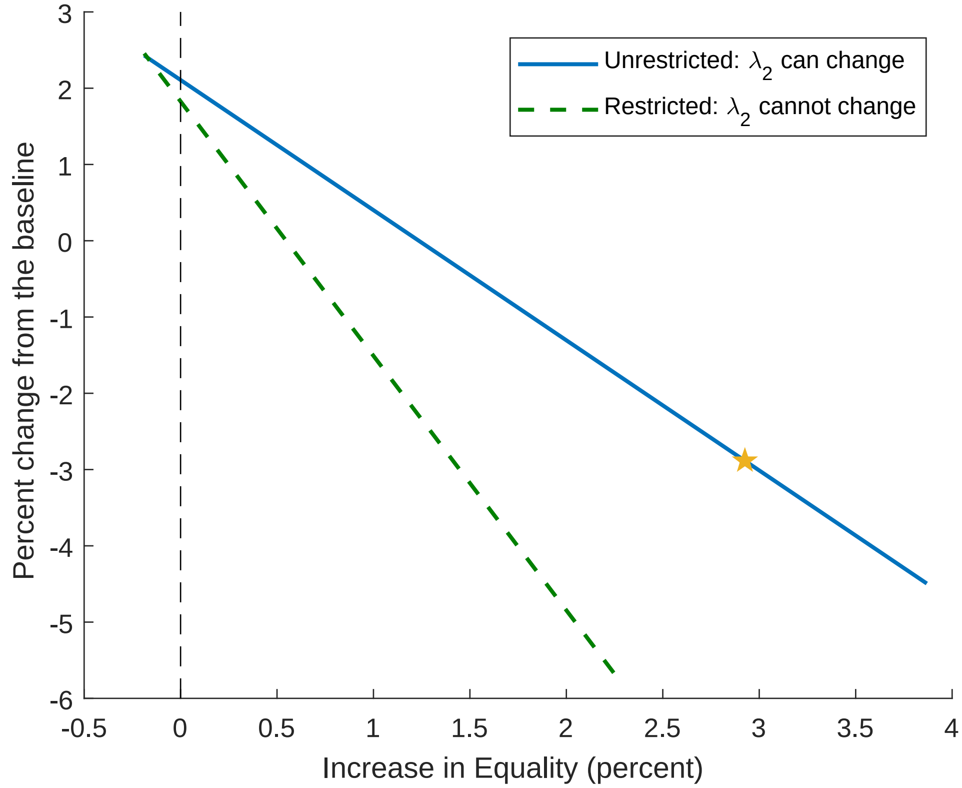

Lower capital primarily operates through the level component of CEV to reduce welfare along the restricted frontier. The dashed green and solid blue lines in the left panel of Figure 8 plot the level component of CEV along the restricted and unrestricted welfare-equality frontiers, respectively. The level component is much steeper on the restricted frontier because the large declines in capital lead to large decreases in consumption (see Appendix B). The dashed green and solid blue lines in the right panel of Figure 8 plot the sum of the age and distribution components along the restricted and unrestricted frontiers, respectively. The difference between the solid blue and dashed green lines in the right panel of Figure 8 is much smaller than in the left panel, implying that the welfare difference between the restricted and unrestricted frontiers is primarily driven by the different savings incentives and their corresponding effects on aggregate capital and the level component.

Figure 8: CEV Components: Unrestricted and Restricted Frontiers

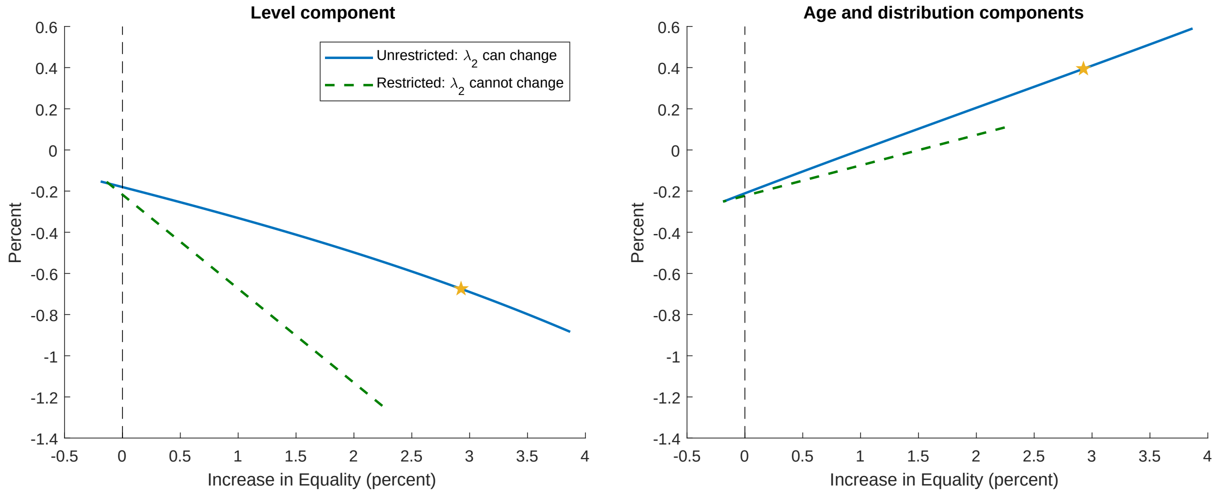

References

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, Roberton C Williams III, and Dallas Burtraw, "Deficit Feduction and Carbon Taxes: Budgetary, Economic, and Distributional Impacts," *Resources for the Future Report*, August 2013.

CBO, "Reducing the Deficit: Spending and Revenue Options," Technical Report, United States Congressional Budget Office March 2011.

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 and Dirk Krueger, "On the Optimal Progressivity of the Income Tax Code," *Journal of Monetary Economics*, 2006, *53*, 1425–1450.

_ , 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 and 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 AC 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.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante, "Optimal tax progressivity: An analytical framework," *The Quarterly Journal of Economics*, 2017, *132* (4), 1693–1754.

Horowitz, John, Julie-Anne Cronin, Hannah Hawkins, Laura Konda, and Alex Yuskavage, "Methodology for Analyzing a Carbon Tax," Technical Report, United States Department of the Treasury, Office of Tax Analysis, Washington D.C. 2017. Working Paper 115.

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 MH 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.

Mathur, Kevin Hassett Aparna and Gilbert Metcalf, "The Incidence of a U.S. Carbon Tax: A Lifetime and Regional Analysis," *Energy Journal*, 2009, *30(2)*, 157–180.

Metcalf, Gilbert, "A Proposal for a US Carbon Tax Swap: An Equitable Tax Reform to Address Global Climate Change," *Hamilton Project, Brookings Institute*, 2007.

Metcalf, Gilbert E, "A Distributional Analysis of Green Tax Reforms," *National Tax Journal*, 1999, pp. 655–681.

Nakajima, Makoto, "Optimal Capital Income Taxation with Housing," *working paper*, 2010.

Parry, Ian WH, "Pollution Taxes and Revenue Recycling," *Journal of Environmental Economics and Management*, 1995, *29* (3), S64–S77.

Pearce, David, "The role of carbon taxes in adjusting to global warming," *The Economic Journal*, 1991, *101* (407), 938–948.

Peterman, William and Erick Sager, "Optimal Public Debt with Life Cycle Motives," Technical Report, Federal Reserve Board of Governors 2018.

Peterman, William B and Kamila Sommer, "How Well Did Social Security Mitigate the Effects of the Great Recession?," *International Economic Review*, Forthcoming.

Repetto, Robert, Roger C Dower, Robin Jenkins, and Jacqueline Geoghegan, *Green fees: how a tax shift can work for the environment and the economy*, World Resources Institutes, 1992.

Terkla, David, "The efficiency value of effluent tax revenues," *Journal of Environmental Economics and Management*, 1984, *11* (2), 107–123.

Trabandt, Mathias and Harald Uhlig, "The Laffer curve revisited," *Journal of Monetary Economics*, 2011, *58*, 305–327.

Williams, Roberton, Hal Gordon, Dalla Burtraw, and Jared Carbone, "The Initial Incidence of a Carbon Tax Across Income Groups," *National Tax Journal*, 2015, *68(1)*, 195–214.

Appendix

A Model and calibration: additional details

A.1 Definition of an equilibrium

We define a sequence-of-markets equilibrium for this economy as a sequence of prices, $\{w_t, r_t, p^e_t\}_{t=0}^{\infty}$, allocations for each household i age j, $\{c_{i,j,t}, e^c_{i,j,t}, a_{i,j,t+1}, h_{i,j,t}\}_{t=0}^{\infty}$, allocations for firms, $\{E^y_t, K^y_t, N^y_t, K^e_t, N^e_t\}_{t=0}^{\infty}$, social security taxes, $\{T^s_{i,j,t}\}_{t=0}^{\infty}$, transfers, $\{T^a_t, T^c_t\}_{t=0}^{\infty}$, and the distribution of individuals over the state space, Φ_t, such that the following holds:

1. Given prices, household allocations maximize,

$$\frac{\tilde{c}_{i,j,t}^{1-\theta_1}}{1-\theta_1} - \chi \frac{h_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} + \mathbb{E}\left\{ \sum_{k=j+1}^{J} \beta^{k-j} \prod_{q=j}^{k-1} (\Psi_q) \frac{\tilde{c}_{i,j,t}^{1-\theta_1}}{1-\theta_1} - \chi \frac{h_{i,j,t}^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} \right\}, \tag{A1}$$

   subject to the budget constraint:

$$\begin{aligned} c_{i,j,t} + (p^e_t + \tau^c_t) e^c_{i,j,t} + a_{i,j,t+1} = {} & \mu_{i,j,t} h_{i,j,t} w_t - T^s_{i,j,t} + (1 + r_t(1-\tau^k))(a_{i,j,t} + T^a_t) \\ & - T^h_t\big(\mu_{i,j,t} h_{i,j,t} w_t - 0.5 T^s_{i,j,t}\big) + T^c_t \quad \text{for } j < j^r \end{aligned} \tag{A2}$$

$$c_{i,j,t} + (p^e_t + \tau^c_t) e^c_{i,j,t} + a_{i,j,t+1} = b^{ss}(x_{i,j,t}) + (1 + r(1-\tau^k))(a_{i,j,t} + T^a_t) + T^c_t \text{ for } j \geq j^r$$

   and the non-negativity constraints, $c_t \geq 0$, $a_t \geq 0$, $h_t \geq 0$, and $e^c_t \geq 0$.

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

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

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:

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

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: CHANGE THIS

$$\tau^s = \frac{S \sum_{j \geq j^r} \Phi(x)}{\sum_{j<j^r} [\min(y_{ijt}^h,\, y^{h,max}) \Phi(x)]}$$

A *stationary competitive equilibrium* consists of prices, $\{w, r, p^e\}$, allocations for firms, $\{E^y, K^y, N^y, K^e, N^e\}$, social security taxes, $\{T_{i,j}^s\}$, 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.

A.2 Social security

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.[24]

A.3 Energy calibration targets

Using data from the five most recent years of the Consumer Expenditures Survey (CEX) (2013-2017), we find that energy budget share in the top half of the income distribution is 33.9 percent smaller than energy budget share in the bottom half of the income distribution. However, the variance in expenditures in the CEX is larger than in our model. In particular, the percent difference in total expenditures between the top and bottom half of the distribution is 288.8 percent in the CEX and only X percent in our model.[25] Following Fried et al. (2018), we adjust the energy-share difference so that $\frac{X}{288.8} = \frac{Y}{-33.8}$. Therefore, we choose $\bar{e} = 0.00474$ to target an energy-share difference between the top and bottom halves of the expenditure distribution of Y percent.

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

A.4 CEV decomposition

Define the CEV between two steady states, denoted by a, and, b, as the percent increase in expected consumption an individual would need in every period in steady state a so that she

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

[25]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.

[26]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).

is indifferent between steady state a and steady state b. More formally, the CEV between steady states a and b solves the following equation

$$\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j U((1+\Delta_{CEV})\tilde{c_j}^a, h_j^a)d\Phi(x^a) = \int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j U(\tilde{c_j}^b, h_j^b)d\Phi(x^b), \tag{A3}$$

where superscripts a and b denote the values in steady states a and b respectively and $\Psi_j \equiv \prod_{j=1}^{j-1}\psi_j$. Hence, positive values of the CEV imply that individuals are better off in steady state b and negative values imply that they are better off in steady state a. Following Peterman and Sager (2018), we divide the CEV into level, age, and distribution components,

$$(1+\Delta_{CEV}) = (1+\Delta_{level})(1+\Delta_{age})(1+\Delta_{dist}). \tag{A4}$$

Since utility depends on both consumption and hours, we further divide each CEV component into a consumption part and an hours part,

$$(1+\Delta_{CEV}) = [(1+\Delta_{C_{level}})(1+\Delta_{H_{level}})][(1+\Delta_{C_{age}})(1+\Delta_{H_{age}})][(1+\Delta_{C_{dist}})(1+\Delta_{H_{dist}})].$$

We provide the definition and intuition for each component below. See Peterman and Sager (2018) for a complete proof that the decomposition holds.

The level-consumption component measures the effect of a change in the average level of consumption between the two steady states,

$$1+\Delta_{C_{level}} = \frac{\bar{c}^b}{\bar{c}^a},$$

where $\bar{c}^a$ and $\bar{c}^b$ denote average consumption in steady states a and b, respectively. The level component is equivalent to the total CEV in a model without age and income heterogeneity (e.g., an infinitely lived representative individual).

The age-consumption component measures the welfare effects of changes in an individual's profile of consumption over the life-cycle, controlling for any changes in the average level of consumption (the level component). Mathematically, the age-component equals,

$$1+\Delta_{C_{age}} = \frac{\left(\sum_{j=1}^{J}\beta^{j-1}\Psi_j(\bar{c}_j^b)^{1-\theta_1}\right)^{\frac{1}{1-\theta_1}}/\bar{c}^b}{\left(\sum_{j=1}^{J}\beta^{j-1}\Psi_j(\bar{c}_j^a)^{1-\theta_1}\right)^{\frac{1}{1-\theta_1}}/\bar{c}^a} \tag{A5}$$

where $\bar{c}_j^a$ and $\bar{c}_j^b$ denote the average consumption at age j in steady states a and b respectively. Thus, the numerator in equation (A5) denotes the expected lifetime welfare, measured in terms of consumption-equivalence, for an average individual in steady state b, normalized by the level of consumption in steady state b. The denominator is the analogous value for steady state a.

Finally, the distribution effect measures the welfare effects of changes in the variance of consumption within an age cohort, holding constant the average level of consumption at each age. The distribution component equals,

$$1+\Delta_{C_{dist}} = \left[\frac{\left(\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j(\tilde{c_j}^b)^{1-\theta_1} d\Phi(x^b)\right) / \left(\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{c_j}^b)^{1-\theta_1}\right)}{\left(\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j(\tilde{c_j}^a)^{1-\theta_1} d\Phi(x^a)\right) / \left(\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{c_j}^a)^{1-\theta_1}\right)}\right]^{\frac{1}{1-\theta_1}}. \tag{A6}$$

The left term in the numerator equals the consumption component of expected lifetime welfare for an individual in steady state b. We divide this value by the consumption component of expected lifetime welfare for an individual with the average consumption profile in steady state b. This division removes the welfare effects of the age and level components. What remains is the distribution component. The denominator is the equivalent expression for steady state a.

The intuition for the hours components of CEV is same as for the consumption-components. However, the expressions are more complex because we measure the utility effects of changes in hours in terms of consumption, instead of hours. The hours pieces of the level, age, and distribution CEV components are,

$$1+\Delta_{H_{level}} = \left[1+\left(1-\frac{\bar{h}^b}{\bar{h}^a}\right)^{\frac{1}{1+\frac{1}{\theta_2}}} \times \left(\frac{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j\chi\left(\frac{(h_j{}^a)^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}}\right) d\Phi(x^a)}{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j\left(\frac{(\tilde{c_j}^b)^{1-\theta_1}}{1-\theta_1}\right) d\Phi(x^b)}\right)\right]^{\frac{1}{1-\theta_1}}$$

$$1+\Delta_{H_{age}} = \left[1+\left(\left(\frac{\bar{h}^b}{\bar{h}^a}\right)^{\frac{1}{1+\frac{1}{\theta_2}}} - \frac{\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{h}_j^b)^{1+\frac{1}{\theta_2}}}{\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{h}_j^a)^{1+\frac{1}{\theta_2}}}\right) \times \left(\frac{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j\chi\left(\frac{(h_j{}^a)^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}}\right) d\Phi(x^a)}{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j\left(\frac{(\tilde{c_j}^b)^{1-\theta_1}}{1-\theta_1}\right) d\Phi(x^b)}\right)\right]^{\frac{1}{1-\theta_1}}$$

$$1+\Delta_{H_{dist}} = \left[1 + \left(\frac{\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{h}_j^b)^{1+\frac{1}{\theta_2}}}{\sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{h}_j^a)^{1+\frac{1}{\theta_2}}} \times \frac{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j({h_j}^b)^{1+\frac{1}{\theta_2}} d\Phi(x^b)}{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j(\bar{h}_j^a)^{1+\frac{1}{\theta_2}} d\Phi(x^a)} \right) - \right. \tag{A7}$$

$$\left. \left(\frac{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j \chi \left(\frac{({h_j}^a)^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} \right) d\Phi(x^a)}{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j \left(\frac{(\tilde{c_j}^b)^{1-\theta_1}}{1-\theta_1} \right) d\Phi(x^b)} \right) \right]^{\frac{1}{1-\theta_1}} \tag{A8}$$

The term,

$$\frac{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j \chi \left(\frac{({h_j}^a)^{1+\frac{1}{\theta_2}}}{1+\frac{1}{\theta_2}} \right) d\Phi(x^a)}{\int_0^1 \sum_{j=1}^{J} \beta^{j-1}\Psi_j \left(\frac{(\tilde{c_j}^b)^{1-\theta_1}}{1-\theta_1} \right) d\Phi(x^b)}$$

converts the welfare effect of a change in hours into consumption units.

B Additional results

Figure A1 displays the discrete policies simulations on the welfare-equality frontier as well as the polynomial fit to the simulated data.

Figure A1: Polynomial fit to the welfare-equality frontier

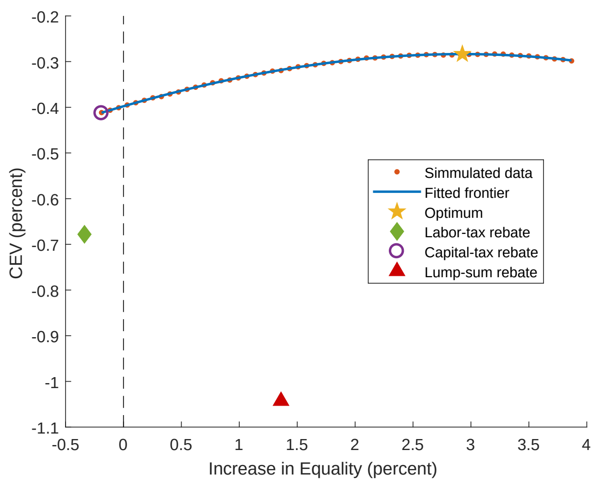

Figure A2 plots the rebate instruments on the restricted frontier. Like the unrestricted frontier, the restricted frontier relies only on the capital-tax rebate, the labor-tax rebate always equals zero. Instead of relying on changes in the progressivity of the labor-tax rebate to achieve a give level of equality, the policymaker on the restricted frontier increases the level of the lump-sum rebate function and decreases its slope.

Figure A2: Rebate instruments on the restricted frontier

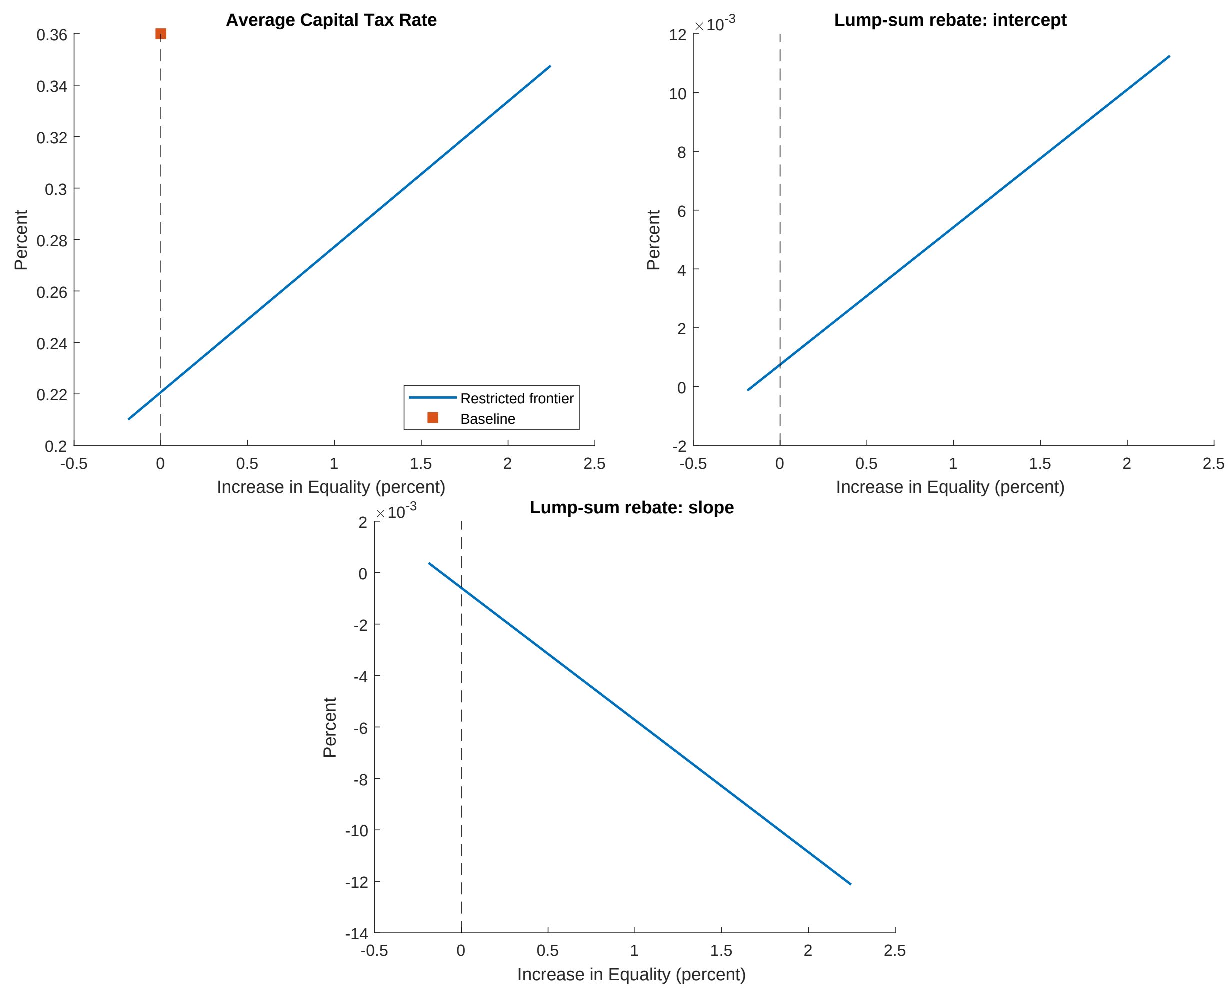

Figure A3 plots output, consumption, and hours on the restricted and unrestricted frontier.

Figure A3: Macro aggregates: restricted and unrestricted frontiers

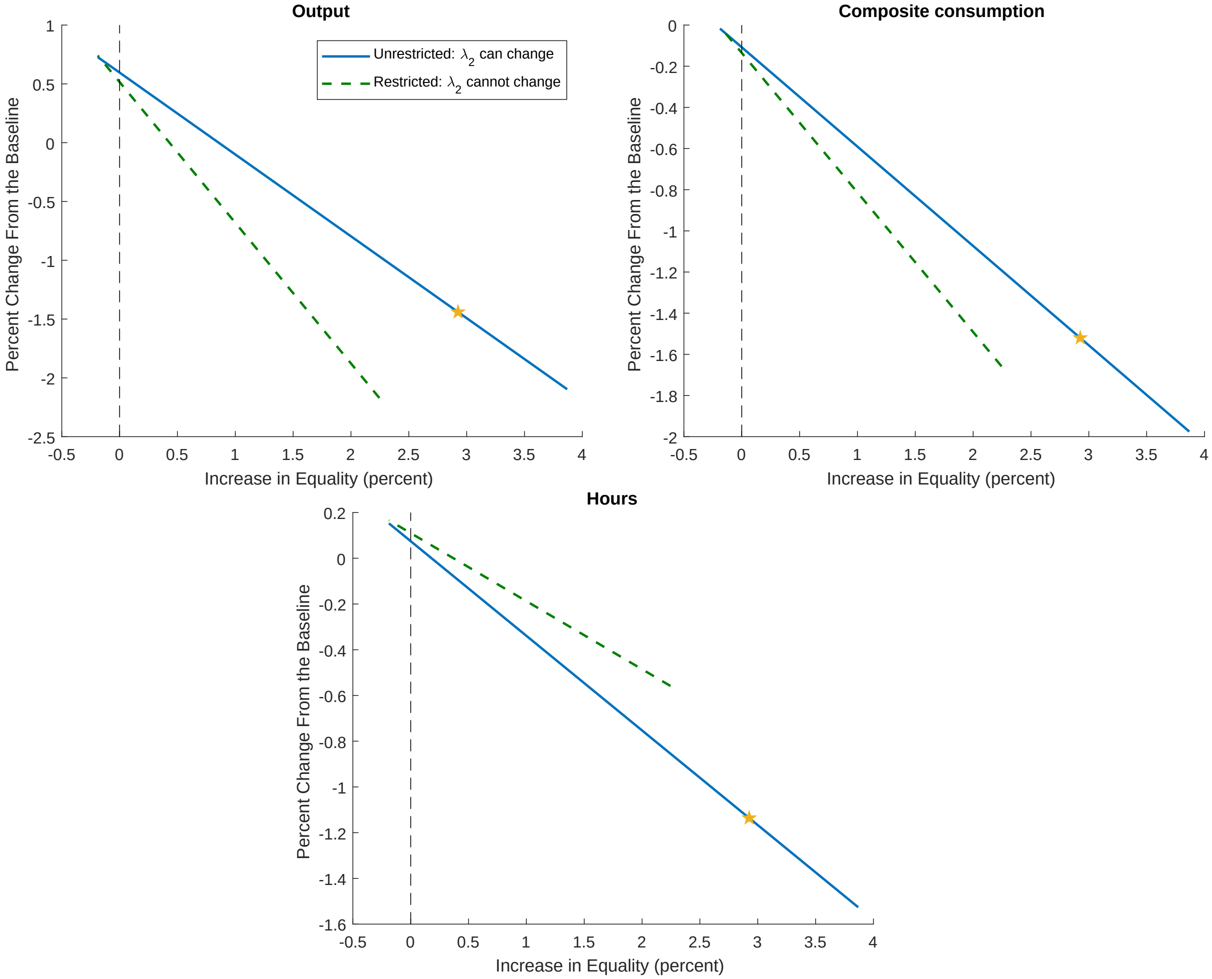

Figure A4 plots capital and welfare on the restricted and unrestricted frontier. I.

Figure A4: Capital and welfare along the frontier – with and without retirees receiving lump-sum payments