Using Consumption Data to Derive Optimal Income and Capital Tax Rates

Christian Hellwig and Nicolas Werquin

REVISED
May 16, 2024

WP 2022-03

https://doi.org/10.21033/wp-2022-03

FEDERAL RESERVE BANK *of* CHICAGO

Using Consumption Data to Derive Optimal Income and Capital Tax Rates*

Christian Hellwig

Toulouse School of Economics and CEPR

Nicolas Werquin

Federal Reserve Bank of Chicago, Toulouse School of Economics, and CEPR

May 16, 2024

Abstract

We study a Mirrleesian economy with labor income, consumption, and retirement savings and derive a novel representation of optimal non-linear income and savings distortions that highlights the role of consumption inequality and consumption responses to tax changes. Our representation establishes a close connection between the formula for top income taxes of Saez (2001) and the uniform commodity taxation theorem of Atkinson and Stiglitz (1976): One cannot be valid without the other, and departures from this joint benchmark lead to a clear trade-off between income and savings taxes. Consumption data in turn discipline the optimal departure from this benchmark. Because consumption is more evenly distributed than income, it is optimal to shift a substantial fraction of the top earners' tax burden from income to savings.

*This paper supersedes our earlier paper, "A Fair Day's Pay for a Fair Day's Work: Optimal Tax Design as Redistributional Arbitrage." We thank Gadi Barlevy, Marco Bassetto, Charlie Brendon, Ashley Craig, Antoine Ferey, Lucie Gadenne, Alexandre Gaillard, Aart Gerritsen, Mike Golosov, Martin Hellwig, Luca Micheletto, Serdar Ozkan, Yena Park, Florian Scheuer, Karl Schulz, Emmanuel Saez, Stefanie Stantcheva, Aleh Tsyvinski, and Philipp Wangner for useful comments. Opinions expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Christian Hellwig acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010).

1 Introduction

Optimal Top Income Tax Rate: A Puzzle. Consider a static economy in which workers differ in productivity and choose consumption C and labor income Y given a non-linear income tax schedule. In a milestone contribution to the modern theory of income taxation, Saez (2001) shows that the revenue-maximizing income tax rate on top earners is given by

$$\overline{\tau}_Y^{Saez} \quad = \quad \frac{1}{1-\zeta_Y^I+\rho_Y\zeta_Y^H}, \tag{1}$$

where $\rho_Y > 1$ represents the upper Pareto coefficient of taxable income, $\zeta_Y^H > 0$ is the compensated elasticity of taxable income with respect to the retention rate (i.e., one minus the marginal tax rate), and $\zeta_Y^I > 0$ is the income effect of a lump-sum tax levy on taxable income.[1] The appeal of this formula is that it readily lends itself to quantitative evaluations of the optimal top tax rate. For example, using standard empirical estimates of the taxable income elasticities, $\zeta_Y^H = 0.33$ and $\zeta_Y^I = 0.25$ (see Section 3.1) and the Pareto coefficient of annual total income $\rho_Y = 1.5$ (Diamond and Saez 2011), we obtain $\overline{\tau}_Y^{Saez} = 80\%$.

A striking feature of this optimal income tax formula is that it completely abstracts from consumption to focus exclusively on income inequality and labor supply responses to taxes. This is puzzling, since the marginal benefit of higher income taxes should intuitively depend on how much redistribution the existing tax system already achieves, that is, on the distribution of consumption rather than that of pre-tax income. More generally, focusing exclusively on measures of income inequality may paint an incomplete picture of the link from allocations to welfare, since in practice agents may insure against labor market risks through other means than the government, such as private insurance, precautionary savings, or intra-family transfers.[2]

The static model studied by Saez (2001) in fact admits an alternative, consumption-based,

[1] Formally, these variables are defined as follows. We let $\rho_Y \equiv -\lim_{Y\to\infty} \partial \ln(1-F_Y(Y))/\partial \ln Y$, where $F_Y(\cdot)$ is the CDF of the income distribution. A lower value of ρ_Y indicates a more unequal distribution, that is, a thicker right tail. We also let $\zeta_Y^H \equiv \partial \ln Y^H/\partial \ln(1-\tau_Y)$, where τ_Y is the marginal income tax rate and Y^H is the compensated labor supply. Finally, $\zeta_Y^I \equiv \partial \ln Y/\partial \overline{T}_Y$, where $\overline{T}_Y \equiv T_Y(Y)/((1-\tau_Y)Y)$ is the (normalized) average tax rate and Y the uncompensated labor supply. Both Y^H and Y are expressed in terms of earnings rather than hours.

[2] Echoing this discussion, equation (1) contrasts with a vast literature on optimal risk sharing that uses consumption inequality and consumption responses to income shocks to assess the optimality of risk-sharing arrangements or redistributive policies—including in models with asymmetric information of the kind studied in Saez (2001). The empirical emphasis on consumption for testing efficiency of risk-sharing arrangements is well known since (at least) Townsend (1994). See, e.g., Ligon (1998) and Kocherlakota and Pistaferri (2009) for applications of this idea in a hidden information context.

representation of revenue-maximizing top income taxes:

$$\overline{\tau}_Y^{Cons} = \frac{1}{\zeta_C^I + \rho_C \zeta_C^H}, \tag{2}$$

where $\rho_C > 1$ is the Pareto tail coefficient of the consumption distribution, $\zeta_C^H > 0$ is the compensated elasticity of consumption with respect to the retention rate, and $\zeta_C^I > 0$ is (minus) the income effect of a lump-sum tax levy on consumption.[3] Equation (2) follows immediately from equation (1) and three model-implied identities, $\rho_C = \rho_Y$, $\zeta_C^H = \zeta_Y^H$, and $\zeta_C^I = 1 - \zeta_Y^I$, which all follow from the observation that consumption equals after-tax income in a static setting. Hence, the role of consumption inequality and consumption responses for optimal income taxes is not separately identified from the role of income inequality and labor supply. While the literature following Saez (2001) has systematically focused on equation (1), one could equally well use the consumption-based formula (2), or any combination of the two, to characterize the optimal policy, assuming that reliable estimates of the corresponding sufficient statistics are available.

Using consumption rather than income inequality measures radically changes the quantitative evaluation of the optimal top tax rate, however. Recent empirical evidence shows that consumption is significantly *more evenly* distributed than income among top earners, with a Pareto coefficient for consumption of $\rho_C \approx 3.1$, more than twice the Pareto coefficient for total income (Toda and Walsh 2015; Buda et al. 2022; Gaillard et al. 2023). Using the same taxable income elasticities as in the first paragraph, the optimal top income tax rate falls to 55% if we use consumption inequality rather than income inequality in formula (1), more than doubling the after-tax income of top earners compared to the income-based estimate given in the opening paragraph.

Thus, the static model unfortunately provides no guidance about which measures are the most appropriate for estimating income taxes. Moreover, it is inconsistent with the empirical discrepancy between consumption and income inequality. Interpreting the model-implied identities instead as testable over-identifying restrictions, their violation suggests that the static model of Saez (2001) is not well suited to address how consumption or income inequality matter for optimal taxes, or to offer sound, empirically grounded policy prescriptions.

Our Contribution. Motivated by these observations, in this paper we study to what extent consumption inequality and consumption responses to income shocks influence optimal tax design, independently of income inequality and labor supply responses. We consider a Mirrleesian

[3] Formally, we define $\rho_C \equiv -\lim_{Y\to\infty} \partial \log(1 - F_C(Y)) / \partial \log C$, $\zeta_C^H \equiv \partial \ln C^H / \partial \ln(1 - \tau_Y)$, and $\zeta_C^I \equiv -\partial \ln C / \partial \overline{T}_Y$.

economy in which agents with heterogeneous labor productivities work, consume and save for retirement.[4] The additional consumption-savings margin allows us to separate consumption from after-tax income, and also introduces capital taxes as a second margin for redistributive taxation. This connects our analysis with another milestone result, the "*zero capital taxation*" (also known as "*uniform commodity taxation*", henceforth UCT) theorem of Atkinson and Stiglitz (1976), which states that it is optimal to leave consumption choices undistorted—and, hence, not tax savings—if preferences for savings are independent of labor productivity. savings taxes must thus be rationalized by departures from this benchmark, for example if agents with higher earnings capacity also have a stronger preference for savings (see Saez 2002).

Our main result (Theorem 2) shows that $\overline{\tau}_Y^{Saez}$ and UCT are two sides of the same coin. Formally, we show that the revenue-maximizing top income tax is equal to $\overline{\tau}_Y^{Saez}$ *if and only if* UCT applies and the optimal savings tax is 0, except in the trivial case where savings vanish at the top of the income distribution. Conversely, departures from this joint benchmark identify a tradeoff between labor income and savings taxes: Either it is optimal to tax savings and reduce the income tax below $\overline{\tau}_Y^{Saez}$, or it is optimal to subsidize savings and raise the income tax above $\overline{\tau}_Y^{Saez}$.

Moreover, data on *consumption* inequality and *consumption* responses to income tax changes allow us to test the empirical relevance of UCT and characterize the optimal savings tax and the departure of the optimal income tax from $\overline{\tau}_Y^{Saez}$. We re-evaluate revenue-maximizing taxes on top income earners and argue that it is optimal to shift a significant part of redistributive taxes from income to savings: the optimal policy prescribes a strictly positive and quantitatively significant savings tax, while substantially lowering the income tax below the level prescribed by $\overline{\tau}_Y^{Saez}$.

What connects the income tax formula (1) of Saez (2001) to UCT? Saez (2001)'s characterization applies to a static economy with a single labor supply margin, and only requires estimates of income inequality and labor supply elasticities. Hence, it continues to apply in a dynamic setting, and the optimal income tax is equal to $\overline{\tau}_Y^{Saez}$, *if and only if* the design of income taxes can be reduced to a static tradeoff between labor supply and after-tax earnings, with no information needed about how the latter is allocated between consumption and savings. But this condition is met precisely when the UCT theorem applies and it is optimal not to tax savings. Conversely, when the UCT theorem fails to hold, the design of optimal income taxes can no longer ignore how after-tax income is allocated to consumption and savings, since preferences for savings and incentives to work are no longer independent.

Why is it optimal to shift part of the tax burden from income to savings? The fact that

[4] Our baseline model is kept deliberately as simple as possible. In Section 5, we show that our results carry over to more general settings that include life-cycle and stochastic economies.

consumption is far less concentrated than income in the upper tail suggests that top earners have a vanishing propensity to consume out of earned income. This could arise either because savings have a higher income elasticity than consumption for given preferences, i.e., agents view savings as a luxury good relative to consumption, or because preferences for savings correlate positively with labor productivity. But existing empirical evidence on consumption responses to tax changes suggests that the income elasticity of consumption cannot be so low as to explain the gap between income and consumption inequality for top earners. Hence we reject UCT in favor of a shift towards positive savings and lower income taxes.

Formally, we first provide a novel representation of optimal income and savings taxes as a function of the Pareto tail coefficients of income, consumption and savings, as well as behavioral responses of these variables to income tax changes (Theorem 1). This result generalizes the income tax formula (1) to economies with consumption and savings, and provides an analogous formula for optimal savings taxes. Our representation highlights that when agents work, consume, and save, inequality and behavioral elasticities of all three variables matter independently for the design of optimal income and savings taxes, even if they are jointly determined by the optimal allocation and must satisfy similar over-identifying restrictions as in the static model.

As a direct corollary of Theorem 1, we then show that it is optimal to tax savings and set the income tax below $\overline{\tau}_Y^{Saez}$ if and only if $\rho_C \zeta_C^H > \rho_Y \zeta_Y^H$, i.e., whenever the product of the Pareto coefficient and the compensated elasticity is higher for consumption than for income; in the opposite case, we should tax savings and raise the income tax above the static optimum. This is precisely the parameter condition that underlies the empirical argument in support of positive savings taxes when comparing the concentration of consumption vs. income, summarized by the ratio of Pareto coefficients ρ_Y / ρ_C, to the pass-through of income changes to consumption, summarized by the ratio of compensated elasticities ζ_C^H / ζ_Y^H.

Finally, the tradeoff between income and savings taxes described by Theorem 2 becomes especially stark when consumption has a thinner tail than income, and the highest income earners save almost all of their income. In that case, which appears to be the empirically relevant one, the analysis reduces again to a static tradeoff at the top, but now between labor supply and *savings*, rather than between labor supply and consumption as in the static model of Saez (2001). Therefore, the static wedge $\overline{\tau}_Y^{Saez}$ continues to determine the optimal policy, but it now accounts for the optimal *combined* wedge between income and savings. Consumption data then identify how this combined wedge should be decomposed into an income tax and a savings tax, i.e., how much one should depart from the joint benchmark of Saez (2001) and Atkinson and Stiglitz (1976).

Policy Implications. Our results have important consequences for policy design. The income tax formula (1) of Saez (2001) and the UCT theorem of Atkinson and Stiglitz (1976) have both been highly influential in shaping tax policies, but they are typically invoked independently from each other: The former features prominently in discussions about income tax design, while debate about savings taxes focuses on the policy relevance of UCT. For example, a highly cited review article by Diamond and Saez (2011) simultaneously makes a case for high top income taxes based on the formula for $\overline{\tau}_Y^{Saez}$ and for positive capital taxes by questioning the assumptions underlying UCT. Theorem 2 instead implies that policy discussion of income and savings taxes cannot be conducted in isolation from each other, but both are part of a policy mix that optimally trades off between multiple margins of redistribution. Hence, the two policy recommendations by Diamond and Saez (2011) are mutually inconsistent.[5]

Our characterization of optimal income and savings taxes in Theorem 1 in turn provides empirical guidance on how the policy maker should resolve this tradeoff: government revenue is maximized by shifting part of the tax burden from income to savings, and consumption data allow us to identify the magnitude of this shift. Hence, we provide empirical support for the second recommendation by Diamond and Saez (2011), while at the same time invalidating their first recommendation.

Relationship to the Literature. Our paper relates to the optimal taxation literature originating with Mirrlees (1971), as well as the sufficient statistics approach towards estimating optimal tax rates that was pioneered by Saez (2001). Our model is based on Atkinson and Stiglitz (1976). Because we allow for arbitrary preferences, their uniform commodity taxation theorem only applies as a special case of our framework.[6] Mirrlees (1976), Saez (2002), and Golosov, Troshkin, Tsyvinski, and Weinzierl (2013) study a similar problem as ours but do not characterize the optimal top tax rates analytically nor express the formulas in terms of empirically observable sufficient statistics. Our paper also relates to a large literature that uses consumption to discipline models of insurance or risk-sharing contracts, e.g., Townsend (1994), Ligon (1998), Kocherlakota and Pistaferri (2009), and Heathcote, Storesletten, and Violante (2014).

The most closely related papers are Scheuer and Slemrod (2021), Gerritsen et al. (2020), Schulz

[5] The tradeoff between income and savings taxes is intuitive from a revenue maximization perspective: since savings taxes reduce incentives to work, they negatively impact the government revenue from income taxes. Likewise, a higher income tax will induce agents to work and save less, and therefore reduces savings tax revenues. Therefore, the revenue-maximizing mix of tax policies is no longer at the peak of a one-dimensional Laffer curve (except when UCT applies), but trades off between higher revenues from one instrument versus the other. Discussing income and savings tax policies independently from each other instead creates the false impression that the government "can have its cake and eat it too".

[6] Christiansen (1984), Jacobs and Boadway (2014), and Gauthier and Henriet (2018) generalize Atkinson and Stiglitz (1976) to non-homothetic preferences, but typically constrain commodity or savings taxes to being linear.

(2021), and Ferey, Lockwood, and Taubinsky (2021). They recently characterized optimal savings taxes in models that are similar to ours. In contrast to these papers, we focus on the joint the design of income and savings taxes for top income earners, emphasize the use of consumption data rather than savings to characterize and quantify optimal income and capital taxes, and do not impose strong *a priori* restrictions on preferences or the sources of underlying return heterogeneity that drives the departure from Atkinson and Stiglitz (1976).[7]

First, by studying jointly the design of income and savings taxes, we identify the joint benchmark that consists of Saez (2001) for the former and Atkinson and Stiglitz (1976) for the latter (Theorem 2), and highlight a sharp tradeoff between these instruments away from this benchmark. To our knowledge, these theoretical results are new to the literature.[8] Second, the emphasis on consumption data for characterizing optimal income and savings taxes (Theorem 1 and Corollary 1) allows us to identify the underlying structure of preferences that pins down the relevant departure from Atkinson and Stiglitz (1976) in a particularly simple and transparent form. Moreover, consumption data are essential for identifying optimal taxes on *top* income earners, who are the main focus of our analysis. In particular, Ferey, Lockwood, and Taubinsky (2021) estimate optimal savings taxes along the income distribution, using the causal effect of income on savings as their main sufficient statistic. As we show formally in Appendix B.1, their identification breaks down at the top of the income distribution, when consumption is strictly less concentrated at the top than income: The savings of top earners then respond one-to-one to income changes, which in turn renders their sufficient statistic uninformative. Our paper and theirs are thus complementary.

Outline of the Paper. In Section 2, we set up our baseline model and derive our main theoretical results on optimal labor and savings taxes (Theorem 1, Corollary 1) and on the relationship between Saez (2001) and Atkinson and Stiglitz (1976) (Theorem 2). In Section 3, we study the quantitative implications of these results. In Section 4, we provide a proof of our first theorem. In Section 5 shows that our results carry over to more general environments. Appendices A and B contain the proofs and additional results.

[7]Scheuer and Slemrod (2021) derive a characterization of the savings tax rates on top earners when agents have exogenous endowments in addition to labor income and marginal utilities of consumption are independent of productivity types. In Appendix B.2, we extend our results to a model that nests their framework and discuss in detail the relationship between both papers. Gerritsen et al. (2020) and Schulz (2021) focus on a model with heterogeneous returns, assuming that preferences satisfy the Atkinson-Stiglitz restrictions. As we explain in Section 5.1, our model nests this framework as a special case. On the other hand, they explore various microfoundations of return heterogeneity that are beyond the scope of our analysis.

[8]Gerritsen et al. (2020) and Schulz (2021) also mention a trade-off between labor and savings taxes, but our result holds at a much higher level of generality, and we prove and characterize this trade-off analytically rather than discussing it only verbally or quantitatively.

2 Optimal Taxation of Top Earners

In this section, we set up the simplest environment that allows us to decouple consumption from after-tax income and derive our optimal labor and savings tax formulas. In Section 5, we show that our results extend directly to much richer environments, with an arbitrary number of periods and goods, heterogeneous rates of return, and stochastic income shocks.

2.1 Environment

There is a continuum of measure 1 of heterogeneous agents indexed by a rank $r \in [0,1]$ uniformly distributed over the unit interval. The preferences of agents of rank r are defined over "consumption" C, "savings" S, and "labor income" Y.[9] They are represented by

$$U(C, Y; r) + V(S; r),$$

where for any r, the functions U and V are twice continuously differentiable with $U_C > 0$, $U_{CC} < 0$, $U_Y < 0$, $U_{YY} < 0$, $V_S > 0$, $V_{SS} < 0$ and satisfy the usual Inada conditions as C, Y or S approach 0 or ∞. We further assume that $U(C, Y; r) + V(S; r)$ is non-decreasing in r, for any (C, Y, S). We interpret U and V as the first- and second-period utility functions, respectively.

Assumption 1 (**Single-Crossing Conditions**). *(i) The marginal rates of substitution (MRS) between consumption and income, $-U_Y/U_C$, and between saving and income, $-U_Y/V_S$, are strictly decreasing in r for all (C, S, Y), i.e.,*

$$\frac{\partial \ln(-U_Y/U_C)}{\partial r} \equiv \frac{U_{Yr}}{U_Y} - \frac{U_{Cr}}{U_C} < 0. \tag{3}$$

$$\frac{\partial \ln(-U_Y/V_S)}{\partial r} \equiv \frac{U_{Yr}}{U_Y} - \frac{V_{Sr}}{V_S} < 0. \tag{4}$$

Furthermore, the marginal disutility of effort is decreasing in r, $U_{Yr}/U_Y < 0$.

(ii) The MRS between consumption and savings, V_S/U_C, is monotonic in r for all (C, Y, S), i.e.,

$$\frac{\partial \ln(V_S/U_C)}{\partial r} \equiv \frac{V_{Sr}}{V_S} - \frac{U_{Cr}}{U_C} \tag{5}$$

is either non-positive or non-negative everywhere.

The single-crossing condition (3) is standard (Mirrlees 1971). Condition (4) introduces its

[9] While it is convenient for the analysis to define preferences in terms of the observables C, Y, and S, it is straightforward to map the type-contingent preference over income into a preference over leisure or labor supply.

analogue with regards to savings. These conditions rank agents according to their preferences over leisure, on the one hand, and consumption or savings, on the other hand. On the margin, agents with higher rank r are more willing to work for a given consumption or savings increase. The additional restriction $U_{Yr}/U_Y < 0$ implies that higher ranks r find it less costly to attain a given income level Y; that is, we can associate agents' ranks with their labor productivity.

The second part of Assumption 1 imposes that the inter-temporal MRS is monotonic. If it is increasing, so that (5) is positive, then higher ranks have a stronger taste for saving (relative to current consumption) than lower ranks. In other words, given the same allocation, those who are the most inclined to work—the higher ranks—are also the most inclined to save. If instead (5) is negative, then those who are the most inclined to work are also those who are the most inclined to spend their incomes on current consumption.

Social Planner's Problem. Consumption, income, and savings are assumed to be observable, but an individual's productivity rank r is their private information. We assume that the social planner is Rawlsian and wishes to maximize the utility of the lowest-ranked agent; for our results on optimal top taxes, this assumption is without loss of generality.[10] The optimal allocation $\{C(r), Y(r), S(r)\}$ maximizes the net present value of tax revenue

$$\int_0^1 \{Y(r) - C(r) - S(r)\} dr \tag{6}$$

subject to the incentive compatibility constraint

$$U(C(r), Y(r); r) + V(S(r); r) \;\geq\; U(C(r'), Y(r'); r) + V(S(r'); r) \tag{7}$$

for all types r and announcements r', and a lower bound constraint on the lowest rank's utility

$$U(C(0), Y(0); 0) + V(S(0); 0) \;\geq\; W_0. \tag{8}$$

To ease notation, we write $X(r) \equiv X(C(r), Y(r), S(r); r)$ for any function X of both the allocation $(C(r), Y(r), S(r))$ and the type r.

Under the Rawlsian objective and the monotonicity of utilities with respect to r, the social planner values transferring resources towards lower-ranked or less productive agents. On the mar-

[10] We generalize our results to arbitrary planner preferences in the Appendix. Our top income tax results remain unaffected so long as the marginal utility converges to 0 at infinity, so that top-ranked agents receive vanishing weight in the planner's objective function.

gin, this redistribution can occur through three channels: by redistributing effort from less to more productive agents, or equivalently leisure towards less productive agents—that is, redistribution *"from each according to his ability"*; by redistributing consumption or savings towards lower-income households whose marginal utilities are the highest—that is, redistribution *"to each according to his needs"*. The respective signs and magnitudes of U_{Yr}/U_Y, U_{Cr}/U_C, and V_{Sr}/V_S govern how preferences over these various channels of redistribution vary with agents' types, and therefore how the corresponding policy instruments interact with their incentives to work and to save.

Labor and Savings Wedges. Let $\tau_Y(r) \equiv U_Y(r)/U_C(r) + 1$ denote the *labor wedge* at rank r implied by the optimal allocation $\{C(\cdot), Y(\cdot), S(\cdot)\}$, i.e., the intra-temporal distortion between the marginal product and the marginal rate of substitution between consumption and income. Let $\tau_S(r) \equiv V_S(r)/U_C(r) - 1$ denote the *savings wedge* at rank r, i.e., the inter-temporal distortion in the agent's first-order condition for savings.

An interpretation of our optimal tax system is a combination of income taxes, social security contributions and pension payments ("savings") that are indexed to labor income, without any additional private savings. The savings wedge then represents the marginal shortfall or excess of social security contributions relative to pension payments. Alternatively, we could relabel S in our model as "bequests", let C and Y stand for life-time earnings and consumption, and reinterpret the savings wedge as a tax on bequests.

2.2 Optimal Top Income and Savings Tax Rates

When should income and savings be taxed? It is well-known since Mirrlees (1976) that the optimal labor and savings wedges inherit the signs of $U_{Cr}(r)/U_C(r) - U_{Yr}(r)/U_Y(r)$ and $V_{Sr}(r)/V_S(r) - U_{Cr}(r)/U_C(r)$, respectively. Assumption 1 implies that the former is positive, and hence that it is optimal to tax labor income. Analogously, it is optimal to tax (respectively, subsidize) savings whenever the latter is positive (resp., negative), that is, if higher ranks have a higher (resp., lower) intertemporal MRS V_S/U_C and are thus more inclined to save (resp., consume) their current income than lower ranks. Intuitively, if the more productive ranks have a stronger taste for savings, the planner can screen them—i.e., deter them from mimicking lower ranks—by taxing the savings of lower ranks. This general result nests the uniform commodity taxation setting of Atkinson and Stiglitz (1976) as a special case. When all ranks r have the same intertemporal MRS V_S/U_C, savings taxes are unable to affect the low-productivity workers differently than the more productive ones who mimick them. It is then optimal to set $\tau_S(r) = 0$ for all r, so that redistribution should only be achieved through the income tax without further distorting the intertemporal consumption margin.

Our first main result, Theorem 1, extends these insights by providing a novel analytical representation of the rank-dependence of the intra- and inter-temporal MRS $U_{Cr}/U_C - U_{Yr}/U_Y$ and $V_{Sr}/V_S - U_{Cr}/U_C$, and hence the optimal income and savings taxes. Our optimal tax formulas depend on two sets of sufficient statistics: distributional parameters, such as the Pareto coefficients, and preference parameters or, equivalently, behavioral elasticities with respect to tax changes.

We denote by $\rho_Y(r), \rho_C(r), \rho_S(r)$ the local Pareto coefficients of the distributions of labor income, consumption, and savings, respectively, and by $s_C(r)$ the share of consumption in retained income at rank r. For any $X \in \{Y, C, S\}$,

$$\rho_X(r) \equiv \frac{d\ln(1-F_X(X(r)))}{d\ln X(r)}, \quad \text{and} \quad s_C(r) \equiv \frac{C(r)}{(1-\tau_Y(r))Y(r)}$$

where F_X denotes the CDF of the distribution of X. We assume that these Pareto coefficients and the consumption share converge to constants for top earners; we denote $\rho_X = \lim_{r\to1}\rho_X(r)$ and $s_C = \lim_{r\to1} s_C(r)$.

In addition, we define the following preference elasticities:

$$\zeta_C(r) \equiv -\frac{C(r)U_{CC}(r)}{U_C(r)},\ \zeta_S(r) \equiv -\frac{S(r)V_{SS}(r)}{V_S(r)},\ \zeta_Y(r) \equiv \frac{Y(r)U_{YY}(r)}{U_Y(r)},\ \zeta_{CY}(r) \equiv \frac{Y(r)U_{CY}(r)}{U_C(r)}.$$

We assume that these elasticities converge to finite limits $\zeta_C, \zeta_S, \zeta_Y, \zeta_{CY}$ as $r \to 1$. Below, we represent these parameters in terms of the substitution and income effects of labor income taxes on taxable income (ζ_Y^H, ζ_Y^I), consumption (ζ_C^H, ζ_C^I), and savings (ζ_S^H, ζ_S^I).

Theorem 1 characterizes optimal income and savings wedges in terms of these Pareto coefficients and preference elasticities:

Theorem 1. *Suppose that Assumption 1 holds and that the optimal allocation $\{C(\cdot), Y(\cdot), S(\cdot)\}$ is co-monotonic. Suppose moreover that as $r \to 1$ we have $\zeta_C/\rho_C < 1 + \zeta_{CY}/\rho_Y$ and $\zeta_S/\rho_S < 1$. Then the optimal labor wedge on top income earners, $\overline{\tau}_Y \equiv \lim_{r\to1}\tau_Y(r)$, satisfies*

$$\overline{\tau}_Y = \frac{\zeta_Y/\rho_Y + \zeta_C/\rho_C - (1 + s_C\rho_Y/\rho_C)\zeta_{CY}/\rho_Y}{1 + \zeta_Y/\rho_Y - s_C\zeta_{CY}/\rho_C} \tag{9}$$

and the optimal savings wedge on top income earners, $\overline{\tau}_S \equiv \lim_{r\to1}\tau_S(r)$, satisfies

$$\overline{\tau}_S = \frac{\zeta_S/\rho_S - \zeta_C/\rho_C + \zeta_{CY}/\rho_Y}{1 - \zeta_S/\rho_S}. \tag{10}$$

Equation (9) generalizes the top income tax rate formula of Saez (2001) to a dynamic environment. Equation (10) provides an analogous formula for savings taxes. We derive these equations from a pair of perturbations that identify possible welfare improvements starting from a sub-optimal tax schedule: If the observed income or savings wedge is higher (lower) than the theoretical optimum described by equations (9) and (10), then a marginal reduction (increase) of the corresponding wedge is strictly welfare-improving. Moreover, the characterization of the optimal labor wedge in equation (9) applies even if the savings wedge is not at its optimal level, and *vice versa*.

The Role of Consumption Inequality for Optimal Taxes. Theorem 1 represents optimal income and savings taxes explicitly as a function of the Pareto tail coefficients and preference elasticities of consumption and savings in addition to those of labor income. *Ceteris paribus*, high income and consumption inequality (i.e., small Pareto coefficients) both lead to high optimal top tax rates on labor income, while high wealth inequality but low consumption inequality lead to high optimal top tax rates on savings.

The reason why income, consumption, and savings inequality all matter for optimal taxes appears clearly in the proof of Theorem 1 in Section 4 below. Proposition 3 shows that the rank-dependence of the intra- and inter-temporal MRS, $U_{Cr}/U_C - U_{Yr}/U_Y$ and $V_{Sr}/V_S - U_{Cr}/U_C$, which determines the sign and magnitude of optimal income and savings taxes, can be identified from the ratios of preference elasticities and Pareto coefficients, along with the progressivity of the income and savings taxes in place. More specifically, if τ_S converges to a constant at the top, we obtain

$$(1-r)\left(\frac{V_{Sr}}{V_S} - \frac{U_{Cr}}{U_C}\right) \;=\; \frac{\zeta_S}{\rho_S} - \frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y}. \tag{11}$$

With constant top savings taxes, agents' consumption and savings must grow with after-tax income at rates that keep the inter-temporal MRS constant, hence $1/V_S$ and $1/U_C$ must have identical upper Pareto tails. With rank-independent inter-temporal MRS ($V_{Sr}/V_S - U_{Cr}/U_C = 0$), this in turn implies that $\zeta_S/\rho_S = \zeta_C/\rho_C - \zeta_{CY}/\rho_Y$. Conversely, if preferences are rank-dependent ($V_{Sr}/V_S \neq U_{Cr}/U_C$), the upper Pareto coefficients of $1/V_S$ and $1/U_C$ augment the terms ζ_S/ρ_S and $\zeta_C/\rho_C - \zeta_{CY}/\rho_Y$, which capture the role of diminishing marginal utilities at a given rank, with additional terms $(1-r)\,V_{Sr}/V_S$ and $(1-r)\,U_{Cr}/U_C$ that capture the rank-dependence of $1/V_S$ and $1/U_C$ at a given allocation. Equating the two Pareto coefficients then implies that any difference between ζ_S/ρ_S and $\zeta_C/\rho_C - \zeta_{CY}/\rho_Y$ on the RHS of equation (11) mirrors one-for-one the rank-dependence of the inter-temporal MRS on the LHS, and can thus be used to fully identify the latter. Along similar lines, we identify the rank-dependence of the intratemporal MRS by equating the Pareto tail coefficients of $1/(-U_Y)$ and $1/U_C$.

To build intuition, consider first the case without preference complementarities ($\zeta_{CY} = 0$) and suppose that savings are more concentrated at the top than consumption ($\rho_C > \rho_S$). This could occur either because top-ranked earners face more strongly diminishing marginal utilities of consumption than of savings ($\zeta_C > \zeta_S$), or because agents' intrinsic preferences for savings are rank-dependent. The ratio ζ_S/ζ_C identifies the concentration of savings relative to consumption that is consistent with rank-independent preferences, while the ratio ρ_S/ρ_C identifies the concentration of savings relative to consumption observed in the data. Rank-independence is rejected when the former differs from the latter. By comparing ρ_S/ρ_C to ζ_C/ζ_S we identify to what extent preference heterogeneity is required to rationalize the gap between consumption and savings inequality, and thus to what extent savings taxes or subsidies are useful to screen and redistribute from higher- to lower-ranked agents, on top of redistributive income taxation. When $\zeta_{CY} > 0$, the preference complementarity with earnings reduces the rate at which the marginal utility of consumption vanishes at the top, but the argument is otherwise the same. The ratios of preference elasticities and Pareto coefficients are therefore natural and transparent sufficient statistics for intrinsic rank-dependence of preferences and hence optimal taxes.

2.3 From Theory to Observables

To characterize optimal taxes in terms of observables, we first identify additional model-implied restrictions that they must satisfy. We then relate preference elasticities to observable counterparts.

Additional Model-Implied Restrictions: A Tale of Three Tails. First, note that the three Pareto coefficients and the consumption share s_C are linked through the agents' inter-temporal budget constraint. As $r \to 1$, the spending on consumption or savings of top-ranked agents cannot grow faster than their after-tax income, so we have $\rho_C \geq \rho_Y$ and $\rho_S \geq \rho_Y$. Hence the cross-sectional distribution of labor income must have a weakly thicker upper tail than those of consumption and savings. But total spending on consumption and savings must grow at the same rate as after-tax income, which in turn implies that $\rho_Y = \min\{\rho_C, \rho_S\}$. This relationship is the analogue of the condition $\rho_Y = \rho_C$ in the static setting. Moreover, $s_C = 0$ if $\rho_C > \rho_Y$ and $s_C = 1$ if $\rho_S > \rho_Y$, i.e., the share of consumption (resp., savings) in retained earnings must be vanishing if consumption (resp., savings) has a strictly thinner tail than income. On the other hand, $s_C \in [0, 1]$ remains unrestricted if $\rho_Y = \rho_C = \rho_S$. Hence only ρ_Y and one of the variables ρ_C, ρ_S, s_C are unrestricted, with the other two statistics in each case pinned down by these additional restrictions.

Second, the substitution and income effects on taxable income (ζ_Y^H, ζ_Y^I), consumption (ζ_C^H, ζ_C^I), and savings (ζ_S^H, ζ_S^I) satisfy two additional identities, $\zeta_Y^H = s_C \zeta_C^H + (1 - s_C) \zeta_S^H$ and $1 - \zeta_Y^I = s_C \zeta_C^I + (1 - s_C) \zeta_S^I$, which are the analogue of the conditions $\zeta_Y^H = \zeta_C^H$ and $1 - \zeta_Y^I = \zeta_C^I$ in the static model. The four preference elasticities are therefore exactly identified from these six behavioral responses to tax changes. However, the behavioral responses that matter for identification depend on the limit of the consumption share s_C: If $s_C = 1$, the income and substitution effects on consumption (ζ_C^H, ζ_C^I) mirror those on taxable income (ζ_Y^H, ζ_Y^I), and therefore identification requires using the response of savings to income tax changes (ζ_S^H, ζ_S^I). If $s_C \in (0,1)$, any two pairs of income and substitution effects identify the third, along with the four preference elasticities. Finally, if $s_C = 0$, then the income and substitution effects on savings (ζ_S^H, ζ_S^I) mirror those on taxable income (ζ_Y^H, ζ_Y^I), and therefore identification requires using the response of *consumption* to income tax changes (ζ_C^H, ζ_C^I).

Our model thus admits three possible scenarios that we summarize as follows:

1. If $\rho_Y = \rho_C < \rho_S$ and $s_C = 1$, savings are strictly less concentrated at the top than income and consumption and top earners consume (almost) all their income. The behavioral responses of consumption and labor supply then coincide, thus making information from consumption redundant, relative to income data. Preference elasticities and optimal taxes are then identified from Pareto coefficients and behavioral responses of income and *savings*.

2. If $\rho_Y = \rho_C = \rho_S$ and $s_C \in [0,1]$, income, consumption, and savings are equally concentrated at the top, and top earners consume and save non-vanishing fractions of their income. The substitution and income effects of any two of the three variables (income, consumption, savings) can then be used to identify our four preference elasticities and determine optimal income and savings taxes.

3. If $\rho_Y = \rho_S < \rho_C$ and $s_C = 0$, consumption is strictly less concentrated at the top than income and savings, and top earners save (almost) all their income. The behavioral elasticities of savings and labor supply coincide, making savings information redundant relative to income data. Preference elasticities and optimal taxes are then identified from Pareto coefficients and behavioral responses of income and *consumption*.

In the sequel, we refer to these cases as *Case 1*, *Case 2*, and *Case 3*, respectively.

Linking Preference Elasticities to Standard Behavioral Responses. Following Chetty (2006), we break the agent's decision into a first stage in which the agent trades off between labor supply and after-tax earnings $M \equiv Y - T_Y(Y)$, and a second stage in which after-tax earnings are allocated to consumption and savings. Define $\mathcal{U}(M, Y; r) \equiv \max_{C,S} \{U(C, Y; r) + V(S, r)\}$ s.t. $C + S + T_S(S) \leq M$ as the indirect utility function that characterizes the solution to the second-stage problem. The first-stage problem then maximizes $\mathcal{U}(Y - T_Y(Y), Y; r)$ with respect to Y, and is identical to the static optimal tax problem solved by Saez (2001).

First-Stage: Which Parameters do the Labor Supply Responses Identify? Applying the same arguments as in the static model (equations (23) and (24) in Saez (2001)), we thus obtain that the substitution and income effects of taxable income with respect to tax changes, (ζ_Y^H, ζ_Y^I), identify the following parameters:

$$\frac{1 - \zeta_Y^I}{\zeta_Y^H} = \tilde{\zeta}_Y - \tilde{\zeta}_{MY}, \quad \text{and} \quad \frac{\zeta_Y^I}{\zeta_Y^H} = \tilde{\zeta}_M - \tilde{\zeta}_{MY}. \tag{12}$$

where $\tilde{\zeta}_M \equiv -M\mathcal{U}_{MM}/\mathcal{U}_M$ is the coefficient of relative risk aversion over after-tax earnings, $\tilde{\zeta}_Y \equiv Y\mathcal{U}_{YY}/\mathcal{U}_Y$ is the first-stage curvature over labor supply, and $\tilde{\zeta}_{MY} \equiv Y\mathcal{U}_{MY}/\mathcal{U}_M$ is the complementarity between after-tax earnings and labor supply. Suppose first that the first-period utility function is separable between consumption and leisure, so that $\zeta_{CY} = 0$. In this case, we have

$$\frac{1}{\tilde{\zeta}_M} = \frac{s_C}{\zeta_C} + \frac{1 - s_C}{\zeta_S}, \quad \tilde{\zeta}_Y = \zeta_Y, \quad \tilde{\zeta}_{MY} = 0.$$

Thus, income and substitution effects of taxable income identify the labor supply parameter ζ_Y and the relative risk aversion coefficient over after-tax income $\tilde{\zeta}_M$ via the relationships $(1-\zeta_Y^I)/\zeta_Y^H = \zeta_Y$ and $\zeta_Y^I/\zeta_Y^H = \tilde{\zeta}_M$. More generally, preference complementarities $\zeta_{CY} > 0$ lead to $\tilde{\zeta}_{MY} > 0$ and modify the relationship between $\tilde{\zeta}_Y$ and ζ_Y. The corresponding expressions are given in Appendix A.2 (equation (29)). We obtain that, fixing the preference complementarity ζ_{CY} and the curvature in consumption ζ_C, income and substitution effects of taxable income continue to identify the labor supply parameter ζ_Y and the relative risk aversion coefficient over after-tax income $\tilde{\zeta}_M$.

Second Stage: Which Parameters do Consumption and Savings Responses Identify? The consumption or savings responses to income tax changes then serve to identify the curvatures on consumption and savings ζ_C, ζ_S, and the preference complementarity ζ_{CY}. Suppose again for simplicity that $\zeta_{CY} = 0$. This restriction implies that the relative strength of income and substitution effects for taxable income, consumption and savings must be the same: $(1 - \zeta_Y^I)/\zeta_Y^H = \zeta_C^I/\zeta_C^H = \zeta_S^I/\zeta_S^H$. In addition, we have that $\zeta_C\zeta_C^H = \zeta_S\zeta_S^H$ and $\zeta_C\zeta_C^I = \zeta_S\zeta_S^I$, i.e., income and substitution effects for consumption and savings are inversely proportional to preference elasticities. Intuitively, in the absence of preference complementarities, agents make no difference between earned and unearned income changes and pass those changes through to consumption and savings with elasticities that are inversely proportional to the respective preference elasticities ζ_C and ζ_S. We thus obtain that the ratio of preference elasticities ζ_C/ζ_S is naturally linked the ratios of income and substitution effects on consumption and savings:

$$\frac{\zeta_C}{\zeta_S} = \frac{1/\zeta_C^I}{1/\zeta_S^I} = \frac{1/\zeta_C^H}{1/\zeta_S^H}.$$

More generally, when $\zeta_{CY} > 0$, we show in Appendix A.2 that the ratio ζ_C/ζ_S is adjusted downwards compared to ζ_S^I/ζ_C^I and upwards compared to ζ_S^H/ζ_C^H. Thus, consumption and income are complements if and only if $\zeta_S^I/\zeta_C^I > \zeta_S^H/\zeta_C^H$, and since the adjustments depend on the value of ζ_{CY}, we can then infer the degree of preference complementarity by comparing these two ratios of income and substitution effects. Intuitively, substitution effects for taxable income, consumption and savings all have the same sign (a higher income tax causes agents to work, consume and save less). The complementarity between consumption and earnings increases the response of consumption relative to savings for given ζ_C/ζ_S. This in turn implies that for given ζ_S^H/ζ_C^H, the inferred ratio of elasticities ζ_C/ζ_S increases. On the other hand, an increase in unearned income reduces earned income, while increasing consumption and savings. In this case the complementarity weakens the income effects of consumption relative to savings for given ζ_C/ζ_S, which in turn implies that the inferred ratio of elasticities ζ_C/ζ_S is lower than ζ_S^I/ζ_C^I with complementarities.

Elasticity of Intertemporal Substitution. We finally define the elasticity of inter-temporal substitution (EIS) as follows:

$$\zeta_{IS} \equiv -\frac{\partial \ln (S/C)}{\partial \ln (1+\tau_S)}\bigg|_{Y,\mathcal{U}\text{ constant}} = \frac{1}{s_C\zeta_S + (1-s_C)\,\zeta_C}.$$

Thus, the EIS ζ_{IS} and risk aversion $\tilde{\zeta}_M$ jointly depend on the preference elasticities ζ_C and ζ_S. In *Case 1*, $s_C = 1$ implies that $\zeta_{IS} = 1/\zeta_S$. That is, while consumption mirrors after-tax income at the top and the first-stage elasticities fully determine the consumption-earnings tradeoff ($\tilde{\zeta}_M = \zeta_C$, $\tilde{\zeta}_Y = \zeta_Y$, and $\tilde{\zeta}_{MY} = \zeta_{CY}$), the EIS governs the response of *savings* to tax changes (and is thus identified from the latter). The opposite holds in *Case 3* ($s_C = 0$): the first-stage elasticities imply $\tilde{\zeta}_M = \zeta_S$, $\tilde{\zeta}_Y = \zeta_Y$, $\tilde{\zeta}_{MY} = 0$, and the EIS $\zeta_{IS} = 1/\zeta_C$ is identified from the response of *consumption* to tax changes. In this case, the top agents face a static trade-off between earnings and savings, with consumption responses to tax changes determined by inter-temporal substitution.

This discussion also highlights that risk aversion $\tilde{\zeta}_M$ and intertemporal substitution ζ_{IS} play conceptually different roles for determining optimal income and savings taxes: the former governs variation in marginal utilities of after-tax income, which enters the social planner's redistribution motive. The latter affects how much consumption and savings respond to tax changes, which in turn determines the relative importance of diminishing marginal utilities and rank-dependent preferences in accounting for consumption and savings inequality—and thus how much redistribution should be shifted from income to savings taxes.

Optimal Taxes in Terms of Observables. We are now in a position to answer our central motivating question: To what extent do consumption inequality and consumption responses to income taxes matter for optimal tax design, independently of income inequality and labor income responses? The following corollary to Theorem 1 provides alternative expressions for the optimal income and savings tax rates in terms of behavioral elasticities. This result pinpoints the precise role that consumption data play in disciplining optimal income and savings taxes in the expressions given in Theorem 1. For simplicity, we focus again on the case of a separable utility, and treat the general case in Appendix A.2.

Corollary 1. *Suppose that $\zeta_{CY} = 0$. The optimal income tax can be rewritten as:*

$$\overline{\tau}_Y \;=\; \overline{\tau}_Y^{Saez}\left[1 - \zeta_Y^I\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)\right] \tag{13}$$

where $\overline{\tau}_Y^{Saez}$ is given by (1), and the optimal savings tax can be rewritten as:

$$\overline{\tau}_S \;=\; \frac{\zeta_Y^I}{\rho_S \zeta_S^H - \zeta_Y^I}\left(1 - \frac{\rho_S \zeta_S^H}{\rho_C \zeta_C^H}\right). \tag{14}$$

Hence, $\overline{\tau}_Y \lesseqqgtr \overline{\tau}_Y^{Saez}$ if and only if $\rho_C\,\zeta_C^H \gtreqqless \rho_Y\,\zeta_Y^H$, and $\overline{\tau}_S \gtreqqless 0$ if and only if $\rho_C\,\zeta_C^H \gtreqqless \rho_S\,\zeta_S^H$. Suppose in addition that $s_C < 1$ (Cases 2 and 3). It is then optimal to tax savings if and only if the product of the consumption-Pareto and the EIS is large enough:

$$\zeta_{IS} \;>\; \frac{\rho_Y}{\rho_C}\frac{\zeta_S^H}{\zeta_Y^I}. \tag{15}$$

Corollary 1 yields empirically testable necessary and sufficient conditions for the deviation from the benchmark optimal tax results of Saez (2001) and Atkinson and Stiglitz (1976), obtained by comparing the substitution effects of consumption with those of earnings and savings, respectively. As long as the utility function is not quasilinear in consumption (i.e., $\zeta_Y^I > 0$), the optimal labor income tax is strictly lower than the static benchmark whenever the product of the Pareto coefficient and the Hicksian elasticity is larger for consumption than for labor income, $\rho_C \zeta_C^H > \rho_Y \zeta_Y^H$; analogously, the optimal savings tax is strictly positive if and only if $\rho_C \zeta_C^H > \rho_S \zeta_S^H$. In the next section, we show that these departures from the Saez (2001) and Atkinson and Stiglitz (1976) benchmarks, respectively, are in fact closely linked.

Furthermore, condition (15) gives a particularly simple necessary and sufficient condition under which taxing savings is optimal, in the form of a lower bound on the elasticity of intertemporal substitution, or an upper bound on the ratio of Pareto coefficients on income vs. consumption. If the consumption share of income converges to zero at the top, which we argue below is the empirically relevant case, this condition simply reads $\zeta_{IS} > (\rho_Y \zeta_Y^H)/(\rho_C \zeta_Y^I)$, which depends on parameters that all have readily available estimates in the empirical literature. In Section 3, we use condition (15) to argue that for empirically plausible values of the EIS the optimal savings tax is indeed strictly positive.

Equations (30), (31), and (32)-(33) in the Appendix generalize the formulas of Corollary 1 to the case $\zeta_{CY} > 0$. Importantly, the results that $\overline{\tau}_Y \lesseqqgtr \overline{\tau}_Y^{Saez}$ if and only if $\rho_C \, \zeta_C^H \gtreqqless \rho_Y \, \zeta_Y^H$, and $\overline{\tau}_S \gtreqqless 0$ if and only if $\rho_C \, \zeta_C^H \gtreqqless \rho_S \, \zeta_S^H$, continue to hold. When it is optimal to tax savings, however, the magnitude of $\overline{\tau}_S$ (and, correspondingly, the lower bound on the EIS) are adjusted upwards if $\zeta_{CY} > 0$. These comparative statics follow from the familiar logic of Corlett and Hague (1953): When preferences are non-separable, it is optimal to tax less heavily the goods that are complementary to labor. Thus, a higher degree of complementarity between consumption and labor income unambiguously raises the optimal top savings tax rate and, symmetrically, lowers the optimal top income tax rate.

2.4 Deviations from Seminal Taxation Results

Our second theorem connects our representation of optimal income and savings taxes in Theorem 1 and Corollary 1 to the seminal results of Saez (2001) and Atkinson and Stiglitz (1976). It highlights a deep connection between these two benchmarks, which to our knowledge has not been recognized before: One cannot hold without the other—outside of the "trivial" case where savings vanish (as a fraction of earnings) at the top of the income distribution.[11] In other words, the static optimal income tax rate and the uniform commodity taxation theorem are two sides of the same coin.

[11] In a simpler model with return heterogeneity as the only departure from Atkinson and Stiglitz (1976), Schulz (2021) shows that the static optimal formula of Saez (2001) applies when UCT is satisfied, but he does not characterize analytically the monotone trade-off between labor and savings taxes away from this benchmark.

Theorem 2. *1. Consider Case 1, in which $s_C = 1$ and $\rho_C = \rho_Y < \rho_S$. Then the optimal top income tax rate $\overline{\tau}_Y$ is equal to the static optimum $\overline{\tau}_Y^{Saez} = \overline{\tau}_Y^{Cons}$ given by equations (1) or (2).*

2. Consider next Cases 2 and 3; i.e., suppose that $s_C < 1$. Then the optimal savings tax rate on top earners is positive (resp., negative), $\overline{\tau}_S \gtreqless 0$, if and only if the optimal top income tax rate is lower (resp., higher) than the static optimum, $\overline{\tau}_Y \lesseqgtr \overline{\tau}_Y^{Saez}$.

3. In particular, in Case 3, where $s_C = 0$, the static wedge is equal to the combined wedge on income and savings:

$$1 - \overline{\tau}_Y^{Saez} \;=\; \frac{1 - \overline{\tau}_Y}{1 + \overline{\tau}_S}. \tag{16}$$

Consider first *Case 1*, with $s_C = 1$. In this case, savings simply do not matter at the top and the model converges to the standard static trade-off between consumption and income; hence, $\overline{\tau}_Y^{Saez}$ and $\overline{\tau}_Y^{Cons}$ continue to provide correct representations of optimal income taxes in the dynamic model. Both consumption and income inequality still matter for optimal income taxes even in this case, but this distinction only arises at a conceptual level—it does not affect quantative prescriptions, because Pareto coefficients and behavioral responses of income and consumption are observationally equivalent in the upper tail. Thus, the optimal income tax $\overline{\tau}_Y$ can be inferred from either $(\rho_Y, \zeta_Y^H, \zeta_Y^I)$ or $(\rho_C, \zeta_C^H, \zeta_C^I)$. On the other hand, savings taxes may or may not satisfy the uniform commodity taxation theorem and still be either positive or negative, but this departure cannot be inferred by comparing income and consumption data—it depends instead on the comparison between the ratio of Pareto tails for income and *savings* ρ_Y/ρ_S with a ratio of compensated labor supply and *savings* elasticities, ζ_S^H/ζ_Y^H.

This raises the question under what conditions consumption inequality and consumption responses to tax changes matter for optimal tax policies independently of income inequality and income responses, or conversely, under what conditions the income-based representation $\overline{\tau}_Y^{Saez}$ correctly identifies optimal income taxes in the dynamic model. The second part of Theorem 2 answers this question by showing that whenever $s_C < 1$, the optimal income tax equals $\overline{\tau}_Y^{Saez}$ *if and only if* the Atkinson-Stiglitz theorem applies and $\overline{\tau}_S = 0$. To understand this insight, recall that we can represent the agent's labor supply and consumption-savings decision as a two-stage problem, where the second stage allocates after-tax income M to consumption and savings, and the first-stage decision determines labor supply and after-tax income. Therefore, the optimal income tax is equal to $\overline{\tau}_Y^{Saez}$ if and only if the optimal tax design can be reduced to a static problem with preferences directly defined over earnings Y and after-tax income M, with no information needed about how the latter is allocated between consumption and savings. But this condition is met precisely when preferences satisfy the necessary and sufficient condition for the uniform commodity taxation theorem of Atkinson and Stiglitz (1976); that is, if the MRS between consumption and savings is independent of rank r. In this case, income and substitution effects of labor supply, along with the Pareto coefficient for earnings, are sufficient to estimate the optimal income tax, and it is optimal not to tax savings. Measures of consumption inequality, along with income and substitution effects for consumption, remain therefore irrelevant for optimal income and savings taxes.

By contrast, the reduction to a single decision margin and single tax distortion characterized by Saez (2001) no longer applies when the preference restrictions of Atkinson and Stiglitz (1976) fail to hold. The design of optimal income taxes can no longer ignore how after-tax income is allocated to consumption and savings, since preferences for savings and incentives to work are no longer independent. If higher-ranked agents are more inclined to save, the taxation of savings facilitates redistribution towards lower-ranked agents and allows the social planner to reduce labor supply distortions. Hence, the optimal income tax is strictly lower than the static benchmark $\overline{\tau}_Y^{Saez}$, and *vice versa* if higher-ranked agents are less inclined to save and it becomes optimal to subsidize savings. The third part of Theorem 2 shows that this trade-off between income and savings taxes is especially stark when $s_C = 0$ and consumption vanishes at the top. In that case, the optimal top income tax is again governed by a static trade-off—but it is now between income and savings, rather than between income and consumption. The static wedge $\overline{\tau}_Y^{Saez}$ therefore equals the *combined* wedge between labor earnings and savings, as formalized by equation (16).

It follows from this discussion that, when $s_C < 1$, income and consumption inequality, as well as the income and consumption responses to tax changes, *both* matter independently for determining optimal income and savings taxes. If $s_C = 0$, then the income-Pareto ρ_Y and taxable income elasticities uniquely pin down $\overline{\tau}_Y^{Saez}$, while the moments from consumption data—specifically, the consumption-Pareto and the elasticity of intertemporal substitution—serve to decompose the combined wedge $\overline{\tau}_Y^{Saez}$ into $\overline{\tau}_Y$ and $\overline{\tau}_S$, as discussed in Section 2.3. More generally, Corollary 1 shows that savings should be taxed, and income should be taxed at a lower rate than the static wedge, iff the product of compensated elasticities and Pareto tail coefficients is larger for consumption than for savings, i.e., $\rho_C \zeta_C^H > \rho_S \zeta_S^H$. Using standard identities, this condition can equivalently be restated as $\rho_C \zeta_C^H > \rho_Y \zeta_Y^H$ as long as the marginal propensity to save does not go to zero.

Consequences for Policy Recommendations. The static optimal tax formula (1) of Saez (2001) features prominently in discussions about income tax design, while debates about savings taxes typically focus on the policy relevance of the uniform commodity taxation theorem of Atkinson and Stiglitz (1976). For example, a highly cited review article by Diamond and Saez (2011) simultaneously makes a case for high top income taxes based on the formula for $\overline{\tau}_Y^{Saez}$ and for positive savings taxes by questioning the assumptions underlying the Atkinson-Stiglitz theorem. We showed, however, that the static income tax formula and the Atkinson-Stiglitz theorem cannot be studied in isolation—one does not apply without the other. In other words, the two policy recommendations made by Diamond and Saez (2011) are mutually inconsistent: An expert calling for high income taxes based on $\overline{\tau}_Y^{Saez}$ must also support the recommendation of zero capital taxation, and support for positive savings taxes must be accompanied by less extreme recommendations for the income tax. More broadly, policy recommendations about income and savings taxes cannot be made in isolation from each other, but are part of a policy mix that optimally trades off between multiple margins of redistribution.

3 Quantitative Implications

3.1 Calibration

Pareto Tail Coefficients (ρ_Y, ρ_S, ρ_C). In Gaillard et al. (2023), we document that labor income and consumption both have Pareto tails and estimate their respective coefficients in the U.S. using the 2005 to 2021 waves of the PSID. We find an average estimate of the Pareto coefficient on labor income equal to $\rho_Y = 2.3$. Karahan, Ozkan, and Song (2022) find an even smaller value for the Pareto coefficient of lifetime earnings, $\rho_Y = 2.13$. By contrast, we obtain a much larger average estimate of the Pareto coefficient on consumption.[12] Figure 1, taken from Gaillard et al. (2023), plots the upper tail of the distribution of consumption for each wave of the data set, along with the best Pareto fit. We obtain an average value of the Pareto coefficient equal to $\rho_C = 3.1$, with a lower and an upper bound equal to 2.87 and 3.47, respectively. We further show that the gap between the two tails is statistically significant and robust across numerous specifications, formal tests, and alternative data sets. These results are also consistent with those obtained by Buda et al. (2022) using a rich dataset of credit card transactions in Spain. In light of these findings, there is little doubt that the relevant empirical scenario is *Case 3*, in which consumption has a substantially thinner tail than labor income, $\rho_C > \rho_Y$. In turn, this implies that the consumption

[12] The permanent income hypothesis suggests that it is preferable to use consumption rather than income data to calibrate the Pareto coefficient ρ_Y in the static Mirrlees setting, since annual consumption may be a better predictor of permanent income than annual income. While this only reinforces the critique we raised in the first section of this paper, according to which one could (and perhaps should) use consumption rather than income inequality data to evaluate optimal taxes in the static framework, this is not the main point of our paper. Instead, our argument is that, to the extent that (lifetime) income and consumption inequality measures do not coincide, they *both* matter independently for optimal taxes.

Figure 1: Pareto coefficient of consumption

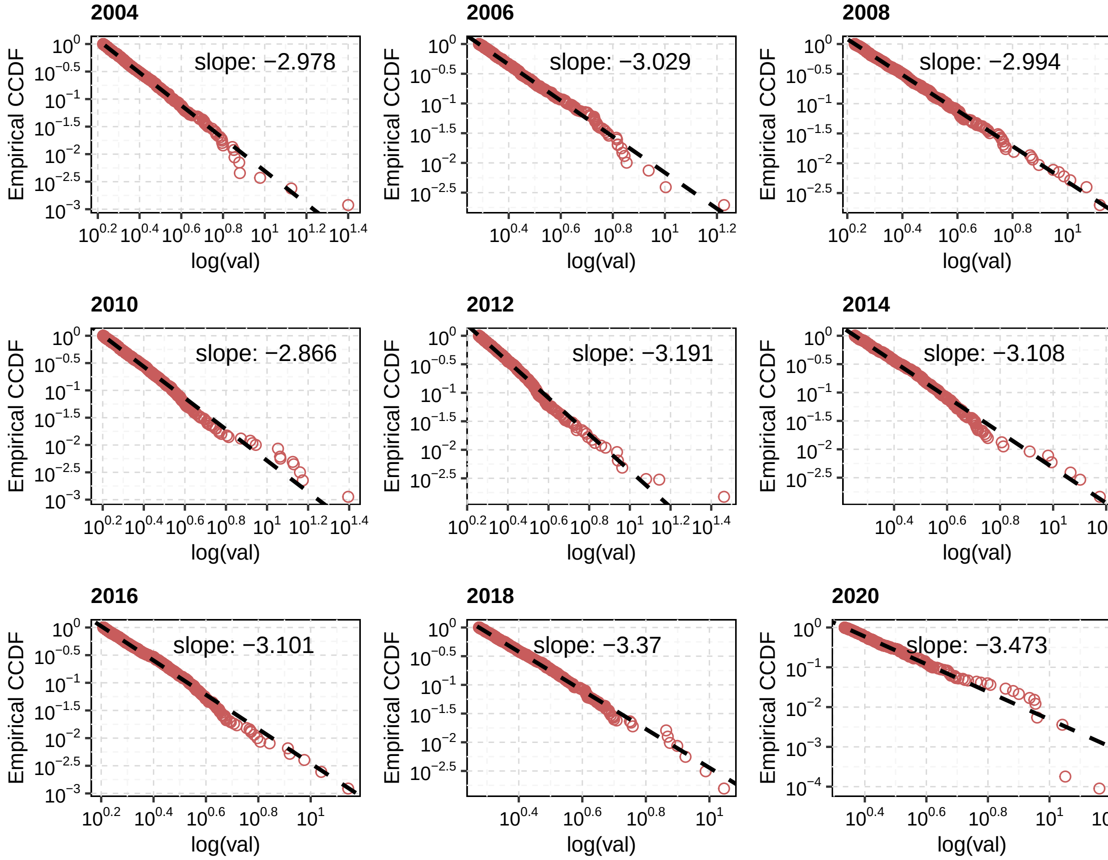

share of income s_C converges to 0 as $r \to 1$.[13]

Note, however, that our model imposes perfect co-monotonicity between labor income and consumption. To make our quantitative analysis consistent with our model, we compute the mean log-consumption of workers within each labor income quantile: By averaging, we remove the variation of consumption conditional on labor income rank. Each graph of Figure 2 plots the following quantiles: 70; 70.5; 71; . . . ; 99; 99.5. If the tail of labor income is indeed Pareto distributed, the fact that the data align along a straight line confirms that consumption is also Pareto-distributed. Moreover, the slope of the relationship yields an estimate of the ratio of Pareto coefficients ρ_Y/ρ_C between 0.55 and 0.71, with an average estimate of 0.66. In our quantitative exercise, we use these values as our lower bound, upper bound, and baseline estimate, respectively. Using our preferred estimate of $\rho_Y = 2.13$, these correspond to Pareto coefficients for consumption equal to $\rho_C \in \{3.0, 3.22, 3.87\}$.

[13] Recall moreover that $\rho_Y = \rho_S$ in our model. While this might appear counterfactual, note that ρ_S is the Pareto coefficient for *savings*, not wealth. As we show in Section 5.1, our model nests the case where asset returns are increasing in savings, implying that the wealth tail is strictly thicker than the savings—and hence the income—tail.

Figure 2: Ratio of Pareto Coefficients: Consumption vs. Income

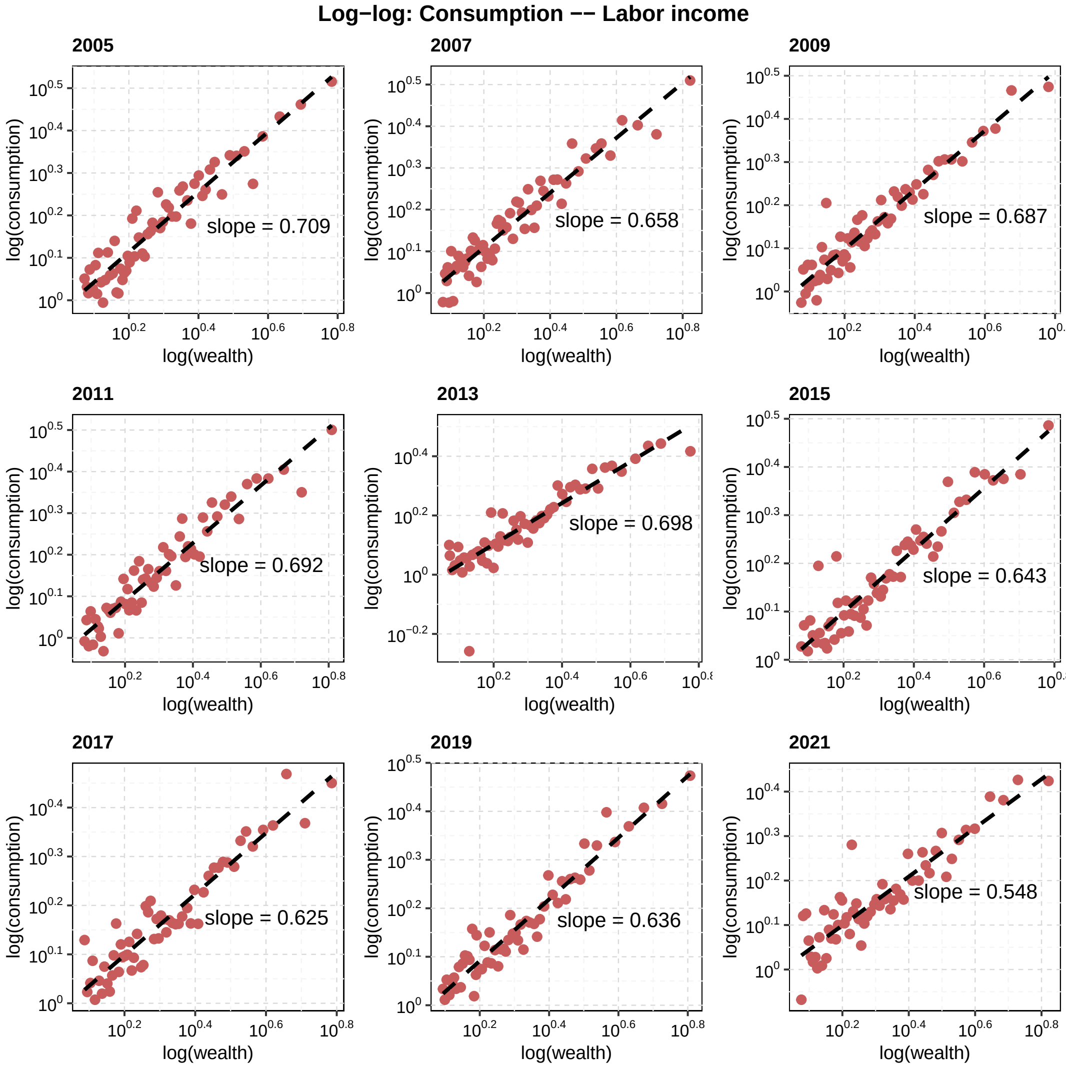

Labor Supply Elasticities (ζ_Y^H, ζ_Y^I). When $s_C = 0$, equation (12) establishes a one-to-one map between the substitution and income effects on labor supply (ζ_Y^H, ζ_Y^I) and the preference elasticities (ζ_Y, ζ_S). A vast literature estimates the elasticities of taxable income with respect to marginal tax rates and lump-sum transfers. The meta-analysis of Chetty (2012) yields a preferred estimate of the Hicksian elasticity of $\zeta_Y^H = 0.33$, while Gruber and Saez (2002) estimate a value of $\zeta_Y^H = 0.5$ for top income earners. Empirical evidence about the size of the income effects ζ_Y^I is mixed (see, e.g., Keane 2011). Gruber and Saez (2002) find negligible income effects, while Golosov, Graber, et al. (2021) estimate that \$1 of additional unearned income reduces the pre-tax income by 67 cents in the highest income quartile, which for a top marginal tax rate of 50% translates into an income effect of 0.33. For our baseline calibration, we choose $\zeta_Y^H = 0.33$ for the Hicksian elasticity and an intermediate value $\zeta_Y^I = 0.25$ for the income effect. These values imply $\tilde{\zeta}_Y = \zeta_Y = (1-\zeta_Y^I)/\zeta_Y^H = 2.27$ and $\tilde{\zeta}_M = \zeta_S = \zeta_Y^I/\zeta_Y^H = 0.76$, reasonable values for the inverse Frisch elasticity of labor supply and the relative risk aversion of top earners; in particular, the latter is consistent with the estimates of Chetty (2006). We evaluate the robustness of our quantitative results to the alternative parameter values $\zeta_Y^H = 0.5$ and $\zeta_Y^I = 0.33$.

Preference Complementarity (ζ_{CY}). Chetty (2006) shows that $\tilde{\zeta}_{MY}$ can be bounded as a function of the coefficient of relative risk aversion $\tilde{\zeta}_M$ by $\tilde{\zeta}_{MY}/\tilde{\zeta}_M \leq \Delta \ln C/\Delta \ln L$, where $\Delta \ln C/\Delta \ln L$ is the change in consumption that results from an exogenous variation in labor supply across various states, e.g., due to job loss or disability. This empirical moment turns out to be small, with $\tilde{\zeta}_{MY}/\tilde{\zeta}_M \leq 0.15$ representing a very conservative upper bound. While this finding is perfectly consistent with *Case 3* ($s_C = 0$), under which $\tilde{\zeta}_{MY}$ should converge to 0 for top earners (see equation (29) in the Appendix), it does not inform us about the underlying value of the preference complementarity ζ_{CY}. For our baseline calibration, we use the value $\zeta_{CY} = 0$ (separable utility), which is both a natural benchmark and consistent with the baseline estimate of the Marshallian elasticity of consumption with respect to permanent before-tax wage changes for men in Blundell, Pistaferri, and Saporta-Eksten (2016). We evaluate the robustness of our quantitative results to a wide range of values of the complementarity.

Elasticity of Intertemporal Substitution (ζ_{IS}). When $s_C = 0$, the preference elasticity ζ_C is related to the elasticity of intertemporal substitution ζ_{IS} through the relationship $\zeta_C = 1/\zeta_{IS}$. To calibrate the EIS, we use Jakobsen et al. (2020) who focus specifically on the behavior of the wealthiest households. Through the lens of a life-cycle model similar to ours, they show that an EIS as large as 2.6, and even higher for the very wealthy, is necessary to replicate the quasi-experimental evidence on the effects of a large wealth tax reform in Denmark on wealth accumulation. This value is consistent with that of Holm et al. (2024), who estimate an EIS of 2 by evaluating the spending response to a dividend tax news shock. Our baseline calibration thus uses $\zeta_{IS} = 2$, which in turn leads to $\zeta_C = 0.5$. As an alternative to the EIS, we could also use direct empirical estimates of the consumption response to a permanent wage shock to calibrate ζ_C. Attanasio and Davis (1996) estimate that between 60% and 80% of relative wage changes among birth cohorts and education groups are passed through to relative consumption changes. Using these estimates for the Marshallian (uncompensated) elasticity $\zeta_C^H + \zeta_C^I$, and the relationship $\zeta_C^I/\zeta_C^H = \zeta_Y$ when $\zeta_{CY} = 0$, we obtain a range of values $\zeta_C = \zeta_Y^I/\zeta_C^H \in (1.02, 1.36)$ in our baseline calibration, corresponding to an EIS in $(0.73, 0.98)$. We thus evaluate the robustness of our results to a wide range of values for the EIS.

A Cautionary Note. Our identification does not rely on any specific functional-form assumption on preferences. Instead, we infer the relevant preference elasticities directly from behavioral responses to income and tax changes. In particular, we do not calibrate our elasticity parameters using empirical estimates obtained by estimating structural models that rely on a particular choice of utility function if the latter is not flexible enough to accommodate the ordering of Pareto tails and the value of s_C to which we calibrate our model. By contrast, many papers impose strong *a priori* assumptions on the utility function to structurally estimate preference parameters. As emphasized by Chetty (2009), a potential pitfall of using these structural estimates to discipline behavioral elasticities and evaluate tax policies is that the structural model on which these empirical estimates are based may not be compatible with restrictions imposed by the underlying model that led to the optimal tax formula. For instance, suppose that we derived optimal taxes under the assumption that preferences take the form $U = u\left(g\left(C\right) - v\left(Y/\theta\left(r\right)\right)\right)$ for some concave constant-elasticity functions u and g and convex function v. While this utility function implies $U_{CY} \geq 0$ and appears to have sufficient flexibility to parametrize ζ_C, ζ_Y and ζ_{CY} via the functions u, g, and v, we can show that optimal consumption choices resulting from this functional form impose that s_C is bounded away from zero at the top; it can therefore not be consistent with the empirically relevant scenario of *Case 3*. A similar comment applies to the widely used estimates of French (2005), whose preference structure is compatible with s_C converging to zero only if labor supply is completely inelastic ($\zeta_Y \to \infty$) at the very top, contradicting empirical evidence on labor supply elasticities. One could make these preferences consistent with *Case 3* by incorporating additional preference heterogeneity—for example by making g rank-dependent. But by doing so, we would mechanically attribute the discrepancy between income and consumption Pareto tails to preference heterogeneity, and therefore hardwire into our preference assumptions the conclusion that it is optimal to tax savings, rather than letting the data identify how much of this gap is attributable to preference heterogeneity rather than income elasticities.

3.2 Results

Table 1 below summarizes our quantitative results for the optimal top tax rates on labor income and savings. While $\overline{\tau}_Y$ represents a marginal labor income tax on gross income, $\overline{\tau}_S$ represents the savings wedge as a proportion of net savings S. For constant top savings wedges, this translates into a top marginal tax on gross savings equal to $\overline{\tau}_S/(1+\overline{\tau}_S)$; this is the variable we report in the

Table 1: Optimal top labor income and savings taxes

	$\rho_Y/\rho_C = 0.55$		$\rho_Y/\rho_C = 0.66$		$\rho_Y/\rho_C = 0.71$		
	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y^{Saez}$
Baseline values***	58%	26%	**59%**	**24%**	60%	23%	**69%**
$\zeta_Y^H = 0.5$	49%	12%	50%	9%	51%	8%	55%
$\zeta_Y^I = 0.33$	55%	39%	57%	37%	57%	36%	73%
$\zeta_{IS} = 0.75$	68%	2%	72%	−10%	73%	−16%	69%
$\zeta_{CY} = 0.5$	47%	42%	48%	40%	48%	40%	69%

*Baseline values: $\rho_Y = 2.13$, $\zeta_Y^H = 0.33$, $\zeta_Y^I = 0.25$, $\zeta_{IS} = 2$, $\zeta_{CY} = 0$.

table. We also report the static optimum $\overline{\tau}_Y^{Saez}$. The first row gives the results for our baseline parameter values $\rho_Y = 2.13$, $\zeta_Y^H = 0.33$, $\zeta_Y^I = 0.25$, $\zeta_{CY} = 0$, $\zeta_{IS} = 2$, and for three values of the ratio of Pareto coefficients of labor income and consumption: $\rho_Y/\rho_C \in \{0.55, 0.66, 0.71\}$. The remaining rows of the table vary one parameter of our baseline calibration at a time.

Our preferred calibration yields a combined wedge $\overline{\tau}_Y^{Saez}$ of 69%, split between a top labor income tax rate of 59% and a top savings tax rate of 24%. To interpret the values of the savings wedge, it is useful to translate them into a tax on annualized returns. In our model, the first period represents a 30-year gap between the beginning of the working period and retirement. If the annual return on savings is 5% (resp., 3%), a savings tax of $\overline{\tau}_S/(1 + \overline{\tau}_S) = 24\%$ corresponds to a 0.9% annual tax on accumulated wealth, or a 19% (resp., 31%) capital income tax. Alternatively, if we interpret our model as one of retirement savings, a wedge of 24% means that top income earners will receive a present value of \$0.81 of additional pension payments for each additional dollar in social security contributions.

Higher values of the compensated elasticity of taxable income ζ_Y^H reduce the combined wedge on labor income and savings $\overline{\tau}_Y^{Saez}$, but also shift the distortions away from savings and towards labor income. Intuitively, a higher compensated income elasticity reduces the pass-through from income to consumption inequality, which in turn reduces the role of heterogeneity in preferences in accounting for the observed gap between the Pareto coefficients of income and consumption, and thus reduces the scope for savings taxes. Conversely, higher values of the income effect on labor supply ζ_Y^I raise the combined wedge, but also shift distortions from income towards savings: a higher ζ_Y^I maps one-for-one to higher risk aversion over after-tax income and thus a stronger redistribution motive, making both the under-lying tax distortions and the shift from income to savings taxes more salient. As a result, a higher ζ_Y^I slightly reduces the optimal income tax but significantly increases optimal savings taxes.

Changes to the elasticity of inter-temporal substitution ζ_{IS} and the complementarity between

Figure 3: Optimal Income and Savings Tax Rates

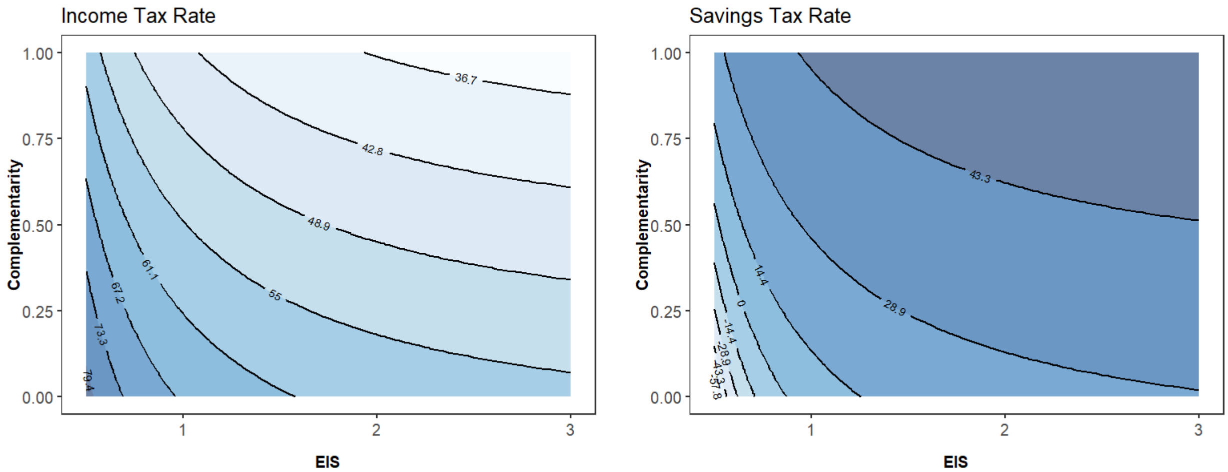

consumption and labor income ζ_{CY} leave the combined labor and savings wedge unchanged, but shift the break-down between labor and savings taxes. Equation (15) in Theorem 2 implies that savings should be taxed whenever the elasticity of intertemporal substitution ζ_{IS} is larger than $(\rho_Y/\rho_C)(\zeta_Y^H/\zeta_Y^I) = 0.87$ in our baseline calibration. The large EIS estimated by Jakobsen et al. (2020) and Holm et al. (2024) thus suggest that savings should optimally be taxed—and, correspondingly, labor income should be taxed at a strictly lower rate than the static optimum $\overline{\tau}_Y^{Saez}$. Nevertheless, lower values of the EIS, along with more unequal distributions of consumption, may lead to optimal savings subsidies, as shown in the fourth row of Table 1 for the case $\zeta_{IS} = 0.75$.

The Corlett-Hague rule implies that the planner should reduce the tax rate on labor income and raise the tax rate on savings when income and consumption are complements ($\zeta_{CY} > 0$). Quantitatively, a value of ζ_{CY} equal to 0.5 rather than 0 adds around 15 percentage points to the optimal savings tax rate while reducing the labor income tax by another 11 percentage points.

Figure 3 gives the values of the labor income and savings tax rates for a wide range of values of the EIS ζ_{IS} and complementarity ζ_{CY} around our baseline calibration.

4 Proof of Theorem 1

Before turning to generalizations of our baseline results, we provide a sketch of the proof of our main result, Theorem 1. We also give an alternative representation of the optimum, based on the idea of *redistributional arbitrage*, that formalizes the tradeoff between different tax distortions.

4.1 Step 1: Sufficient Conditions for Incentive Compatibility

We characterize the solution to the planner's problem of maximizing (6) subject to (7) and (8) in several steps using standard tools from mechanism design. First, we replace the set of global incentive compatibility conditions (7) by the local incentive compatibility or envelope condition

$$W'(r) = U_r(C(r), Y(r); r) + V_r(S(r); r) \tag{17}$$

where $W(r) \equiv U(C(r), Y(r); r) + V(S(r); r)$, and a co-monotonicity condition on allocations:

Proposition 1 (**Sufficient Condition for Global Incentive Compatibility**). *Suppose that Assumption 1 is satisfied. Then an allocation $\{C(r), Y(r), S(r)\}$ satisfies incentive compatibility (7) whenever it satisfies (i) local incentive compatibility (17); (ii) monotonicity of income: $Y(\cdot)$ is non-decreasing in r; and (iii) monotonicity of either savings or consumption, depending on the sign of $V_{Sr}/V_S - U_{Cr}/U_C$, namely: $S(\cdot)$ is non-decreasing in r if this term is strictly positive, and $C(\cdot)$ is non-decreasing in r if it is strictly negative. Conversely, an allocation $\{C(r), Y(r), S(r)\}$ that violates both (ii) and (iii) over some open interval of ranks $(r_1, r_2) \subset (0, 1)$ violates incentive compatibility (7).*

Proposition 1 shows that local incentive compatibility combined with co-monotonicity of income and either consumption or savings is sufficient for global incentive compatibility (7). These conditions are weaker than the co-monotonicity condition of Theorem 1. If $V_{Sr}/V_S = U_{Cr}/U_C$ for all r, condition (iii) holds automatically, and the monotonicity of $Y(\cdot)$ and local incentive compatibility are both necessary and sufficient for incentive compatibility—analogous to the well-known necessary and sufficient conditions in mechanism design problems with a single decision margin. If $V_{Sr}/V_S \neq U_{Cr}/U_C$ for some r, conditions (ii) and (iii) together are sufficient, but not necessary, for incentive compatibility. With two decision margins (labor supply and consumption-savings), incentive compatibility (7) holds if, for any pair of types, the sum of information rents across both margins is non-negative. Condition (ii) guarantees that information rents along the labor supply margin are always non-negative, while condition (iii) guarantees that the same holds along the consumption-savings margin. However, (7) may still hold if information rents are positive along only one of the two margins, i.e., if one of the two conditions is violated. The partial converse then says that incentive compatibility must be violated if both conditions are violated over some interval of types, that is, if information rents are negative along both margins.

4.2 Step 2: Solution to the Relaxed Planner's Problem

Next, we analyze the "relaxed" planner's problem of maximizing (6) subject to (8) and (17). We derive a representation of optimal income and savings wedges in terms of the model primitives, i.e., the agent's preferences. Proposition 2 provides a characterization of optimal labor and savings wedges akin to the well-known "ABC" formula (Diamond 1998). We derive it using a perturbation-based argument, and give a formal proof in the Appendix.

Proposition 2 (**Optimal Tax System**). *The optimal labor and savings wedges satisfy*

$$\frac{\tau_Y(r)}{1-\tau_Y(r)} = (1-r)\left(\frac{U_{Cr}(r)}{U_C(r)} - \frac{U_{Yr}(r)}{U_Y(r)}\right) U_C(r)\,\gamma(r) \tag{18}$$

and

$$\frac{\tau_S(r)}{1+\tau_S(r)} = (1-r)\left(\frac{V_{Sr}(r)}{V_S(r)} - \frac{U_{Cr}(r)}{U_C(r)}\right) U_C(r)\,\gamma(r) \tag{19}$$

where

$$\gamma(r) = \mathbb{E}\left[\frac{1}{U_C(r')}\exp\left(\int_r^{r'} \frac{U_{Cr}(r'')}{U_C(r'')}dr''\right)\middle| r' \geq r\right]. \tag{20}$$

Thus, Assumption 1 implies that $\tau_Y(r) > 0$ for all r, and $\tau_S(r) \lesseqgtr 0$ depending on the sign of (5).[14]

Equations (18) and (19) state that the optimal labor and savings wedges equalize the marginal resource gain from reducing labor or savings distortions at rank r, to the marginal resource cost from raising the information rent at rank r, thus hindering redistribution from ranks $r' \geq r$ to lower-ranked agents. More specifically, fix $r \in (0,1)$ and consider the following (infinitesimal) perturbation: simultaneously raise the consumption and the earnings (i.e., reduce the leisure) of rank r while keeping her utility unchanged, so that $U_C(r)\Delta C(r) = -U_Y(r)\Delta Y(r)$. On the one hand, the net resource gain associated with this perturbation is $\Delta Y(r) - \Delta C(r) = (\tau_Y(r)/(1-\tau_Y(r)))\Delta C(r)$, corresponding to the left-hand side of equation (18). On the other hand, this perturbation makes it strictly more attractive for ranks $r' > r$ to mimick rank r, since by Assumption 1 they are more inclined than rank r to give up a given amount of leisure for an extra unit of consumption. Formally, the information rent $W'(r)$ at rank r (i.e., the additional utility that must be awarded to agents with rank slightly above r to deter them from mimicking rank r) increases by $\Delta W'(r) = U_{Cr}(r)\Delta C(r) + U_{Yr}(r)\Delta Y(r) = [U_{Cr}(r)/U_C(r) - U_{Yr}(r)/U_Y(r)]\,U_C(r)\Delta C(r)$. The planner's marginal resource cost of ceding these additional information rents is then obtained

[14] We use Assumption 1 only to guarantee that local incentive compatibility and monotonicity are sufficient for global incentive compatibility. The result that τ_S inherits the sign of $V_{Sr}/V_S - U_{Cr}/U_C$ for each r applies even without Assumption 1 provided that incentive compatibility is satisfied.

by multiplying $\Delta W'(r)$ with the "shadow value of redistribution", derived in the next paragraph, $(1-r)\gamma(r)$. Equating the marginal benefit and cost of the perturbation yields equation (18). A similar perturbation of $\{\Delta C(r), \Delta S(r)\}$ of consumption and savings delivers equation (19).

It remains to validate the interpretation of $(1-r)\gamma(r)$ as the shadow value of redistribution at rank r. Consider a perturbation $\{\Delta C(r')\}$ that raises the consumption of ranks $r' > r$ while preserving local incentive compatibility. Let $\Delta W(r') = U_C(r')\Delta C(r')$ denote the associated increase in welfare for types $r' > r$, and $\Delta W'(r') = U_{Cr}(r')\Delta C(r')$ the increase in information rent at r'. Therefore, we have $\Delta W'(r') = (U_{Cr}(r')/U_C(r'))\Delta W(r')$. That is, the change in utility at rank r' causes a change in information rents that must be passed on to the utility of all higher ranks r'', thus further changing information rents, etc. Integrating up this ODE yields the cumulative utility changes for higher ranks that are required as a result of preserving local incentive compatibility at all lower ranks:[15] $\Delta W(r') = \Delta W(r) \cdot \exp \int_r^{r'} [U_{Cr}(r'')/U_C(r'')]dr''$. To interpret this expression, suppose that higher ranks have lower consumption needs, i.e., $U_{Cr} < 0$. Then the utility of higher ranks does not need to increase by as much as that of lower ranks to maintain incentive compatibility, because the higher level of consumption at rank r is not so attractive for higher ranks $r' > r$; in this case, a smaller increase in utility at r' is sufficient to deter them from mimicking lower ranks. Finally, this expression allows us to compute $\Delta C(r') = \Delta W(r')/U_C(r')$ and hence the resource cost of the perturbation, $(1-r)\mathbb{E}[\Delta C(r')|r' \geq r] = (1-r)\gamma(r)\Delta W(r)$. This concludes the proof.

4.3 Step 3: Identification

Our third and final step consists of mapping the model primitives that appear on the right-hand sides of equations (18) and (19) to empirically observable sufficient statistics—standard elasticities, Pareto coefficients, and measures of tax progressivity. We do so in Proposition 3, which generalizes Lemma 1 in Saez (2001) to our economy. We discussed its interpretation—in particular, the role that the Pareto coefficients and preference elasticities play in the identification of the intra- and inter-temporal MRS—in the text following the statement of Theorem 1.

We can then easily complete the proof of Theorem 1 by substituting these expressions into those of Proposition 2, taking the limit as $r \to 1$, and invoking the assumption that the relevant sufficient statistics all converge to finite limits.

[15] In the special case where $U_{Cr} = 0$, i.e., consumption utilities are independent of type, the perturbation raises the utility of ranks $r' > r$ *uniformly* by $\Delta W(r') = \Delta W(r) > 0$. It is well known that this perturbation of consumption preserves incentive compatibility for all $r' > r$. That is, the increase in the consumption of rank r' does not modify $U_r(r')$, and hence does not require any additional change in utility for types above r'. By contrast, with general non-separable preferences $U_{Cr} \neq 0$, a uniform increase in utility no longer preserves local incentive compatibility.

Proposition 3 (Identification). *For any given (observed) system of tax distortions $\{\hat{\tau}_Y(r), \hat{\tau}_S(r)\}$, the variables $U_{Cr}/U_C - U_{Yr}/U_Y$ and $V_{Sr}/V_S - U_{Cr}/U_C$ can be expressed as*

$$(1-r)\left(\frac{U_{Cr}(r)}{U_C(r)} - \frac{U_{Yr}(r)}{U_Y(r)}\right) = \frac{\zeta_C(r)}{\rho_C(r)} + \frac{\zeta_Y(r)}{\rho_Y(r)} - \frac{\zeta_{CY}(r)}{\rho_Y(r)}\left(1 + s_C(r)\frac{\rho_Y(r)}{\rho_C(r)}\right) + \frac{d\ln(1-\hat{\tau}_Y(r))}{d\ln(1-r)} \tag{21}$$

and

$$(1-r)\left(\frac{V_{Sr}(r)}{V_S(r)} - \frac{U_{Cr}(r)}{U_C(r)}\right) = \frac{\zeta_S(r)}{\rho_S(r)} - \frac{\zeta_C(r)}{\rho_C(r)} + \frac{\zeta_{CY}(r)}{\rho_Y(r)} - \frac{d\ln(1+\hat{\tau}_S(r))}{d\ln(1-r)}. \tag{22}$$

Moreover, the shadow value of redistribution can be expressed as:

$$U_C(r)\gamma(r) = \mathbb{E}\left[\exp\left(\int_r^{r'}\left\{\frac{\zeta_C(r'')}{\rho_C(r'')} - \frac{\zeta_{CY}(r'')}{\rho_Y(r'')}\right\}\frac{1}{1-r''}dr''\right)\middle| r' \geq r\right]. \tag{23}$$

Equations (21), (22), and (23) follow by totally differentiating $1-\hat{\tau}_Y(r) = -U_Y(r)/U_C(r)$, $1+\hat{\tau}_S(r) = V_S(r)/U_C(r)$, and $U_C(r)$, respectively, and noting that in each case the differentiation can be decomposed into a component that captures the rank-dependence of preferences for a given allocation, and a component that captures the variation in allocations at a given rank. The latter is fully identified from preference elasticities and local Pareto tail coefficients, and can thus be used to identify the former, which is given by the left-hand side of (21) to (23).

4.4 Redistributional Arbitrage

Equation (20) expresses the shadow value of redistribution $(1-r)\gamma(r)$ in terms of redistributing *consumption* from agents ranked $r' > r$ towards lower-ranked ones. Applying analogous perturbation arguments as above, we obtain two additional, equivalent representations of $\gamma(r)$ based on redistributing earnings (or leisure, that is, asking the top-ranked agents to work more) or savings from higher to lower ranks:

$$\gamma(r) = \mathbb{E}\left[\frac{1}{-U_Y(r')}\exp\left(\int_r^{r'}\frac{U_{Yr}(r'')}{U_Y(r'')}dr''\right)\middle| r' \geq r\right] = \mathbb{E}\left[\frac{1}{V_S(r')}\exp\left(\int_r^{r'}\frac{V_{Sr}(r'')}{V_S(r'')}dr''\right)\middle| r' \geq r\right].$$

Combined with equation (20), these representations articulate the presence of three margins of redistribution: consumption, earnings (or leisure), and savings. The optimal allocation must then obey a principle of *"no redistributional arbitrage"*, according to which there are no gains from shifting redistribution from one margin (say, consumption) to a different one (say, leisure or savings). Taking ratios of $-U_Y(r)\gamma(r)$, $V_S(r)\gamma(r)$, and $U_C(r)\gamma(r)$, we obtain the following alternative, yet equivalent, representation of optimal labor and savings taxes:

$$1-\tau_Y(r) \;=\; \frac{B_Y(r)}{B_C(r)} \;\equiv\; \frac{\mathbb{E}\left[\frac{U_Y(r)}{U_Y(r')}\exp\left(\int_r^{r'}\frac{U_{Yr}(r'')}{U_Y(r'')}dr''\right)\mid r'\geq r\right]}{\mathbb{E}\left[\frac{U_C(r)}{U_C(r')}\exp\left(\int_r^{r'}\frac{U_{Cr}(r'')}{U_C(r'')}dr''\right)\mid r'\geq r\right]} \tag{24}$$

and

$$1+\tau_S(r) \;=\; \frac{B_S(r)}{B_C(r)} \;\equiv\; \frac{\mathbb{E}\left[\frac{V_S(r)}{V_S(r')}\exp\left(\int_r^{r'}\frac{V_{Sr}(r'')}{V_S(r'')}dr''\right)\mid r'\geq r\right]}{\mathbb{E}\left[\frac{U_C(r)}{U_C(r')}\exp\left(\int_r^{r'}\frac{U_{Cr}(r'')}{U_C(r'')}dr''\right)\mid r'\geq r\right]}. \tag{25}$$

These optimal tax formulas formalize the idea that, at the optimal allocation, the planner is indifferent between redistributing slightly less along one margin of inequality—consumption, leisure, or wealth—and slightly more along another. The variables B_C, B_Y and B_S represent the marginal (resource) benefits of reducing the consumption, leisure, and savings of agents with rank above r, respectively. Thus, the ratio B_Y/B_C describes the trade-off between redistributing resources from the top via income or via consumption—or in other words, how the social planner maximizes the extraction of resources from top earners by asking them to work more versus reducing their consumption. Similarly, the ratio B_S/B_C describes the trade-off between redistributing consumption or savings. Comparing equations (24) and (25) with the individual's first-order conditions $1-\tau_Y = -U_Y/U_C$ and $1+\tau_S = V_S/U_C$ then leads to the following interpretation of optimal taxes: The optimal income (resp., savings) wedge equalizes the agent's private trade-off between consumption and leisure (resp., savings), to the social trade-off in redistributing from the top via consumption or leisure (resp., savings). Sufficient statistics representation of these expressions can then be obtained as in Proposition 3, which would then lead us to recover the formulas given in Theorem 1.[16] The interpretation of optimal top income and savings taxes as equalizing the shadow values of redistributing leisure, savings, and consumption clarifies why all of these margins of inequality matter independently for the design of redistributive policies.

5 Extensions

In this last section of our paper, we show that our baseline results can be easily extended to much richer settings, including heterogeneous rates of return, multiple-period, multiple-good, life-cycle, and dynamic stochastic economies. We provide formal derivations and study an additional

[16] It is straight-forward to check that $-U_Y(r)\gamma(r) = \mathbb{E}_{r'\geq r}[\exp\int_r^{r'}\{-\frac{\zeta_Y(r'')}{\rho_Y(r'')} + s_C(r'')\frac{\zeta_{CY}(r'')}{\rho_C(r'')}\}\frac{dr''}{1-r''}]$ and $V_S(r)\gamma(r) = \mathbb{E}_{r'\geq r}[\exp\int_r^{r'}\frac{\zeta_S(r'')}{\rho_S(r'')}\frac{dr''}{1-r''}]$. Therefore, $\lim_{r\to1}B_Y(r) = [1+\zeta_Y/\rho_Y - s_C\zeta_{CY}/\rho_C]^{-1}$, $\lim_{r\to1}B_S(r) = [1-\zeta_S/\rho_S]^{-1}$, and $\lim_{r\to1}B_C(r) = [1-\zeta_C/\rho_C+\zeta_{CY}/\rho_Y]^{-1}$.

extension (heterogeneous initial capital) in the Appendix.

5.1 Return Heterogeneity

Recent empirical evidence suggests that heterogeneous rates of return, whereby wealthier agents earn higher returns on their savings, are an important component of the observed concentration of wealth at the top; see, e.g., Bach, Calvet, and Sodini (2020), Fagereng et al. (2020), and Gaillard et al. (2023). There are two potential sources of such heterogeneity: scale-dependence (returns increase with wealth, regardless of an individual's rank r) and type-dependence (returns increase with an individual's exogenous rank r, for any level of wealth). In the Appendix, we show that the generic utility function $V(S;r)$ introduced in our baseline framework of Section 2 nests both cases. To see this, interpret $V(\cdot;r)$ as an indirect utility function over *initial* savings, rather than over second-period consumption. Specifically, define $V(S,r) = \beta(r)\, v(R(S;r)\, S)$, where $R(S;r)$ denotes the returns on savings, which can be scale-dependent through their dependence on S or type-dependent through their dependence on r. Second-period consumption is then given by $C_2(r) = R(S;r)\, S(r)$.

This argument implies that our optimal tax formulas continue to hold, except that the relevant savings elasticity ζ_S and Pareto coefficient ρ_S should be those of initial savings. In particular, as explained in Section 2.3, we have $\rho_S = \rho_Y$ by construction. Since $1/\rho_{C_2} = 1/\rho_S + 1/\rho_R$, where ρ_{C_2} and ρ_R denote respectively the Pareto coefficients on second-period consumption and rates of return, we obtain that wealth has a strictly thicker tail than labor income—as is the case in the data. Note finally an important advantage of the calibration in Section 3 of top income and savings taxes: It identifies the sufficient statistic ζ_S directly from income and substitution effects on labor supply, without taking a stand on return heterogeneity. That is, conditional on the usual labor supply elasticities ζ_Y^H, ζ_Y^I, the expressions for optimal taxes we derived above hold for any underlying heterogeneity in rates of return, and any combination of type- and scale-dependence.[17]

[17] Return heterogeneity would instead enter the characterization of ζ_S in terms of primitives. It is straightforward to check that $\zeta_S = \zeta_{C_2} - \eta(1-\zeta_{C_2})$, where $\zeta_{C_2} = -C_2 v''(C_2)/v'(C_2)$ and $\eta = SR_S(S;r)/R(S;r)$ is the scale-dependence parameter. Hence, scale dependence of returns affects the savings elasticity ζ_S through the parameter η whenever $\zeta_{C_2} \neq 1$. Increasing returns to savings ($\eta > 0$) lower ζ_S and thus optimal savings taxes when $\zeta_{C_2} < 1$, and increase ζ_S and optimal savings taxes when $\zeta_{C_2} > 1$. The opposite result holds if savings have decreasing returns ($\eta < 0$). Note finally that our optimal tax formulas use the Pareto coefficient of *savings* $\rho_S = \rho_Y$ as a sufficient statistic. Since type dependence only affects the Pareto coefficient of ex-post wealth ρ_{C_2} (via ρ_R), it does not affect optimal taxes. By contrast, optimal tax formulas expressed in terms of the Pareto coefficient of *wealth* as the relevant sufficient statistic, such as Schulz (2021), are explicitly affected by type-dependence.

5.2 General Preferences and Multiple Commodities

In our baseline model of Section 2, we assumed that preferences were additively separable between "savings", on the one hand, and "consumption" and "labor income", on the other hand. We now extend Theorem 1 to general preferences $U(C_1, \ldots, C_N, Y; r)$ over efficiency units of labor Y and an arbitrary set of consumption goods $C_1, \ldots, C_N$, nesting our baseline model as a special case with two consumption goods. The separability assumption imposed some structure on income and substitution effects of the different commodities, which simplified the identification of sufficient statistics leading to Theorem 1: The computation of the top income and savings taxes required estimates of four preference parameters—three elasticities and an adjustment for complementarity between consumption and income.

Suppose that agents' preferences are defined as $U(\boldsymbol{C}, Y; r)$, where $\boldsymbol{C}$ is an N-dimensional commodity vector and $r \in [0,1]$. Let $U_n \equiv \partial U/\partial C_n$ and $U_{rn} \equiv \partial U_r/\partial C_n$ and assume that U_{rn}/U_n is increasing in n. Hence, U_m/U_n is increasing in r whenever $m > n$. As before, suppose that $U_Y \equiv \partial U/\partial Y < 0$, and $U_{rY}/U_Y < U_{rn}/U_n$ for all $n = 1, ..., N$. It will be convenient write $U_{r0} \equiv U_{rY}$ and $U_0 \equiv U_Y$, i.e., to index efficiency units of labor as the lowest-indexed good 0.

The planner can produce good n at a constant marginal cost of p_n efficiency units of labor. We consider a general social welfare objective, with rank-dependent Pareto weights $\omega(\cdot)$ and a concave Bergson-Samuelson function $G(\cdot)$ that captures the planner's aversion to inequality. The planner's problem reads

$$\max_{\boldsymbol{C}(\cdot), Y(\cdot)} \int_0^1 \left[\omega(r) G(U(\boldsymbol{C}(r), Y(r); r)) + Y(r) - \sum_{n=1}^{N} p_n C_n(r) \right] dr$$

subject to the agents' incentive compatibility constraints. Let $\hat{\omega}(r) \equiv \omega(r) G'(U(r))$ denote the marginal social welfare weight on rank r. The optimal wedge of any good between any pair of goods m and n (including labor, $n = 0$) is defined by $1 + \tau_{m,n}(r) = (U_m(r)/p_m)/(U_n(r)/p_n)$ and takes the form

$$1 + \tau_{m,n}(r) = \frac{B_m(r)}{B_n(r)} = \frac{\mathbb{E}\left[\left\{1 - p_m^{-1} U_m(r') \hat{\omega}(r')\right\} \frac{U_m(r)}{U_m(r')} \exp\left(\int_r^{r'} \frac{U_{mr}(r'')}{U_m(r'')} dr'' \right) \middle| r' \geq r \right]}{\mathbb{E}\left[\left\{1 - p_n^{-1} U_n(r') \hat{\omega}(r')\right\} \frac{U_n(r)}{U_n(r')} \exp\left(\int_r^{r'} \frac{U_{nr}(r'')}{U_n(r'')} dr'' \right) \middle| r' \geq r \right]}. \quad (26)$$

We further have that $\tau_{m,n}(r) > 0$ whenever $m > n$. This characterization also applies to the optimal wedge between efficiency units of labor and any consumption good n, for which we have $1 + \tau_{0,n}(r) = B_0(r)/B_n(r) < 1$ and therefore $\tau_{0,n}(r) < 0$. That is, leisure should be subsidized relative to other consumption goods, and correspondingly, labor income be taxed.

This expression has the same structure and interpretation as equations (24) and (25), except that the welfare weight $\hat{\omega}(r') U_n(r')$ modifies the expression of each $B_n(r)$; this term account for the fact that, when the planner is not Rawlsian, perturbing the allocation as described in Section 4 has a direct effect on social welfare in addition to the resource costs and benefits. As described in Section 4.4, $B_n(r)$ represents the marginal benefits of reducing the consumption of commodity n for ranks above r while preserving incentive-compatibility for $r' \geq r$.[18] Formula (26) thus characterizes the optimal relative price distortions as arbitraging between redistribution through one commodity vs. another.

We can finally identify $\lim_{r\to 1} B_n(r)$ in terms of observables, following a similar proof as that of Proposition 3. Assuming that the preference elasticities $\zeta_{nk}(r) \equiv C_k(r) U_{nk}(r) / U_n(r)$, $\zeta_n(r) \equiv -C_n(r) U_{nn}(r) / U_n(r)$, $\zeta_{n0}(r) \equiv Y(r) U_{nY}(r) / U_n(r)$, and $\zeta_0(r) \equiv Y(r) U_{YY}(r) / U_Y(r)$, the spending shares $s_n(r) \equiv U_n(r) C_n(r) / (-U_Y(r) Y(r))$ and the local tail coefficients $\rho_k(r) \equiv -\partial \ln C_k(r) / \partial \ln(1-r)$ converge to constants ζ_{nk}, ζ_n, ζ_{n0}, ζ_0, s_n, and $\rho_k > 1$ as $r \to 1$, we obtain

$$\lim_{r\to 1} B_n(r) = \left[1 + \sum_{k=0}^{N} \frac{\zeta_{nk}}{\rho_k}\right]^{-1} \text{ and } \lim_{r\to 1} B_0(r) = \left[1 + \frac{\zeta_0}{\rho_0} - \sum_{k=1}^{N} s_k \frac{\zeta_{k0}}{\rho_k}\right]^{-1} \tag{27}$$

Although the terms $B_n(r)$ are expressed in terms of cardinal preference elasticities, we show that they only depend on ordinal preferences; that is, if $\hat{U}$ is a monotone transformation of U, then they generate the same marginal benefits of redistribution $B_n(r)$. It follows that these terms can be fully identified from the income and substitution effects of the various goods. Specifically, with $N+1$ different commodities (including efficiency units of labor in the count), we have N independent income effects ζ_n^I and $N(N+1)/2$ independent substitution effects $\zeta_{n,k}^H$ that can—in principle—be estimated as behavioral responses to income, price or tax changes.[19] On the other hand, there are $(N+1)^2$ preference elasticities in $\{\zeta_n, \zeta_{nk}; \zeta_0, \zeta_{n0}\}$, with $(N+1)^2 - N(N+1)/2$ degrees of freedom, since the off-diagonal elements must satisfy $s_n \zeta_{nk} = s_k \zeta_{kn}$. Hence, income and substitution effects $\{\zeta_n^I; \zeta_{n,k}^H\}$ together identify the preference elasticities $\{\zeta_n, \zeta_{nk}; \zeta_0, \zeta_{n0}\}$ up to a single degree of freedom, which accounts for the fact that marginal benefits of redistribution and optimal wedges at the top are invariant to monotone transformations of cardinal utilities and only

[18] Note that the Inada conditions ensure that the term in brackets converges to 1 at the top of the type distribution: If $\lim_{r\to 1} \hat{\omega}(r) U_n(r) = 0$ for any n, we recover the Rawlsian representation of Section 4.4. In the Appendix, we show that the relative price of goods m and n should be undistorted everywhere, i.e., it is optimal to tax the two goods uniformly, if and only if the marginal rate of substitution $U_m(r) / U_n(r)$ is uniform across preference ranks r. More generally, it is optimal to tax good m at a higher rate than good n, so that $\tau_{m,n}(r) > 0$ for all r, whenever $U_m(r) / U_n(r)$ increases with r.

[19] The Slutsky matrix has $(N+1)^2$ entries, but symmetry of the off-diagonal elements imposes $N(N+1)/2$ restrictions and the usual adding-up constraint on substitution effects for a given price change another $N+1$ restrictions, leaving us with $N(N+1)/2$ degrees of freedom.

depend on ordinal preferences.

5.3 Income and Savings Taxes over the Life Cycle

As an application of the general framework of the previous section, we now study how income and savings taxes should vary over the life cycle. Consider a Mirrleesian economy in which households work and consume over a fixed number of periods, indexed by $t = 1, ..., T$. Their initial preference rank is drawn prior to date $t = 1$, and is their private information. The households' preferences are given by $\sum_{t=1}^{T} \beta^t U_t (C_t, Y_t; r)$, where the within-period utility function is allowed to vary deterministically over time (for example to capture age-dependence of preferences over consumption or work productivity), but otherwise satisfies the same restrictions as in our baseline economy. The age-dependent labor and savings taxes on top earners are then given by:

$$\overline{\tau}_Y (t) = \frac{\zeta_{Y_t}/\rho_{Y_t} + \zeta_{C_t}/\rho_{C_t} - \zeta_{C_tY_t}/\rho_{Y_t}(1 + s_{C_t}\rho_{Y_t}/\rho_{C_t})}{1 + \zeta_{Y_t}/\rho_{Y_t} - s_{C_t}\zeta_{C_tY_t}/\rho_{C_t}}$$

and

$$\overline{\tau}_S (t) \;=\; \frac{\zeta_{C_{t+1}}/\rho_{C_{t+1}} - \zeta_{C_t}/\rho_{C_t} - (\zeta_{C_{t+1}Y_{t+1}}/\rho_{Y_{t+1}} - \zeta_{C_tY_t}/\rho_{Y_t})}{1 - \zeta_{C_{t+1}}/\rho_{C_{t+1}} + \zeta_{C_{t+1}Y_{t+1}}/\rho_{Y_{t+1}}}.$$

These formulas are analogous to those found in the two-period environment, but they are now based on age-specific rather than unconditional preference elasticities and Pareto tail coefficients. Following the same procedure as described in Section 3.1, we compute the age-specific Pareto coefficients for consumption and income, as well as their ratio, by birth cohort from different PSID waves, and then plot them against age in Figure 4. The decline of these estimates with age until retirement illustrate the growth of income and consumption inequality over the first half of the life cycle. Their ratio is remarkably stable across ages, with values between 0.75 and 0.8.[20]

What do these age-specific Pareto coefficients imply for the evolution of income and savings taxes? Assuming that the preference parameters do not vary too much with age, the rising income inequality over the life cycle suggests that income taxes should be increasing with age. At the same time, the fact that age-specific Pareto coefficients are uniformly lower than their unconditional counterpart also results in uniformly lower income taxes. Using $\zeta_{Y_t}^{-1} = 0.44$, $\zeta_{C_t} = 0.75$, $\zeta_{C_tY_t} = 0$, along with $\rho_{C_t}/\rho_{Y_t} = 0.75$ for all t yields top optimal labor income taxes that increase from $\bar{\tau}_Y (t) = 60.5\%$ at age 20 to 68.5% for ages 50 and above. Moreover, the gradual increase in

[20]Our estimates of the Pareto coefficients for income by age are consistent with those found by Karahan, Ozkan, and Song (2022).

Figure 4: Pareto Coefficients conditional on Age

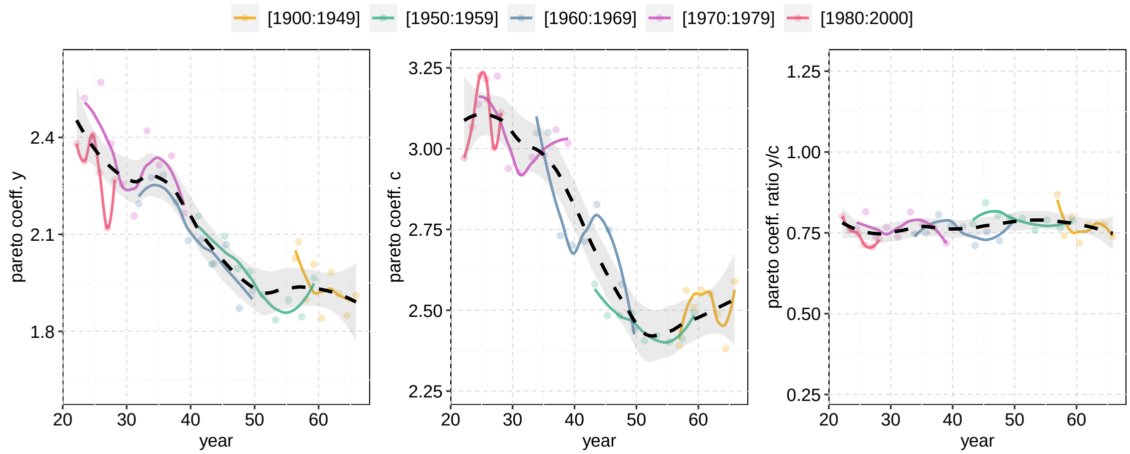

consumption inequality introduces a rationale for back-loading redistribution, or taxing savings. The previous calibration implies a cumulative savings tax between ages 20 and 50 of 7.7%, or equivalently about 0.25% per annum, before dropping to 0 beyond age 50. These estimates are smaller than in our baseline economy, but stem from an entirely different channel, namely the growth in income and consumption inequality with age, rather than the difference between consumption and income inequality in the cross-section.

5.4 Inverse Euler Equation

We can finally link our results to the "Inverse Euler Equation" that emerges in dynamic Mirrleesian economies with stochastically evolving types (Golosov, Kocherlakota, and Tsyvinski 2003; Farhi and Werning 2013; Golosov, Troshkin, and Tsyvinski 2016) by re-interpreting the second-period utility function V in our model. Except for Hellwig (2021), this literature abstracts from both heterogeneity in preferences for savings and complementarities between consumption and labor, which are the two key channels that drive savings taxes in our setting.

Suppose that agents' preferences over second-period consumption C_2 and second-period income Y_2 are given by $\beta v(C_2, Y_2; r_2)$, where the second period rank $r_2 \in [0,1]$ is uniform and i.i.d. across agents and independent of the first period rank r. First-period savings S generate a return $R > 0$. The social planner then sets second-period allocations $\{C_2(\cdot), Y_2(\cdot)\}$ to maximize

$$V(S) \equiv \beta \int_0^1 v(C_2(r_2), Y_2(r_2); r_2)\, dr_2$$

subject to the break-even constraint

$$RS \;\geq\; \int \left(C_2\left(r_2\right) - Y_2\left(r_2\right)\right) dr_2$$

and incentive-compatibility constraints

$$v\left(C_2\left(r_2\right), Y_2\left(r_2\right), r_2\right) \;\geq\; v\left(C_2\left(r_2'\right), Y_2\left(r_2'\right), r_2\right)$$

for all $r_2, r_2' \in [0,1]$. That is, our second-period utility function $V\left(S\right)$ simply stands for the expected present value of future utility. We can then characterize exactly as in our baseline model the optimal labor distortions in both periods (equation (24)), and the wedge between first-period consumption and savings, $\tau_S\left(r\right)$ (equation (25)). In addition, a perturbation argument analogous to that of Section 4 implies that

$$\frac{1}{\left(1+\tau_S\left(r\right)\right) U_C\left(r\right)} \;\equiv\; \frac{1}{V_S\left(S\right)} \;=\; \mathbb{E}\left[\frac{1}{\beta R v_{C_2}\left(r_2\right)} \cdot \frac{\exp\left(\int_0^{r_2}\{v_{C_2 r}\left(r'\right)/v_{C_2}\left(r'\right)\}dr'\right)}{\mathbb{E}\left[\exp\left(\int_0^{r_2}\{v_{C_2 r}\left(r'\right)/v_{C_2}\left(r'\right)\}dr'\right)\right]}\right]. \tag{28}$$

Thus, the inverse marginal utility of savings $1/V_S\left(S\right)$ is equal to an expected inverse marginal utility of second-period consumption, weighted by an adjustment factor that accounts for the non-separability of preferences and is analogous to that derived in Proposition 2. Combining this expression for $V_S\left(S\right)$ with our characterization of the savings wedge $1-\tau_S\left(r\right) = B_S\left(r\right)/B_C\left(r\right)$ thus yields a generalization of the Inverse Euler Equation.

In other words, our characterization of optimal savings wedges naturally extends to a dynamic Mirrleesian economy, which now combines two separate rationales for taxing savings: the heterogeneity in inter-temporal marginal rates of substitution or departure from uniform commodity taxation that is captured by $\tau_S\left(r\right)$, and the adverse incentive effect of savings that the Inverse Euler equation emphasizes by characterizing the marginal value of savings as a harmonic expectation of second-period marginal utilities. Specifically, to preserve incentive compatibility in the second period, returns to savings are reweighted by an adjustment factor that is proportional to inverse marginal utilities $1/v_{C_2}$, and a further adjustment to factor in non-separability of preferences akin to the adjustment factors in the static marginal benefits of redistribution formulas (24)-(25).[21] That said, recall that, as for the extension in Section 5.1, ζ_S is identified directly from income and substitution effects on labor supply in period 1, without taking a stand on preferences and labor

[21] Equation (5.4) can equivalently be written as $U_C\left(r\right) = \{\beta R/\left(1+\tau_S\left(r\right)\right)\}\mathbb{E}[v_{C_2}\left(r_2\right) \cdot M\left(r_2\right)]/\mathbb{E}[M\left(r_2\right)]$, where $M\left(r_2\right) = \left(1/v_{C_2}\left(r_2\right)\right) \cdot \exp(\int_0^{r_2}\left(v_{C_2 r}\left(r'\right)/v_{C_2}\left(r'\right)\right)dr')$. In this representation, $M\left(r_2\right)$ represents the type-contingent returns to savings that are required to preserve incentive compatibility of a small perturbation to consumption-savings decisions.

productivities in period 2. Conditional on ζ_Y^H, ζ_Y^I, the expressions for optimal taxes we derived in Section 2 therefore continue to apply.

The present discussion was kept deliberately simple by assuming that ranks were i.i.d. across time and across agents. Hellwig (2021) analyzes a dynamic Mirrleesian economy with arbitrary Markovian shock processes and non-separable preferences that integrates motives for savings taxes due to preference heterogeneity with adverse incentive effects of savings.[22] The key difference with the representation of Theorem 1 is that the sufficient statistics required to compute optimal taxes are now based on the Pareto coefficients of income, consumption and savings *conditional on* the entire prior sequence of types, or equivalently the entire earnings history. Just as age-dependence alters the level of Pareto coefficients in Section 5.3, conditioning on past income histories further refines and reduces the within-cohort measures of inequality, thus resulting in lower levels of optimal income and savings taxes at the top.

Conclusion

This paper argues that labor income and savings taxes cannot be studied in isolation; they form instead an optimal policy mix that must be characterized jointly. Doing so leads to a stark tradeoff between raising one tax instrument versus the other. If the marginal rate of substitution between consumption and saving is homogeneous across agents, it is optimal to leave savings undistorted, as is well known since Atkinson and Stiglitz (1976), *but also* to set the level of the labor income tax rate at the static optimum given by Saez (2001). Away from this joint benchmark, it is optimal to raise (resp., lower) the savings tax rate above zero *and* simultaneously reduce (resp., increase) the labor income tax rate below the static optimum, if and only if more productive agents have a higher taste for saving relative to current consumption. The central message of our paper is that *consumption* data naturally determine the direction and the magnitude of the income and savings taxes away from the Saez (2001) and Atkinson and Stiglitz (1976) benchmark. Our novel optimal tax formulas, expressed in terms of the Pareto coefficients and elasticities of income and consumption, suggest that it is optimal to shift a significant share of the burden of taxes from income to savings. Given such crucial importance of measures of inequality and behavioral responses of consumption for optimal taxes, we believe empirical research should devote as much attention estimating them as has been given to their income counterparts.

[22] Brendon, Hellwig, and Maideu Morera (2024) extend the present analysis to persistent productivity shocks and arbitrary, exogenous, returns to savings and show that if shocks are persistent and agents have precautionary savings incentives, then it is optimal to subsidize savings.

References

Atkinson, Anthony and Joseph Stiglitz (1976). "The design of tax structure: direct versus indirect taxation". In: *Journal of public Economics* 6.1-2, pp. 55–75.

Attanasio, Orazio and Steven J Davis (1996). "Relative wage movements and the distribution of consumption". In: *Journal of political Economy* 104.6, pp. 1227–1262.

Bach, Laurent, Laurent E Calvet, and Paolo Sodini (2020). "Rich pickings? Risk, return, and skill in household wealth". In: *American Economic Review* 110.9, pp. 2703–47.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten (2016). "Consumption inequality and family labor supply". In: *American Economic Review* 106.2, pp. 387–435.

Brendon, Charles, Christian Hellwig, and Gerard Maideu Morera (2024). "Subsidizing Self-Insurance". In: *Working Paper*.

Buda, Gergely, Vasco M Carvalho, Stephen Hansen, Alvaro Ortiz, Tomasa Rodrigo, and José V Rodríguez Mora (2022). "National Accounts in a World of Naturally Occurring Data: A Proof of Concept for Consumption". In.

Chetty, Raj (2006). "A new method of estimating risk aversion". In: *American Economic Review* 96.5, pp. 1821–1834.

— (2009). "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods". In: *Annu. Rev. Econ.* 1.1, pp. 451–488.

— (2012). "Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply". In: *Econometrica* 80.3, pp. 969–1018.

Christiansen, Vidar (1984). "Which commodity taxes should supplement the income tax?" In: *Journal of Public Economics* 24.2, pp. 195–220.

Corlett, Wilfred J and Douglas C Hague (1953). "Complementarity and the excess burden of taxation". In: *The Review of Economic Studies* 21.1, pp. 21–30.

Diamond, Peter (1998). "Optimal income taxation: an example with a U-shaped pattern of optimal marginal tax rates". In: *American Economic Review*, pp. 83–95.

Diamond, Peter and Emmanuel Saez (2011). "The case for a progressive tax: from basic research to policy recommendations". In: *Journal of Economic Perspectives* 25.4, pp. 165–90.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri (2020). "Heterogeneity and persistence in returns to wealth". In: *Econometrica* 88.1, pp. 115–170.

Farhi, Emmanuel and Iván Werning (2013). "Insurance and taxation over the life cycle". In: *Review of Economic Studies* 80.2, pp. 596–635.

Ferey, Antoine, Benjamin Lockwood, and Dmitry Taubinsky (2021). *Sufficient Statistics for Nonlinear Tax Systems with Preference Heterogeneity.* Working Paper. National Bureau of Economic Research.

French, Eric (2005). "The effects of health, wealth, and wages on labour supply and retirement behaviour". In: *The Review of Economic Studies* 72.2, pp. 395–427.

Gaillard, Alexandre, Philipp Wangner, Christian Hellwig, and Nicolas Werquin (2023). "Consumption, Wealth, and Income Inequality: A Tale of Tails". In: *Available at SSRN.*

Gauthier, Stéphane and Fanny Henriet (2018). "Commodity taxes and taste heterogeneity". In: *European Economic Review* 101, pp. 284–296.

Gerritsen, Aart, Bas Jacobs, Alexandra Victoria Rusu, and Kevin Spiritus (2020). *Optimal taxation of capital income with heterogeneous rates of return.* CESifo Working Paper.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2021). *How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income.* Working Paper. National Bureau of Economic Research.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski (2003). "Optimal indirect and capital taxation". In: *The Review of Economic Studies* 70.3, pp. 569–587.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski (2016). "Redistribution and social insurance". In: *American Economic Review* 106.2, pp. 359–86.

Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl (2013). "Preference heterogeneity and optimal capital income taxation". In: *Journal of Public Economics* 97, pp. 160–175.

Gruber, Jon and Emmanuel Saez (2002). "The elasticity of taxable income: evidence and implications". In: *Journal of public Economics* 84.1, pp. 1–32.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante (2014). "Consumption and labor supply with partial insurance: An analytical framework". In: *American Economic Review* 104.7, pp. 2075–2126.

Hellwig, Christian (2021). *Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach.* TSE Working Paper.

Holm, Martin B, Rustam Jamilov, Marek Jasinski, and Plamen Nenov (2024). "Estimating the elasticity of intertemporal substitution using dividend tax news shocks". In.

Jacobs, Bas and Robin Boadway (2014). "Optimal linear commodity taxation under optimal nonlinear income taxation". In: *Journal of Public Economics* 117, pp. 201–210.

Jakobsen, Katrine, Kristian Jakobsen, Henrik Kleven, and Gabriel Zucman (2020). "Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark". In: *The Quarterly Journal of Economics* 329, p. 388.

Karahan, Fatih, Serdar Ozkan, and Jae Song (2022). "Anatomy of lifetime earnings inequality: Heterogeneity in job ladder risk vs. human capital". In: *FRB St. Louis Working Paper* 2022-2.

Keane, Michael P (2011). "Labor supply and taxes: A survey". In: *Journal of Economic Literature* 49.4, pp. 961–1075.

Kocherlakota, Narayana and Luigi Pistaferri (2009). "Asset pricing implications of Pareto optimality with private information". In: *Journal of Political Economy* 117.3, pp. 555–590.

Ligon, Ethan (1998). "Risk sharing and information in village economies". In: *The Review of Economic Studies* 65.4, pp. 847–864.

Mirrlees, James (1971). "An exploration in the theory of optimum income taxation". In: *The review of economic studies* 38.2, pp. 175–208.

— (1976). "Optimal tax theory: A synthesis". In: *Journal of public Economics* 6.4, pp. 327–358.

Saez, Emmanuel (2001). "Using elasticities to derive optimal income tax rates". In: *The review of economic studies* 68.1, pp. 205–229.

— (2002). "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes". In: *Journal of Public Economics* 83.2, pp. 217–230.

Scheuer, Florian and Joel Slemrod (2021). "Taxing our wealth". In: *Journal of Economic Perspectives* 35.1, pp. 207–30.

Schulz, Karl (2021). *Redistribution of Return Inequality*. CESifo Working Paper.

Toda, Alexis Akira and Kieran Walsh (2015). "The double power law in consumption and implications for testing Euler equations". In: *Journal of Political Economy* 123.5, pp. 1177–1200.

Townsend, Robert M (1994). "Risk and insurance in village India". In: *Econometrica: journal of the Econometric Society*, pp. 539–591.

Online Appendix for

"Using Consumption Data to Derive Optimal Income and Capital Tax Rates"

Christian Hellwig and Nicolas Werquin

A Proofs and Additional Theoretical Results

A.1 Proof of Theorem 1

Proof of Proposition 1. Consider an economy with a utility function $\mathcal{U}$ over N goods $X_1, \ldots, X_N$, and suppose that for all n,

$$\frac{\partial \ln(\mathcal{U}_n/\mathcal{U}_{n+1})}{\partial r} \equiv \frac{\mathcal{U}_{nr}}{\mathcal{U}_n} - \frac{\mathcal{U}_{n+1,r}}{\mathcal{U}_{n+1}} \leq 0.$$

In our baseline economy, we have $N = 3$, and the goods are ordered as follows: Y is indexed by $n = 1$, and C (respectively, S) is indexed by $n = 2$ if $V_{Sr}/V_S - U_{Cr}/U_C$ is positive (resp., negative). Consider an allocation $\boldsymbol{X}(r) = \{X_n(r)\}_{1 \leq n \leq N}$ that is locally incentive compatible and continuously differentiable at every r. We then have

$$\begin{aligned}
\mathcal{U}(\boldsymbol{X}(r); r) - \mathcal{U}(\boldsymbol{X}(r'); r) &= \sum_{n=1}^{N} \int_{r'}^{r} \mathcal{U}_n(\boldsymbol{X}(r''); r) X_n'(r'') \, dr'' \\
&= \sum_{n=1}^{N} \int_{r'}^{r} \frac{\mathcal{U}_n(\boldsymbol{X}(r''); r)}{\mathcal{U}_n(\boldsymbol{X}(r''); r'')} \mathcal{U}_n(\boldsymbol{X}(r''); r'') X_n'(r'') \, dr''
\end{aligned}$$

This expression can be rewritten as

$$\begin{aligned}
&\sum_{n=2}^{N} \int_{r'}^{r} \left[\frac{\mathcal{U}_n(\boldsymbol{X}(r''); r)}{\mathcal{U}_n(\boldsymbol{X}(r''); r'')} - \frac{\mathcal{U}_{n-1}(\boldsymbol{X}(r''); r)}{\mathcal{U}_{n-1}(\boldsymbol{X}(r''); r'')} \right] \sum_{k=n}^{N} \mathcal{U}_k(\boldsymbol{X}(r''); r'') X_k'(r'') \, dr'' \\
&+ \int_{r'}^{r} \frac{\mathcal{U}_1(\boldsymbol{X}(r''); r)}{\mathcal{U}_1(\boldsymbol{X}(r''); r'')} \sum_{k=1}^{N} \mathcal{U}_k(\boldsymbol{X}(r''); r'') X_k'(r'') \, dr''.
\end{aligned}$$

The single-crossing conditions imply that the term in square brackets of the last expression is positive (respectively, negative) when $r > r''$ (resp., $r < r''$). Therefore, we have $\mathcal{U}(\boldsymbol{X}(r)) - \mathcal{U}(\boldsymbol{X}(r'); r) \geq 0$ whenever $\sum_{k=n}^{N} \mathcal{U}_k(\boldsymbol{X}(r''); r'') X_k'(r'') \geq 0$ for all n. Moreover, the local incentive compatibility implies that $\sum_{k=1}^{N} \mathcal{U}_k(\boldsymbol{X}(r''); r'') X_k'(r'') = 0$, and therefore the above condition automatically holds for $n = 1$, and these conditions are all equivalent to $\sum_{k=1}^{n} \mathcal{U}_k(\boldsymbol{X}(r''); r'') X_k'(r'') \leq 0$ for all $n = 1, 2, ..., N-1$. In our baseline economy, assuming wlog that $V_{Sr}/V_S - U_{Cr}/U_C > 0$, these conditions boil down to: (i) $U_Y(r'') Y'(r'') \leq 0$, so that Y is weakly increasing ; and (ii)

$U_Y(r'')\,Y'(r'') + U_C(r'')\,C'(r'') \leq 0$, or equivalently $-U_S(r'')\,S'(r'') \leq 0$, so that S is weakly increasing. (If instead if $V_{Sr}/V_S - U_{Cr}/U_C < 0$, so that good $n = 2$ is savings, condition (ii) becomes (ii′) $-U_C(r'')\,C'(r'') \leq 0$, i.e., C is weakly increasing.)

Conversely, suppose that there exists an interval $(r_1, r_2) \subset (0,1)$ in which both $Y(\cdot)$ and $S(\cdot)$ are strictly decreasing, and assume again wlog that $V_{Sr}/V_S - U_{Cr}/U_C > 0$. Then the same sequence of equalities as above, combined with local incentive compatibility $U_C(r'')\,C'(r'') + V_S(r'')\,S'(r'') = -U_Y(r'')\,Y'(r'')$, implies that

$$\begin{aligned}\mathcal{U}(\boldsymbol{X}(r_2);r_2) - \mathcal{U}(\boldsymbol{X}(r_1);r_2) \quad = \quad & \int_{r_1}^{r_2}\left[\frac{U_C(\boldsymbol{X}(r'');r_2)}{U_C(\boldsymbol{X}(r'');r'')} - \frac{U_Y(\boldsymbol{X}(r'');r_2)}{U_Y(\boldsymbol{X}(r'');r'')}\right]\{-U_Y(r'')\,Y'(r'')\}\,dr'' \\ & + \int_{r_1}^{r_2}\left[\frac{V_S(\boldsymbol{X}(r'');r_2)}{V_S(\boldsymbol{X}(r'');r'')} - \frac{U_C(\boldsymbol{X}(r'');r_2)}{U_C(\boldsymbol{X}(r'');r'')}\right]V_S(r'')\,S'(r'')\,dr''\end{aligned}$$

is strictly negative, contradicting global incentive compatibility. □

Proof (and Generalization) of Proposition 2. Consider a general weighted-utilitarian social welfare objective, with Pareto weights $\omega(r) \geq 0$ assigned to ranks r that satisfy $\mathbb{E}[\omega] = 1$. The social planner minimizes the net present value of transfers:

$$K(v_0) = \min_{\{C(r),Y(r),S(r)\}} \int_0^1 [C(r) - Y(r) + S(r)]\,dr$$

subject to the ex-ante promise-keeping constraint

$$\int_0^1 \omega(r)\,W(r)\,dr \geq v_0$$

the promise-keeping constraint

$$W(r) = U(C(r), Y(r); r) + V(S(r); r)$$

and the local incentive compatibility constraint

$$W'(r) = U_r(C(r), Y(r); r) + V_r(S(r); r).$$

If the utility promise v_0 is chosen so that the net present value of transfers at the optimum equals 0, the solution to the problem corresponds to the allocation that maximizes the expected utility of agents, subject to satisfying an aggregate break-even condition. (The problem studied in the main body of the paper is a special case of this general formulation with $\omega(r) = 0$ for all $r > 0$.)

This is an optimal control problem with $W(\cdot)$ as the state variable, and $C(\cdot)$, $Y(\cdot)$, and $S(\cdot)$ as controls. Defining λ, $\psi(r)$, and $\phi(r)$ as the multipliers on, respectively, the ex-ante promise-keeping constraint and the promise-keeping and local incentive compatibility constraints given r, the Hamiltonian for this problem is given by:

$$\begin{aligned}\mathcal{H} &= \{C(r) - Y(r) + S(r) + \lambda(v_0 - W(r))\omega(r)\} \\ &\quad + \psi(r)\{W(r) - U(C(r), Y(r); r) - V(S(r); r)\} \\ &\quad + \phi(r)\{U_r(C(r), Y(r); r) + V_r(S(r); r)\}.\end{aligned}$$

The first-order conditions with respect to the allocations $C(\cdot)$, $Y(\cdot)$, and $S(\cdot)$ yield:

$$\psi(r) = \frac{1}{U_C(r)} + \phi(r)\frac{U_{Cr}(r)}{U_C(r)} = \frac{1}{-U_Y(r)} + \phi(r)\frac{U_{Yr}(r)}{U_Y(r)} = \frac{1}{V_S(r)} + \phi(r)\frac{V_{Sr}(r)}{V_S(r)}.$$

The first-order conditions for $C(\cdot)$, $Y(\cdot)$, and $S(\cdot)$ define a shadow cost of utility of agents with rank r, $\psi(r)$, which consists of a direct shadow cost $1/U_C(r)$, $1/(-U_Y(r))$, or $1/V_S(r)$ of increasing rank r utility through higher consumption, lower income or higher savings, and a second term that measures how such a consumption or income increase affects $U_r(r)$ and $V_r(r)$ and thereby tightens or relaxes the local incentive compatibility constraint at r by $\frac{U_{Cr}(r)}{U_C(r)}$, $\frac{U_{Yr}(r)}{U_Y(r)}$, or $\frac{V_{Sr}(r)}{V_S(r)}$. The latter is weighted by the multiplier $\phi(r)$ and added to the former.

Now define

$$m_C(r) = \exp\left(-\int_r^1 \frac{U_{Cr}(r')}{U_C(r')}dr'\right), \quad m_Y(r) = \exp\left(-\int_r^1 \frac{U_{Yr}(r')}{U_Y(r')}dr'\right), \quad m_S(r) = \exp\left(-\int_r^1 \frac{V_{Sr}(r')}{V_S(r')}dr'\right)$$

so that $U_{Cr}(r)/U_C(r) = m_C'(r)/m_C(r)$, and analogous expressions for the other variables. Combining the first two first-order conditions and rearranging terms then yields the following static optimality condition:

$$\frac{1}{U_C(r)}\frac{\tau_Y(r)}{1-\tau_Y(r)} = \frac{1}{-U_Y(r)} - \frac{1}{U_C(r)} = \left(\frac{U_{Cr}(r)}{U_C(r)} - \frac{U_{Yr}(r)}{U_Y(r)}\right)\phi(r) \equiv A(r)\phi(r).$$

The multipliers $\phi(\cdot)$ and λ are derived by solving the linear ODE $\phi'(r) = -\partial\mathcal{H}/\partial W$, after substituting out $\psi(r)$ using the first first-order condition:

$$\phi'(r) = -\frac{\partial \mathcal{H}}{\partial W} = \lambda\omega(r) - \psi(r) = \lambda\omega(r) - \frac{1}{U_C(r)} - \phi(r)\frac{U_{Cr}(r)}{U_C(r)},$$

along with the boundary conditions $\phi(0) = \phi(1) = 0$. Substituting into the previous ODE and integrating out yields

$$\phi(1)\,m_C(1)-\phi(r)\,m_C(r)=\int_r^1\left(\lambda\omega(r')-\frac{1}{U_C(r')}\right)m_C(r')\,dr',$$

or

$$\phi(r)=\frac{1-r}{m_C(r)}\left\{\mathbb{E}\left[\frac{1}{U_C(r')}m_C(r')\middle|r'\geq r\right]-\lambda\mathbb{E}\left[\omega(r')\,m_C(r')\middle|\,r'\geq r\right]\right\}.$$

The boundary condition $\phi(0)=0$ then gives $\lambda=\mathbb{E}[m_C U_C^{-1}]/\mathbb{E}[m_C\,\omega]$. Therefore,

$$\frac{\phi(r)}{1-r} \;=\; \mathbb{E}\left[\frac{1}{U_C(r')}\frac{m_C(r')}{m_C(r)}\middle|r'\geq r\right]-\frac{\mathbb{E}\left[\frac{1}{U_C(r')}\frac{m_C(r')}{m_C(r)}\right]}{\mathbb{E}\left[\omega(r')\frac{m_C(r')}{m_C(r)}\right]}\mathbb{E}\left[\omega(r')\frac{m_C(r')}{m_C(r)}\middle|r'\geq r\right] \;\equiv\; \frac{B_C(r)}{U_C(r)}.$$

Substituting this expression into the static optimality condition then yields the first intra-temporal optimality condition ("ABC") $\frac{\tau_Y(r)}{1-\tau_Y(r)}=A(r)\cdot B_C(r)$.

The first-order condition for income yields an analogous ODE,

$$\phi'(r)=\lambda\omega(r)-\frac{1}{-U_Y(r)}-\phi(r)\frac{U_{Yr}(r)}{U_Y(r)}.$$

Apply the same steps as above to get

$$\frac{\phi(r)}{1-r} \;=\; \mathbb{E}\left[\frac{1}{-U_Y(r')}\frac{m_Y(r')}{m_Y(r)}\middle|r'\geq r\right]-\frac{\mathbb{E}\left[\frac{1}{-U_Y(r')}\frac{m_Y(r')}{m_Y(r)}\right]}{\mathbb{E}\left[\omega(r')\frac{m_Y(r')}{m_Y(r)}\right]}\mathbb{E}\left[\omega(r')\frac{m_Y(r')}{m_Y(r)}\middle|r'\geq r\right] \;\equiv\; \frac{B_Y(r)}{-U_Y(r)},$$

and $\lambda=\mathbb{E}[m_Y(-U_Y^{-1})]/\mathbb{E}[m_Y\,\omega]$. We obtain the second intra-temporal optimality condition ("ABC") $\tau_Y(r)=A(r)\cdot B_Y(r)$, and setting $B_Y(r)/(-U_Y(r))$ equal to $B_C(r)/U_C(r)$, the redistributional arbitrage condition (equation (24)):

$$1-\tau_Y(r)=\frac{B_Y(r)}{B_C(r)}.$$

Finally, we solve for the inter-temporal optimality condition. Combining the ODE $\phi'(r)=-\partial\mathcal{H}/\partial W=\lambda\omega(r)-\psi(r)$ with the first-order condition for savings yields

$$\phi'(r)=\lambda\omega(r)-\frac{1}{V_S(r)}-\phi(r)\frac{V_{Sr}(r)}{V_S(r)}.$$

The previous ODE can be integrated and solved along the same lines as above to find

$$\frac{\phi(r)}{1-r} = \mathbb{E}\left[\frac{1}{V_S(r')}\frac{m_S(r')}{m_S(r)}\middle| r' \geq r\right] - \frac{\mathbb{E}\left[\frac{1}{V_S(r')}\frac{m_S(r')}{m_S(r)}\right]}{\mathbb{E}\left[\omega(r')\frac{m_S(r')}{m_S(r)}\right]}\mathbb{E}\left[\omega(r')\frac{m_S(r')}{m_S(r)}\middle| r' \geq r\right] \equiv \frac{B_S(r)}{V_S(r)},$$

with $\lambda = \mathbb{E}[m_S/V_S]/\mathbb{E}[m_S\,\omega]$. Equating this last expression to $B_C(r)/U_C(r)$ then yields the arbitrage representation for the savings wedge (equation (25)):

$$1 + \tau_S(r) \equiv \frac{V_S(r)}{U_C(r)} = \frac{B_S(r)}{B_C(r)}.$$

We finally show that if savings are unbounded above and $\lim_{r\to 1}\tau_Y(r) < 1$, then optimal allocations satisfy the Inada condition $\lim_{r\to 1} U_C(r) = \lim_{r\to 1}(-U_Y(r)) = \lim_{r\to 1} V_S(r) = 0$. The last equality follows from the Inada condition on V. Moreover, $\lim_{r\to 1}(-U_Y(r)) = \lim_{r\to 1}\frac{B_Y(r)}{B_S(r)}V_S(r)$. It is easy to check that $\lim_{r\to 1} B_S(r) \geq 1$ and $\lim_{r\to 1} B_Y(r) \leq 1$, and hence $\lim_{r\to 1}(-U_Y(r)) \leq \lim_{r\to 1} V_S(r) = 0$. Finally, $\lim_{r\to 1} U_C(r) = \lim_{r\to 1}\frac{(-U_Y(r))}{1-\tau_Y(r)} = 0$. □

Proof of the Atkinson and Stiglitz (1976) Theorem. The theorem of Atkinson and Stiglitz (1976) and its converse follow easily from the previous proof: The optimal savings wedge is $\tau_S(r)$ is positive (resp., negative) for all r if and only if $V_{Sr}(r)/V_S(r) - U_{Cr}(r)/U_C(r)$ is positive (resp., negative) for all r. To see this, recall that

$$\frac{1}{V_S(r)} + \phi(r)\frac{V_{Sr}(r)}{V_S(r)} = \frac{1}{U_C(r)} + \phi(r)\frac{U_{Cr}(r)}{U_C(r)},$$

with $\phi(r) > 0$ for all r. Since $U_{Cr}(r)/U_C(r) - V_{Sr}(r)/V_S(r)$ has a constant sign, we get $U_C(r) \lesseqgtr V_S(r)$, or $\tau_S(r) \gtreqless 0$ for all r, if and only if $U_{Cr}(r)/U_C(r) - V_{Sr}(r)/V_S(r) \lesseqgtr 0$ for all r. More generally, consider a framework with multiple goods as in Section 5.2. For any two goods $m < n$, suppose that the marginal rate of substitution $U_m(r)/U_n(r)$ is weakly increasing in r, so that $U_n(r)/U_n(r') \geq U_m(r)/U_m(r')$ for all $r' > r$. Equivalently, $U_{mr}(r)/U_m(r) \geq U_{nr}(r)/U_n(r)$ for all r. The first-order conditions of the planner's problem read

$$\frac{p_m}{U_m(r)} = \frac{p_n}{U_n(r)} + \phi(r)\left(\frac{U_{nr}(r)}{U_n(r)} - \frac{U_{mr}(r)}{U_m(r)}\right),$$

with $\phi(r) > 0$ is the Lagrange multiplier on the local incentive constraint. We immediately obtain that $\tau_{m,n}(r) = 0$ for all r if and only if the MRS $U_m(r)/U_n(r)$ is uniform across types. More generally, we have $U_m(r)/U_n(r) < p_m/p_n$, so that $\tau_{m,n}(r) > 0$ iff $U_{nr}(r)/U_n(r) > U_{mr}(r)/U_m(r)$. □

Proof of Proposition 3. We now rewrite the optimality conditions derived above in terms of sufficient statistics. Totally differentiating $U_C(r)$, $-U_Y(r)$, and $V_S(r)$ yields respectively

$$
\begin{aligned}
\frac{\frac{d}{dr}U_C(r)}{U_C(r)} &= \frac{U_{CC}(r)}{U_C(r)}C'(r) + \frac{U_{CY}(r)}{U_C(r)}Y'(r) + \frac{U_{Cr}(r)}{U_C(r)}, \\
\frac{\frac{d}{dr}(-U_Y(r))}{-U_Y(r)} &= \frac{U_{CY}(r)}{U_Y(r)}C'(r) + \frac{U_{YY}(r)}{U_Y(r)}Y'(r) + \frac{U_{Yr}(r)}{U_Y(r)}, \\
\frac{\frac{d}{dr}V_S(r)}{V_S(r)} &= \frac{V_{SS}(r)}{V_S(r)}S'(r) + \frac{V_{Sr}(r)}{V_S(r)}.
\end{aligned}
$$

Using the two first-order conditions $-U_Y/U_C = 1-\tau_Y$ and $V_S/U_C = 1+\tau_S$, and noting that $CU_{CY}/(-U_Y) = s_C\zeta_{CY}$, we obtain that these three equations can be rewritten in terms of our elasticities and Pareto coefficients as follows:

$$
\begin{aligned}
-\frac{\frac{d\ln U_C(r)}{d\ln(1-r)}}{1-r} &= -\frac{\zeta_C(r)}{(1-r)\rho_C(r)} + \frac{\zeta_{CY}(r)}{(1-r)\rho_Y(r)} + \frac{U_{Cr}(r)}{U_C(r)}, \\
-\frac{\frac{d\ln(1-\tau_Y(r))}{d\ln(1-r)} + \frac{d\ln U_C(r)}{d\ln(1-r)}}{1-r} &= -\frac{s_C(r)\zeta_{CY}(r)}{(1-r)\rho_C(r)} + \frac{\zeta_Y(r)}{(1-r)\rho_Y(r)} + \frac{U_{Yr}(r)}{U_Y(r)}, \\
-\frac{\frac{d\ln(1+\tau_S(r))}{d\ln(1-r)} + \frac{d\ln U_C(r)}{d\ln(1-r)}}{1-r} &= -\frac{\zeta_S(r)}{(1-r)\rho_S(r)} + \frac{V_{Sr}(r)}{V_S(r)}.
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\frac{U_{Yr}}{U_Y} - \frac{U_{Cr}}{U_C} &= -\frac{\zeta_C}{(1-r)\rho_C} - \frac{\zeta_Y}{(1-r)\rho_Y} + \left(1+\frac{s_C\rho_Y}{\rho_C}\right)\frac{\zeta_{CY}}{(1-r)\rho_Y} - \frac{\tau_Y'}{1-\tau_Y} \\
\frac{V_{Sr}}{V_S} - \frac{U_{Cr}}{U_C} &= -\frac{\zeta_C}{(1-r)\rho_C} + \frac{\zeta_S}{(1-r)\rho_S} + \frac{\zeta_{CY}}{(1-r)\rho_Y} + \frac{\tau_S'}{1+\tau_S}.
\end{aligned}
$$

Moreover, let

$$
M_C(r) = \frac{1}{U_C(r)}e^{-\int_r^1 \frac{U_{Cr}(r')}{U_C(r')}dr'},\ M_Y(r) = \frac{1}{-U_Y(r)}e^{-\int_r^1 \frac{U_{Yr}(r')}{U_Y(r')}dr'},\ M_S(r) = \frac{1}{V_S(r)}e^{-\int_r^1 \frac{V_{Sr}(r')}{V_S(r')}dr'}.
$$

We have

$$
\begin{aligned}
M_C(r) &= \frac{1}{U_C(r)} e^{-\int_r^1 \frac{\frac{d}{dr}U_C(r')}{U_C(r')}dr'} e^{\int_r^1\left\{-\zeta_C(r')\frac{C'(r')}{C(r')}+\zeta_{CY}(r')\frac{Y'(r')}{Y(r')}\right\}dr'} = e^{\int_r^1\left\{-\frac{\zeta_C(r')}{\rho_C(r')}+\frac{\zeta_{CY}(r')}{\rho_Y(r')}\right\}\frac{dr'}{1-r'}}, \\
M_Y(r) &= \frac{1}{-U_Y(r)} e^{-\int_r^1 \frac{\frac{d}{dr}(-U_Y(r'))}{-U_Y(r')}dr'} e^{\int_r^1\left\{\zeta_Y(r')\frac{Y'(r')}{Y(r')}-s_C(r')\zeta_{CY}(r')\frac{C'(r')}{C(r')}\right\}dr'} = e^{\int_r^1\left\{\frac{\zeta_Y(r')}{\rho_Y(r')}-\frac{s_C(r')\zeta_{CY}(r')}{\rho_C(r')}\right\}\frac{dr'}{1-r'}}, \\
M_S(r) &= \frac{1}{V_S(r)} e^{-\int_r^1 \frac{\frac{d}{dr}V_S(r')}{V_S(r')}dr'} e^{-\int_r^1 \zeta_S(r')\frac{S'(r')}{S(r')}dr'} = e^{-\int_r^1 \frac{\zeta_S(r')}{\rho_S(r')}\frac{dr'}{1-r'}}.
\end{aligned}
$$

This leads to equation (23), along with the corresponding expressions in Section 4.4.

Finally, we have $\lim_{r\to1}\frac{1-r}{U_C(r)} = 0$ from the boundary condition for tax distortions at the top. This leaves two possibilities. First, if $\lim_{r\to1}\frac{dU_C(r)}{d(1-r)} < \infty$, then $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} = 0$, i.e., the inverse marginal utilities necessarily have a thin upper tail. Second, if $\lim_{r\to1}\frac{dU_C(r)}{d(1-r)} = \infty$, there exists a sequence $\{r_n\} \underset{n\to\infty}{\longrightarrow} 1$, such that $U_C(r_n) > U_C(1) + (1-r_n)\left.\frac{dU_C(r)}{d(1-r)}\right|_{r=r_n}$, where $U_C(1) = \lim_{r\to1}U_C(r)$. Dividing by $U_C(r_n)$ and taking the limit as $n\to\infty$ implies that

$$
\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} \leq 1 - \lim_{r\to1}\frac{U_C(1)}{U_C(r_n)}.
$$

Hence if $U_C(1) > 0$, we obtain $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} = 0$, whereas if $U_C(1) = 0$, $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} \leq 1$. Furthermore, if it were the case that $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} = 1$, then there would exist $A \neq 0$, such that $U_C(r) = A(1-r) + o\left((1-r)^2\right)$. But then $\lim_{r\to1}\frac{1-r}{U_C(r)} = \frac{1}{A} \neq 0$, which would violate the boundary condition. To summarize, $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)}$ is bounded above by 1 (imposing a lower bound on the Pareto tail coefficient of inverse marginal utilities) whenever $U_C(1) = 0$, and $\lim_{r\to1}\frac{d\ln U_C(r)}{d\ln(1-r)} = 0$ (implying that inverse marginal utilities are thin-tailed), whenever $U_C(1) > 0$. $\square$

Proof of Theorem 1. Using the expressions derived in the previous proof, we can write

$$
\begin{aligned}
\lim_{r\to1}\tau_Y(r) &= 1 - \lim_{r\to1}\frac{\mathbb{E}\left[\left.\frac{M_Y(r')}{M_Y(r)}\right| r'\geq r\right]}{\mathbb{E}\left[\left.\frac{M_C(r')}{M_C(r)}\right| r'\geq r\right]} = 1 - \lim_{r\to1}\frac{\mathbb{E}\left[\left. e^{-\int_r^{r'}\zeta_Y\frac{Y'(r'')}{Y(r'')}dr''+\int_r^{r'}s_C\zeta_{CY}\frac{C'(r'')}{C(r'')}dr''}\right| r'\geq r\right]}{\mathbb{E}\left[\left. e^{\int_r^{r'}\zeta_C\frac{C'(r'')}{C(r'')}dr''-\int_r^{r'}\zeta_{CY}\frac{Y'(r'')}{Y(r'')}dr''}\right| r'\geq r\right]} \\
&= 1 - \lim_{r\to1}\frac{\mathbb{E}\left[\left.\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_Y}\left(\frac{C(r')}{C(r)}\right)^{s_C\zeta_{CY}}\right| r'\geq r\right]}{\mathbb{E}\left[\left.\left(\frac{C(r')}{C(r)}\right)^{\zeta_C}\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_{CY}}\right| r'\geq r\right]}.
\end{aligned}
$$

For the numerator, define $X(r) \equiv (Y(r))^{-\zeta_Y}(C(r))^{s_C\zeta_{CY}}$. We wish to compute $\mathbb{E}\left[X(r')/X(r)\,|r'\geq r\right]$, given that $C(r)$, $Y(r)$, and $X(r)$ are perfectly co-monotonic and C and Y are distributed according to a Pareto distribution with tail coefficients ρ_C and ρ_Y. We get

$$-\frac{d\ln X(r)}{d\ln(1-r)} = (1-r)\frac{X'(r)}{X(r)} = -\zeta_Y(1-r)\frac{Y'(r)}{Y(r)} + s_C\zeta_{CY}(1-r)\frac{C'(r)}{C(r)} = -\frac{\zeta_Y}{\rho_Y} + \frac{s_C\zeta_{CY}}{\rho_C},$$

so that $X(r)$ follows a Pareto distribution with tail coefficient $1/[-\zeta_Y/\rho_Y + s_C\zeta_{CY}/\rho_C]$. This implies

$$\lim_{r\to1}\mathbb{E}\left[\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_Y}\left(\frac{C(r')}{C(r)}\right)^{s_C\zeta_{CY}}\middle| r'\geq r\right] = \left[1+\frac{\zeta_Y}{\rho_Y}-\frac{s_C\zeta_{CY}}{\rho_C}\right]^{-1}$$

Along the same lines,

$$\lim_{r\to1}\mathbb{E}\left[\left(\frac{C(r')}{C(r)}\right)^{\zeta_C}\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_{CY}}\middle| r'\geq r\right] = \left[1-\frac{\zeta_C}{\rho_C}+\frac{\zeta_{CY}}{\rho_Y}\right]^{-1}$$

and therefore

$$\lim_{r\to1}\tau_Y(r) = 1-\frac{1-\frac{\zeta_C}{\rho_C}+\frac{\zeta_{CY}}{\rho_Y}}{1+\frac{\zeta_Y}{\rho_Y}-\frac{s_C\zeta_{CY}}{\rho_C}}.$$

At the optimal allocation, we must have $\zeta_C/\rho_C < 1+\zeta_{CY}/\rho_Y$. It then follows automatically that $\lim_{r\to1}\tau_Y(r) < 1$. To prove the second part of Theorem 1, follow analogous steps as above to get

$$\begin{aligned}\lim_{r\to1}B_S(r) &\equiv \lim_{r\to1}\mathbb{E}\left[\frac{M_S(r')}{M_S(r)}\middle| r'\geq r\right] = \lim_{r\to1}\mathbb{E}\left[e^{\int_r^{r'}\zeta_S\frac{S'(r'')}{S(r'')}dr''}\middle| r'\geq r\right]\\ &= \lim_{r\to1}\mathbb{E}\left[\left(\frac{S(r')}{S(r)}\right)^{\zeta_S}\middle| r'\geq r\right] = \left[1-\frac{\zeta_S}{\rho_S}\right]^{-1}\end{aligned}$$

for $\zeta_S/\rho_S < 1$. Combining this result with the previous expressions, we get

$$\lim_{r\to1}\tau_S(r) = \frac{1-\frac{\zeta_C}{\rho_C}+\frac{\zeta_{CY}}{\rho_Y}}{1-\frac{\zeta_S}{\rho_S}} - 1.$$

This concludes the proof. □

A.2 Proof of Corollary 1

Relationship between Preference and Behavioral Elasticities. Consider a labor income tax schedule $T_Y(Y)$ and a savings tax schedule $T_S(S)$. For ease of notation, assume that the tax schedules are locally linear in the top bracket, $T_Y''(Y) = T_S''(S) = 0$. A perturbation of the total tax payment by ∂T_Y and the marginal tax rate by $\partial T_Y'$ leads to behavioral responses $(\partial Y, \partial C, \partial S)$ that satisfy the perturbed first-order conditions

$$-\frac{U_Y\left[C+\partial C,Y+\partial Y;r\right]}{U_C\left[C+\partial C,Y+\partial Y;r\right]}=1-T_Y'\left(Y\right)-\partial T_Y'$$

and

$$\frac{V'\left[S+\partial S\right]}{U_C\left[C+\partial C,Y+\partial Y,r\right]}=1+T_S'\left(S\right)$$

with

$$\partial C+\left(1+T_S'\left(S\right)\right)\partial S=\left(1-T_Y'\left(Y\right)\right)\partial Y-\partial T_Y.$$

We obtain the responses of income, consumption and savings by taking first-order Taylor expansions of the two perturbed FOCs around 0:

$$\hat{\zeta}_Y\frac{\partial Y}{Y}+\hat{\zeta}_C\frac{\partial C}{C}=-\frac{\partial T_Y'}{1-T_Y'}$$

and

$$\hat{\zeta}_S\frac{\partial Y}{Y}-\left[s_S\hat{\zeta}_C+s_C\hat{\zeta}_S\right]\frac{\partial C}{C}=\zeta_S\frac{\partial T_Y}{\left(1-T_Y'\right)Y}$$

where we let $\hat{\zeta}_C=\zeta_C-s_C\zeta_{CY}$, $\hat{\zeta}_Y=\zeta_Y-\zeta_{CY}$, $\hat{\zeta}_S=\zeta_S+s_S\zeta_{CY}$ and $s_C=C/\left(1-T_Y'\right)Y$, $s_S=\left(1+T_S'\right)S/\left(\left(1-T_Y'\right)Y\right)$. Note that as $r\to 1$, so that $Y,S\to\infty$ and T_Y',T_S' converge to constants, we have $s_C+s_S\to 1$. Solving this system leads to

$$\begin{aligned}\frac{\partial Y}{Y}&=-\zeta_Y^H\frac{\partial T_Y'}{1-T_Y'}+\zeta_Y^I\frac{\partial T_Y}{\left(1-T_Y'\right)Y},\\ \frac{\partial C}{C}&=-\zeta_C^H\frac{\partial T_Y'}{1-T_Y'}-\zeta_C^I\frac{\partial T_Y}{\left(1-T_Y'\right)Y},\\ \frac{\partial S}{S}&=-\zeta_S^H\frac{\partial T_Y'}{1-T_Y'}-\zeta_S^I\frac{\partial T_Y}{\left(1-T_Y'\right)Y},\end{aligned}$$

with

$$\zeta_Y^H=\frac{1}{\hat{\zeta}_Y+\frac{\hat{\zeta}_C\hat{\zeta}_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}},\quad \zeta_Y^I=\frac{\frac{\hat{\zeta}_C\zeta_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}}{\hat{\zeta}_Y+\frac{\hat{\zeta}_C\hat{\zeta}_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}},$$

and

$$\zeta_C^H=\frac{\frac{\hat{\zeta}_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}}{\hat{\zeta}_Y+\frac{\hat{\zeta}_C\hat{\zeta}_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}},\quad \zeta_C^I=\frac{\frac{\hat{\zeta}_Y\zeta_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}}{\hat{\zeta}_Y+\frac{\hat{\zeta}_C\hat{\zeta}_S}{s_S\hat{\zeta}_C+s_C\hat{\zeta}_S}}.$$

In particular, when $s_C\to 1$ and $s_S\to 0$, we have $\zeta_Y^H=\zeta_C^H=1/\left(\hat{\zeta}_Y+\hat{\zeta}_C\right)$, and $\zeta_Y^I=1-\zeta_C^I=\hat{\zeta}_C/\left(\hat{\zeta}_Y+\hat{\zeta}_C\right)$. When $s_C\to 0$ and $s_S\to 1$, we have $\zeta_Y^H=\frac{1}{\zeta_Y+\zeta_S}$, $\zeta_Y^I=\frac{\zeta_S}{\zeta_Y+\zeta_S}$, $\zeta_C^H=\frac{\zeta_S+\zeta_{CY}}{\zeta_C}\zeta_Y^H$, and $\zeta_C^I = \frac{\zeta_Y - \zeta_{CY}}{\zeta_C} \zeta_Y^I$. Moreover, the budget constraint $C + S + T_S(S) = Y - T_Y(Y)$ implies that:

$$s_C \zeta_C^H + (1 - s_C) \zeta_S^H = \zeta_Y^H$$

and

$$1 - s_C \zeta_C^I - (1 - s_C) \zeta_S^I = \zeta_Y^I.$$

Conversely, the primitive elasticities $\hat{\zeta}_C, \hat{\zeta}_Y, \zeta_{CY}, \zeta_S$ are given in terms of the behavioral responses $\zeta_Y^H, \zeta_Y^I, \zeta_C^H, \zeta_C^I$ via

$$\hat{\zeta}_Y = \frac{1}{\zeta_Y^H + \zeta_Y^I \zeta_C^H / \zeta_C^I}, \quad \hat{\zeta}_C = \frac{\zeta_Y^I}{\zeta_C^I} \hat{\zeta}_Y, \quad \zeta_{CY} = \frac{\zeta_C^H \hat{\zeta}_C - \zeta_Y^I}{\zeta_Y^H - s_C \zeta_C^H}, \quad \zeta_S = \frac{(1 - s_C) \zeta_Y^I}{\zeta_Y^H - s_C \zeta_C^H}.$$

Finally, the Slutsky equations read

$$\begin{aligned}
\frac{\partial Y}{\partial (1 - T_Y')} &= \left. \frac{\partial Y}{\partial (1 - T_Y')} \right|_U - \frac{\partial Y}{\partial T_Y} Y, \\
\frac{\partial C}{\partial (1 - T_Y')} &= \left. \frac{\partial C}{\partial (1 - T_Y')} \right|_U - \frac{\partial C}{\partial T_Y} Y.
\end{aligned}$$

Therefore, using the above definitions of the substitution and income effects $\zeta_C^H = \partial \ln C / \partial \ln (1 - T_Y')|_U$, $\zeta_C^I = -(1/s_C) \partial C / \partial T_Y$ and denoting by ζ_Y^M, ζ_C^M the Marshallian (uncompensated) elasticities, we have $\zeta_Y^M = \zeta_Y^H - \zeta_Y^I$ and $\zeta_C^M = \zeta_C^H + \zeta_C^I$. □

First-Stage and Second-Stage Elasticities. The FOC to the problem $\max_Y \mathcal{U}(Y - T_Y(Y), Y; r)$ reads $(1 - \tau_Y) \mathcal{U}_M + \mathcal{U}_Y = 0$. We can thus derive the impact of a tax reform $(\partial T_Y, \partial \tau_Y)$ as follows:

$$\begin{aligned}
\mathcal{U}_M \partial \tau_Y &= (1 - \tau_Y) \left[((1 - \tau_Y) \mathcal{U}_{MM} + \mathcal{U}_{MY}) \partial Y - \mathcal{U}_{MM} \partial T_Y \right] \\
&\quad + \left[((1 - \tau_Y) \mathcal{U}_{MY} + \mathcal{U}_{YY}) \partial Y - \mathcal{U}_{MY} \partial T_Y \right]
\end{aligned}$$

Rearranging terms leads to

$$\begin{aligned}
\frac{\partial Y}{Y} &= \frac{-\mathcal{U}_Y / Y}{(\mathcal{U}_Y / \mathcal{U}_M)^2 \mathcal{U}_{MM} - 2 (\mathcal{U}_Y / \mathcal{U}_M) \mathcal{U}_{MY} + \mathcal{U}_{YY}} \frac{\partial \tau_Y}{1 - \tau_Y} \\
&\quad + \frac{(\mathcal{U}_Y / \mathcal{U}_M)^2 \mathcal{U}_{MM} - (\mathcal{U}_Y / \mathcal{U}_M) \mathcal{U}_{MY}}{(\mathcal{U}_Y / \mathcal{U}_M)^2 \mathcal{U}_{MM} - 2 (\mathcal{U}_Y / \mathcal{U}_M) \mathcal{U}_{MY} + \mathcal{U}_{YY}} \frac{\partial T_Y}{(1 - \tau_Y) Y}.
\end{aligned}$$

Thus, we obtain

$$\zeta_Y^H = \frac{\mathcal{U}_Y/Y}{\left(\mathcal{U}_Y/\mathcal{U}_M\right)^2\mathcal{U}_{MM} - 2\left(\mathcal{U}_Y/\mathcal{U}_M\right)\mathcal{U}_{MY} + \mathcal{U}_{YY}} = \frac{1}{-\frac{M\mathcal{U}_{MM}}{\mathcal{U}_M}\frac{(1-\tau_Y)Y}{Y-T_Y(Y)} - 2\frac{Y\mathcal{U}_{MY}}{\mathcal{U}_M} + \frac{Y\mathcal{U}_{YY}}{\mathcal{U}_Y}}$$

$$\zeta_Y^I = \frac{\left(\mathcal{U}_Y/\mathcal{U}_M\right)^2\mathcal{U}_{MM} - \left(\mathcal{U}_Y/\mathcal{U}_M\right)\mathcal{U}_{MY}}{\left(\mathcal{U}_Y/\mathcal{U}_M\right)^2\mathcal{U}_{MM} - 2\left(\mathcal{U}_Y/\mathcal{U}_M\right)\mathcal{U}_{MY} + \mathcal{U}_{YY}} = \frac{-\frac{M\mathcal{U}_{MM}}{\mathcal{U}_M}\frac{(1-\tau_Y)Y}{Y-T_Y(Y)} - \frac{Y\mathcal{U}_{MY}}{\mathcal{U}_M}}{-\frac{M\mathcal{U}_{MM}}{\mathcal{U}_M}\frac{(1-\tau_Y)Y}{Y-T_Y(Y)} - 2\frac{Y\mathcal{U}_{MY}}{\mathcal{U}_M} + \frac{Y\mathcal{U}_{YY}}{\mathcal{U}_Y}}.$$

Hence, using the fact that, for top earners, $\frac{Y\mathcal{U}_Y}{M\mathcal{U}_M} = \frac{(1-\tau_Y(Y))Y}{Y-T_Y(Y)} \to 1$ as the marginal tax rate $\tau_Y(Y)$ converges to a constant strictly below 1,

$$\zeta_Y^H = \frac{1}{\tilde{\zeta}_M - 2\tilde{\zeta}_{MY} + \tilde{\zeta}_Y}, \quad \zeta_Y^I = \frac{\tilde{\zeta}_M - \tilde{\zeta}_{MY}}{\tilde{\zeta}_M - 2\tilde{\zeta}_{MY} + \tilde{\zeta}_Y}.$$

Therefore,

$$\frac{1-\zeta_Y^I}{\zeta_Y^H} = \tilde{\zeta}_Y - \tilde{\zeta}_{MY}, \quad \text{and} \quad \frac{\zeta_Y^I}{\zeta_Y^H} = \tilde{\zeta}_M - \tilde{\zeta}_{MY}.$$

Next, we have $\mathcal{U}(M,Y;r) = U(M - S^* - T_S(S^*), Y; r) + V(S^*, r)$, where S^* is the solution to the first-order condition $(1+\tau_S)U_C(M - S^* - T_S(S^*), Y; r) = V_S(S^*, r)$. By the envelope theorem, we have

$$\mathcal{U}_M(M,Y;r) = U_C(C,Y;r),$$

where $C = M - S^* - T_S(S^*)$. Differentiating this expression leads to

$$\mathcal{U}_{MM}(M,Y;r) = \left[1 - (1+\tau_S)\frac{\partial S^*}{\partial M}\right]U_{CC}(C,Y;r),$$

so that

$$\frac{M\mathcal{U}_{MM}}{\mathcal{U}_M} = \left[1 - (1+\tau_S)\frac{\partial S^*}{\partial M}\right]\frac{Y - T_Y(Y)}{C}\frac{CU_{CC}}{U_C} = -\left[1 - (1+\tau_S)\frac{\partial S^*}{\partial M}\right]\frac{\zeta_C}{s_C}.$$

Differentiating the FOC leads to $(1+\tau_S)\left[1 - (1+\tau_S)\frac{\partial S^*}{\partial M}\right]U_{CC} = V_{SS}\frac{\partial S^*}{\partial M}$, or

$$\frac{\partial S^*}{\partial M} = \frac{U_{CC}}{(1+\tau_S)U_{CC} + \frac{1}{(1+\tau_S)}V_{SS}} = \frac{1}{1+\tau_S}\cdot\frac{\frac{(1+\tau_S)S}{C}\frac{CU_{CC}}{U_C}}{\frac{(1+\tau_S)S}{C}\frac{CU_{CC}}{U_C} + \frac{SV_{SS}}{V_S}} = \frac{1}{1+\tau_S}\cdot\frac{\frac{1-s_C}{s_C}\zeta_C}{\frac{1-s_C}{s_C}\zeta_C + \zeta_S}.$$

Substituting into the previous expression gives

$$\tilde{\zeta}_M = \frac{\zeta_S}{\frac{1-s_C}{s_C}\zeta_C + \zeta_S}\frac{\zeta_C}{s_C} = \frac{1}{\frac{1-s_C}{\zeta_S} + \frac{s_C}{\zeta_C}}.$$

Similar calculations easily lead to

$$\frac{1}{\tilde{\zeta}_M} = \frac{s_C}{\zeta_C} + \frac{1-s_C}{\zeta_S}, \quad \tilde{\zeta}_Y = \zeta_Y - \left(1 - s_C\frac{\tilde{\zeta}_M}{\zeta_C}\right)\frac{s_C}{\zeta_C}\zeta_{CY}^2, \quad \tilde{\zeta}_{MY} = s_C\frac{\tilde{\zeta}_M}{\zeta_C}\zeta_{CY}. \tag{29}$$

Note that, when $s_C = 1$ (*Case 1*), the first-stage elasticities are equal to their primitive counterparts, i.e., $\tilde{\zeta}_M = \zeta_C$, $\tilde{\zeta}_{MY} = \zeta_{CY}$, and $\tilde{\zeta}_Y = \zeta_Y$. Next, using the map derived above between preference and behavioral elasticities, we have

$$\frac{\zeta_Y^I}{\zeta_S^H} = \frac{\frac{\hat{\zeta}_C\zeta_S}{s_S\hat{\zeta}_C + s_C\hat{\zeta}_S}}{\frac{\hat{\zeta}_C}{s_S\hat{\zeta}_C + s_C\hat{\zeta}_S}} = \zeta_S$$

and

$$\frac{\zeta_Y^H}{\zeta_C^H} = \frac{s_S\hat{\zeta}_C + s_C\hat{\zeta}_S}{\hat{\zeta}_S} = \frac{(1-s_C)\,\zeta_C + s_C\zeta_S}{\zeta_S + (1-s_C)\,\zeta_{CY}}$$

so that, using the identity $\zeta_Y^H = s_C\zeta_C^H + (1-s_C)\,\zeta_S^H$

$$\frac{\zeta_C}{\zeta_S} = \frac{1}{1-s_C}\frac{\zeta_Y^H}{\zeta_C^H} + \frac{\zeta_Y^H}{\zeta_C^H}\frac{\zeta_{CY}}{\zeta_S} - \frac{s_C}{1-s_C} = \frac{\zeta_S^H}{\zeta_C^H} + \frac{\zeta_Y^H}{\zeta_C^H}\frac{\zeta_{CY}}{\zeta_S},$$

or, alternatively,

$$\frac{\zeta_C}{\zeta_S} = \left[\frac{s_C}{1-s_C} + \frac{\zeta_S^H}{\zeta_C^H}\right]\left[1 + (1-s_C)\frac{\zeta_{CY}}{\zeta_S}\right] - \frac{s_C}{1-s_C} = \frac{\zeta_S^H}{\zeta_C^H}\left[\frac{1 + (1-s_C)\,\zeta_{CY}/\zeta_S}{1 - s_C\zeta_{CY}/\zeta_C}\right].$$

We can also write

$$\frac{\zeta_C}{\zeta_S} = \frac{\zeta_S^I}{\zeta_C^I}\left[1 - \left(1 - \frac{s_C}{\zeta_C}\zeta_{CY}\right)\zeta_Y^F\zeta_{CY}\right]$$

where ζ_Y^F is the Frisch elasticity of labor supply, which satisfies $1/\zeta_Y^F = \tilde{\zeta}_Y - \tilde{\zeta}_{MY}^2/\tilde{\zeta}_M = \zeta_Y - s_C\zeta_{CY}^2/\zeta_C$. To summarize, we have shown that

$$\zeta_S = \frac{\zeta_Y^I}{\zeta_S^H}, \quad \text{and} \quad \zeta_C = \frac{\zeta_Y^I}{\zeta_C^H} + \frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY}.$$

Finally, the elasticity of intertemporal substitution is defined by

$$\zeta_{IS} \equiv -\left.\frac{\partial\ln(S/C)}{\partial\ln(1+\tau_S)}\right|_{Y,U\text{ constant}}.$$

The expenditure minimization problem $\min_{C,S} C + (1+\tau_S)\,S$ s.t. $U(C,Y) + V(S) \geq \bar{U}$ yields the optimality condition $1+\tau_S = V_S/U_C$ along with the constant utility constraint. Differentiating

$1+\tau_S = V_S/U_C$ leads to

$$\frac{\partial\left(1+\tau_S\right)}{1+\tau_S} = -\zeta_S\frac{\partial S}{S}+\zeta_C\frac{\partial C}{C}$$

Differentiating $U\left(C,Y\right)+V\left(S\right)=\bar{U}$ yields $U_C\partial C = -V_S\partial S$, and therefore

$$\frac{\partial C}{C} = -\frac{1-s_C}{s_C}\frac{\partial S}{S}$$

Substituting the latter into the former yields

$$\frac{\partial\left(1+\tau_S\right)}{1+\tau_S} = -\frac{\partial S}{S}\frac{1}{s_C}\left[\left(1-s_C\right)\zeta_C+s_C\zeta_S\right],$$

and therefore

$$\zeta_{IS} = \left.\frac{\frac{\partial C}{C}-\frac{\partial S}{S}}{\frac{\partial(1+\tau_S)}{1+\tau_S}}\right|_{Y,U\text{ constant}} = \frac{1}{\left(1-s_C\right)\zeta_C+s_C\zeta_S},$$

which is the expression given in the text. $\square$

Proof (and Generalization) of equation (14). Differentiating the budget constraint $C\left(r\right)+S\left(r\right)+T_S\left(S\left(r\right)\right)=Y\left(r\right)-T_Y\left(Y\left(r\right)\right)$ yields $C'\left(r\right)+S'\left(r\right)\left(1+\tau_S\left(r\right)\right)+Y'\left(r\right)\left(1-\tau_Y\left(r\right)\right)$ or, after dividing by $Y\left(r\right)\left(1-\tau_Y\left(r\right)\right)$:

$$\frac{1-s_C\left(r\right)}{\rho_S\left(r\right)}+\frac{s_C\left(r\right)}{\rho_C\left(r\right)}=\frac{1}{\rho_Y\left(r\right)}.$$

We work with the limit as $r\to 1$: $\left(1-s_C\right)/\rho_S+s_C/\rho_C=1/\rho_Y$. Now, the optimal savings tax formula reads

$$\overline{\tau}_S=\frac{\zeta_S/\rho_S-\zeta_C/\rho_C+\zeta_{CY}/\rho_Y}{1-\zeta_S/\rho_S}=\frac{\zeta_S/\rho_S}{1-\zeta_S/\rho_S}\cdot\left[\frac{\zeta_S/\rho_S-\zeta_C/\rho_C+\zeta_{CY}/\rho_Y}{\zeta_S/\rho_S}\right].$$

But recall that

$$\frac{\zeta_C}{\zeta_S}=\frac{\zeta_S^H}{\zeta_C^H}+\frac{\zeta_Y^H}{\zeta_C^H}\frac{\zeta_{CY}}{\zeta_S}$$

and hence

$$\begin{aligned}\overline{\tau}_S=&\frac{\zeta_S/\rho_S}{1-\zeta_S/\rho_S}\left[1-\frac{\rho_S}{\rho_C}\left(\frac{\zeta_S^H}{\zeta_C^H}+\frac{\zeta_Y^H}{\zeta_C^H}\frac{\zeta_{CY}}{\zeta_S}\right)+\frac{\zeta_{CY}/\rho_Y}{\zeta_S/\rho_S}\right]\\=&\frac{\zeta_S/\rho_S}{1-\zeta_S/\rho_S}\left[1-\frac{\rho_S\zeta_S^H}{\rho_C\zeta_C^H}+\frac{\zeta_{CY}/\rho_Y}{\zeta_S/\rho_S}\left(1-\frac{\rho_Y\zeta_Y^H}{\rho_C\zeta_C^H}\right)\right].\end{aligned}$$

Substituting $\zeta_Y^H = s_C \zeta_C^H + (1-s_C)\,\zeta_S^H$ and $\frac{\rho_S}{\rho_Y} = s_C \frac{\rho_S}{\rho_C} + 1 - s_C$ yields $\frac{\rho_S}{\rho_Y} - \frac{\rho_S}{\rho_C}\frac{\zeta_Y^H}{\zeta_C^H} = (1-s_C)\left(1 - \frac{\rho_S \zeta_S^H}{\rho_C \zeta_C^H}\right)$ and therefore, after substituting $\zeta_S = \zeta_Y^I / \zeta_S^H$, we obtain

$$\overline{\tau}_S = \frac{\zeta_Y^I}{\rho_S \zeta_S^H - \zeta_Y^I}\left(1 - \frac{\rho_S \zeta_S^H}{\rho_C \zeta_C^H}\right)\left[1 + (1-s_C)\,\frac{\zeta_{CY}}{\zeta_S}\right]. \tag{30}$$

This expression generalizes equation (14) to the case $\zeta_{CY} \neq 0$. $\square$

Proof (and Generalization) of equation (13). We can rewrite the income tax formula as

$$1 - \overline{\tau}_Y = \frac{1 - \zeta_C/\rho_C + \zeta_{CY}/\rho_Y}{1 + \zeta_Y/\rho_Y - s_C \zeta_{CY}/\rho_C} = \frac{(1 - \zeta_S/\rho_S)\,(1 + \overline{\tau}_S)}{1 + \zeta_Y/\rho_Y - s_C \zeta_{CY}/\rho_C}.$$

The previous proof, along with the equalities $\zeta_Y^H = s_C \zeta_C^H + (1-s_C)\,\zeta_S^H$ and $\frac{\rho_S}{\rho_Y} = s_C \frac{\rho_S}{\rho_C} + 1 - s_C$, imply that the numerator equals

$$\left(1 - \frac{\zeta_S}{\rho_S}\right)(1 + \overline{\tau}_S) = 1 - \frac{\zeta_Y^I}{\rho_C \zeta_C^H} + \left(1 - \frac{\rho_S \zeta_S^H}{\rho_C \zeta_C^H}\right)(1-s_C)\,\frac{\zeta_{CY}}{\rho_S}.$$

Now, using $(1-s_C)\,\frac{1}{\rho_S} = \frac{1}{\rho_Y} - s_C \frac{1}{\rho_C}$ and $\frac{(1-s_C)\zeta_S^H}{\rho_C \zeta_C^H} = \frac{\zeta_Y^H}{\rho_C \zeta_C^H} - s_C \frac{1}{\rho_C}$, we have that $\left(1 - \frac{\rho_S \zeta_S^H}{\rho_C \zeta_C^H}\right)(1-s_C)\,\frac{\zeta_{CY}}{\rho_S} = \frac{\zeta_{CY}}{\rho_Y}\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)$, and therefore

$$\left(1 - \frac{\zeta_S}{\rho_S}\right)(1 + \overline{\tau}_S) = \frac{1}{\rho_Y \zeta_Y^H}\left[\rho_Y \zeta_Y^H - \zeta_Y^I + \left(\zeta_Y^I + \zeta_Y^H \zeta_{CY}\right)\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)\right].$$

For the denominator, we write

$$\zeta_Y = \tilde{\zeta}_Y + s_C \frac{\zeta_{CY}^2}{\zeta_C} - s_C^2 \frac{\zeta_{CY}^2}{\zeta_C^2}\tilde{\zeta}_M = \frac{1 - \zeta_Y^I}{\zeta_Y^H} + s_C \frac{\zeta_{CY}^2}{\zeta_C} + s_C \frac{\zeta_{CY}}{\zeta_C}\frac{\zeta_Y^I}{\zeta_Y^H}$$

so that

$$\begin{aligned} 1 + \frac{\zeta_Y}{\rho_Y} - s_C \frac{\zeta_{CY}}{\rho_C} =& 1 + \frac{1 - \zeta_Y^I}{\rho_Y \zeta_Y^H} + s_C \frac{\zeta_{CY}}{\zeta_C}\left(\frac{\zeta_Y^I + \zeta_Y^H \zeta_{CY}}{\rho_Y \zeta_Y^H} - \frac{\zeta_C}{\rho_C}\right) \\ =& 1 + \frac{1 - \zeta_Y^I}{\rho_Y \zeta_Y^H} + s_C \frac{\zeta_{CY}}{\zeta_C}\frac{\zeta_Y^I + \zeta_Y^H \zeta_{CY}}{\rho_Y \zeta_Y^H}\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right). \end{aligned}$$

Combining these expressions and rearranging terms yields

$$\overline{\tau}_Y = \frac{1 - \left(1 - s_C \frac{\zeta_{CY}}{\zeta_C}\right)\left(\zeta_Y^I + \zeta_Y^H \zeta_{CY}\right)\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)}{1 + \rho_Y \zeta_Y^H - \zeta_Y^I + s_C \frac{\zeta_{CY}}{\zeta_C}\left(\zeta_Y^I + \zeta_Y^H \zeta_{CY}\right)\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)}.$$

This expression generalizes equation (13) to the case $\zeta_{CY} \neq 0$. If $s_C = 1$, then $\rho_C = \rho_Y$ and $\zeta_C^H = \zeta_Y^H$ and we recover $\overline{\tau}_Y = \tau_Y^{Saez}$. If $s_C = 0$, then

$$\overline{\tau}_Y = \overline{\tau}_Y^{Saez}\left[1 - \left(\zeta_Y^I + \zeta_Y^H \zeta_{CY}\right)\left(1 - \frac{\rho_Y \zeta_Y^H}{\rho_C \zeta_C^H}\right)\right], \tag{31}$$

so that the last term serves as a downwards or upwards adjustment factor relative to $\overline{\tau}_Y^{Saez}$; for general s_C, this adjustment factor represents a lower bound on the departure of $\overline{\tau}_Y$ from $\overline{\tau}_Y^{Saez}$. □

Proof (and Generalization) of equation (15). We have shown that $\overline{\tau}_S > 0$ iff $\zeta_S^H/\zeta_C^H < \rho_C/\rho_S$. If $\zeta_{CY} = 0$, then we have $\zeta_S^H/\zeta_C^H = \zeta_C/\zeta_S$, so that the condition can be rewritten as

$$\frac{\zeta_C}{\zeta_S} < \frac{\rho_C}{\rho_S}.$$

If in addition $s_C = 0$, we can use the expressions $\zeta_C = 1/\zeta_{IS}$ and $\zeta_S = \zeta_Y^I/\zeta_Y^H$ to rewrite this condition as

$$\zeta_{IS} > \frac{\rho_Y}{\rho_C}\frac{\zeta_Y^H}{\zeta_Y^I}.$$

Note that, in the other polar case $s_C = 1$, we would get a lower bound on the EIS: we would now have $\zeta_S = 1/\zeta_{IS}$ and $\zeta_C = \zeta_Y^I/\zeta_Y^H$, and hence $\zeta_{IS} < (\rho_Y/\rho_S)\left(\zeta_Y^H/\zeta_Y^I\right)$. More generally, consider arbitrary ζ_{CY} and s_C. Using the expressions $\zeta_S = \zeta_Y^I/\zeta_S^H$ and $\zeta_C = \zeta_Y^I/\zeta_C^H + (\zeta_Y^H/\zeta_C^H)\zeta_{CY}$, the condition $\zeta_S^H/\zeta_C^H < \rho_C/\rho_S$ can be rewritten as

$$\zeta_C - \frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY} < \frac{\rho_C}{\rho_S}\zeta_S.$$

Now recall that the EIS is given by $\zeta_{IS} = 1/\left[s_C \zeta_S + (1 - s_C)\zeta_C\right]$. If $s_C < 1$, we can substitute for ζ_C in the previous expression to get the equivalent condition

$$\frac{1}{(1 - s_C)}\left[\frac{1}{\zeta_{IS}} - s_C \zeta_S\right] < \frac{\rho_C}{\rho_S}\zeta_S + \frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY},$$

which can be restated, using $(1-s_C)\frac{1}{\rho_S}=\frac{1}{\rho_Y}-s_C\frac{1}{\rho_C}$ and $\zeta_S=\zeta_Y^I/\zeta_S^H$, as

$$\zeta_{IS}>\left[\frac{\rho_C}{\rho_Y}\frac{\zeta_Y^I}{\zeta_S^H}+(1-s_C)\frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY}\right]^{-1}. \tag{32}$$

Analogously, if $s_C>0$, we can substitute for ζ_S to get

$$\zeta_C<\frac{\rho_C}{\rho_S}\frac{1}{s_C}\left[\frac{1}{\zeta_{IS}}-(1-s_C)\zeta_C\right]+\frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY},$$

which can be restated, using $\zeta_C=\zeta_Y^I/\zeta_C^H+(\zeta_Y^H/\zeta_C^H)\zeta_{CY}$, as

$$\zeta_{IS}<\left[\frac{\rho_S}{\rho_Y}\frac{\zeta_Y^I}{\zeta_C^H}+(1-s_C)\frac{\zeta_Y^H}{\zeta_C^H}\zeta_{CY}\right]^{-1}. \tag{33}$$

Thus, if $s_C=0$, we obtain a lower bound on ζ_{IS} above which savings should be taxed; this condition coincides with (15) if $\zeta_{CY}=0$, since we then have $\zeta_S^H=\zeta_Y^H$. If $s_C=1$, we obtain an upper bound on ζ_{IS} that generalizes the expression obtained above. If $s_C\in(0,1)$, both of these conditions are equivalent and can be used interchangeably. □

Proof of Theorem 2. Suppose that $\zeta_Y^I+\zeta_Y^H\zeta_{CY}$ is strictly positive (i.e., the first-period utility function is not quasilinear in consumption). Then expression (31) implies that $\overline{\tau}_Y<\overline{\tau}_Y^{Saez}$ if and only if $\frac{\rho_Y\zeta_Y^H}{\rho_C\zeta_C^H}<1$. Moreover, equation (30) implies that $\overline{\tau}_S>0$ if and only if $\frac{\rho_S\zeta_S^H}{\rho_C\zeta_C^H}<1$. Now, recall that $s_C\zeta_C^H+(1-s_C)\zeta_S^H=\zeta_Y^H$ and $(1-s_C)/\rho_S+s_C/\rho_C=1/\rho_Y$. Therefore, we have $\rho_S\zeta_S^H<\rho_C\zeta_C^H$ if and only if

$$\frac{1}{(1-s_C)\frac{1}{\rho_S}}\left[\zeta_Y^H-s_C\zeta_C^H\right]<\rho_C\zeta_C^H,$$

or $\zeta_Y^H-s_C\zeta_C^H<\rho_C\zeta_C^H\left[1/\rho_Y-s_C/\rho_C\right]$, which in turn reduces to $\rho_Y\zeta_Y^H<\rho_C\zeta_C^H$. This concludes the proof. □

B Relationship with the Previous Literature

B.1 Relationship to Ferey, Lockwood, and Taubinsky (2021)

Following Saez (2002) and Gerritsen et al. (2020), a recent paper by Ferey, Lockwood, and Taubinsky (2021, henceforth FLT) estimates optimal savings taxes emphasizing different sufficient statistics, namely the cross-sectional variation of savings with income net of the causal effect of income on savings. This sufficient statistic decomposes the cross-sectional variation in savings into a component due to cross-sectional variation in income and a component due to cross-sectional variation in preferences, and identifies the latter as the key driver of optimal savings taxes, in line with Atkinson and Stiglitz (1976).

We derive below the precise relationship between our optimal tax formulas and this alternative representation. We show that both representations are equivalent *provided that* $s_C(r) > 0$, i.e., consumption takes up a non-negligible fraction of after-tax income. However, if—as we argued is the empirically relevant case—consumption has a strictly thinner tail than savings and $\lim_{r\to 1} s_C(r) = 0$, then the identification of FLT breaks down for top earners; that is, their additional sufficient statistics lose their informational content. This is consistent with the observation that information from savings remains informative only if $s_C(r) > 0$, i.e., consumption is non-vanishing at the top of the income distribution.

Hence, while FLT's representation offers additional insight into the identification of preference elasticities along the bulk of the tax schedule, their identification breaks down towards the top of the income distribution and they cannot offer prescriptions on top savings taxes unless the equality $\rho_C = \rho_Y$ holds empirically. By contrast, our result based on the Pareto tails of consumption offers an alternative that identifies top income taxes even if $\lim_{r\to 1} s_C(r) = 0$. Both papers are therefore complementary, in the sense that we are able to offer prescriptions for income and savings taxes on top earners, on which their sufficient statistics are unable to shed light.

Proof. Formally, given the tax schedule, define $S(Y,r)$ as the optimal savings of a household of rank r given income Y, defined by solving the FOC for savings $(1+\tau_S)U_C = V'$ and the household budget constraint $C + S = Y - T(Y,S)$, where $\tau_S = \frac{\partial T(Y,S)}{\partial S}$ and $\tau_Y = \frac{\partial T(Y,S)}{\partial Y}$, for C and Y. Taking derivatives, we decompose $S'(r)$ as follows:

$$\frac{S'(r)}{S(r)}(1-r) \;=\; \frac{\partial \ln S(Y,r)}{\partial \ln Y}\frac{Y'(r)}{Y(r)}(1-r) - \frac{\partial \ln S(Y,r)}{\partial (1-r)}.$$

Rearranging terms and noting that $\frac{S'(r)}{S(r)}(1-r) = \frac{1}{\rho_S(r)}$ and $\frac{Y'(r)}{Y(r)}(1-r) = \frac{1}{\rho_Y(r)}$ we obtain

$$-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)} \;=\; \frac{1}{\rho_S(r)} - \frac{\partial \ln S(Y,r)}{\partial \ln Y}\frac{1}{\rho_Y(r)}.$$

Hence the elasticity $\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$ captures the effect of preference heterogeneity on savings for a given income and corresponds to $s'_{het}\cdot\frac{(1-r)}{S}$ in FLT, while the elasticity $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$ measures the causal effect of income on savings and corresponds to $s'_{inc}\cdot\frac{Y}{S}$ in FLT. Also recall that $s_C(r) = \frac{C(r)}{(1-\tau_Y(r))Y(r)}$ and define $s_S(r) \equiv \frac{(1+\tau_S(r))S(r)}{(1-\tau_Y(r))Y(r)}$. We characterize $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$ and $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$ using perturbation arguments:[23]

$$\frac{\partial \ln S(Y,r)}{\partial \ln Y} = \frac{\zeta_C(r)\left(1 - s_C(r)\frac{\zeta_{CY}(r)}{\zeta_C(r)}\right)}{s_S(r)\zeta_C(r) + s_C(r)\zeta_S(r)}$$

and

$$-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)} = \frac{s_C(r)}{s_S(r)\zeta_C(r) + s_C(r)\zeta_S(r)}\left(\frac{\zeta_S(r)}{\rho_S(r)} - \frac{\zeta_C(r)}{\rho_C(r)} + \frac{\zeta_{CY}(r)}{\rho_Y(r)}\right).$$

Hence, whenever $s_C(r) > 0$, $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$ is strictly decreasing in $\frac{\zeta_S(r)}{\zeta_C(r)}$ and thus offers an additional identifying moment for the preference elasticities. Likewise $-\frac{\partial \ln S(Y,r)}{\partial(1-r)}$ is strictly increasing in $\frac{\zeta_S(r)}{\zeta_C(r)}$, for given preferences, spending shares, and Pareto tails. However, if $\lim_{r\to 1} s_C(r) = 0 = 1 - \lim_{r\to 1} s_S(r)$, then $\lim_{r\to 1}\frac{\partial \ln S(Y,r)}{\partial \ln Y} = 1$ and $\lim_{r\to 1}\left(-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}\right) = 0$, regardless of the other parameters, which confirms that the identifying power of $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$ vanishes when $\lim_{r\to 1} s_C(r) = 0$ at the top of the income distribution.

The main representation of optimal savings taxes in FLT (equation (19)) can then be translated as follows into the notation of our model:

$$\frac{\tau_S(r)}{1+\tau_S(r)} = \frac{-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}}{-\frac{\partial \ln S(Y,r)}{\partial \ln(1+\tau_S)}\Big|_{Y,T(Y,S)\text{ constant}}}\mathbb{E}\left[1 - \hat{g}(r')\,|r' \geq r\right].$$

Here, $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$ is as defined above, and $-\frac{\partial \ln S(Y,r)}{\partial \ln(1+\tau_S)}\Big|_{Y,T(Y,S)\text{ constant}}$ represents a compensated elasticity of savings to savings taxes, holding constant the households income Y and total tax burden $T(Y,S)$. A simple perturbation argument shows that[24]

$$-\frac{\partial \ln S(Y,r)}{\partial \ln(1+\tau_S)}\Big|_{Y,T(Y,S)\text{ constant}} = \frac{s_C(r)}{s_S(r)\zeta_C(r) + s_C(r)\zeta_S(r)}$$

where $s_S(r)\zeta_C(r) + s_C(r)\zeta_S(r)$ represents the inverse of the inter-temporal elasticity of substitution. Therefore $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$ and $-\frac{\partial \ln S(Y,r)}{\partial \ln(1+\tau_S)}\Big|_{Y,T(Y,S)\text{ constant}}$ both converge to zero if $\lim_{r\to 1} s_C(r)$, but their ratio converges to a finite constant $\frac{\zeta_S(r)}{\rho_S(r)} - \frac{\zeta_C(r)}{\rho_C(r)} + \frac{\zeta_{CY}(r)}{\rho_Y(r)}$, which is the same as $\frac{B_S(r)-B_C(r)}{B_S(r)B_C(r)}$ in our model when $r \to 1$. By contrast, our representation implies $\frac{\tau_S(r)}{1+\tau_S(r)} = \frac{B_S(r)-B_C(r)}{B_S(r)}$. The two representations are therefore identical if the remaining term, $\mathbb{E}\left[1 - \hat{g}(r')\,|r' \geq r\right]$, converges to

[23] Consider a perturbation $(\partial C, \partial Y, \partial S)$ along the households' FOC for savings, $\zeta_C\frac{\partial C}{C} - \zeta_{CY}\frac{\partial Y}{Y} = \zeta_S\frac{\partial S}{S}$, and budget constraint $s_C\frac{\partial C}{C} + s_S\frac{\partial S}{S} = \frac{\partial Y}{Y}$. Solving these two equations for $\frac{\partial S}{S}/\frac{\partial Y}{Y}$ yields $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$. Totally differentiating the FOC for savings $(1+\tau_S)U_C = V'$ and using Lemma ?? to substitute out $\frac{\partial \tau_S}{1+\tau_S} + \frac{U_{Cr}}{U_C}$ yields the expression for $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$.

[24] Consider a perturbation $(\partial C, \partial S, \partial \tau_S)$ along the households' FOC for savings, $\zeta_C\frac{\partial C}{C} + \frac{\partial \tau_S}{1+\tau_S} = \zeta_S\frac{\partial S}{S}$, that keeps household utility unchanged: $U_C\partial C + \beta V'\partial S = 0$, or $s_C\frac{\partial C}{C} = -s_S\frac{\partial S}{S}$. Solving these two equations for $-\frac{\partial S}{S}/\frac{\partial \tau_S}{1+\tau_S}$ yields $-\frac{\partial \ln S(Y,r)}{\partial \ln(1+\tau_S)}\Big|_{Y,T(Y,S)\text{ constant}}$.

$B_C(r)$. The term $\mathbb{E}[1-\hat{g}(r')\,|r'\geq r]$ in FLT captures a mix of Pareto weights (which are vanishing at the top) and changes in tax revenue in response to income tax changes, which do not have a straight-forward mapping to our model. However, both the discussion in FLT and the equivalence between the two models suggests that $\lim_{r\to 1}\mathbb{E}[1-\hat{g}(r')\,|r'\geq r]=\lim_{r\to 1}B_C(r)$. In addition, we can rewrite equation (18) in FLT as

$$\frac{\tau_Y}{1-\tau_Y}=\left\{\frac{1}{\zeta_Y^H}-s_s\frac{\partial\ln S(Y,r)}{\partial\ln Y}\left(\frac{\rho_Y}{\rho_S}\zeta_S-\zeta_C\left(\frac{\rho_Y}{\rho_C}-\frac{\zeta_{CY}}{\zeta_C}\right)\right)\right\}\frac{1}{\rho_Y}\mathbb{E}\left[1-\hat{g}(r')\,|r'\geq r\right]$$

where the compensated income elasticity ζ_Y^H satisfies[25]

$$\frac{1}{\zeta_Y^H}=\zeta_Y-\zeta_{CY}+(\zeta_C-s_C\zeta_{CY})\frac{\zeta_S+s_S\zeta_{CY}}{s_S\zeta_C+s_C\zeta_S}.$$

Substituting $s_s\frac{\partial\ln S(Y,r)}{\partial\ln Y}=\frac{s_S\zeta_C(1-s_C\zeta_{CY}/\zeta_C)}{s_S\zeta_C+s_C\zeta_S}$ then allows us to evaluate the above expression in the limit as $r\to 1$: If $\lim_{r\to 1}s_C=1$ (Case 1), it follows that $\frac{1}{\zeta_Y^H}=\zeta_Y-\zeta_{CY}+\zeta_C-\zeta_{CY}$ and $s_s\frac{\partial\ln S(Y,r)}{\partial\ln Y}\to 0$, so

$$\frac{\tau_Y}{1-\tau_Y}=\{\zeta_Y-\zeta_{CY}+\zeta_C-\zeta_{CY}\}\frac{1}{\rho_Y}\mathbb{E}\left[1-\hat{g}(r')\,|r'\geq r\right].$$

If $\lim_{r\to 1}s_C=0$ (Case 2), it follows that $\frac{1}{\zeta_Y^H}=\zeta_Y+\zeta_S$ and $s_s\frac{\partial\ln S(Y,r)}{\partial\ln Y}\to 1$, so

$$\frac{\tau_Y}{1-\tau_Y}=\left\{\zeta_Y+\zeta_C\left(\frac{\rho_Y}{\rho_C}-\frac{\zeta_{CY}}{\zeta_C}\right)\right\}\frac{1}{\rho_Y}\mathbb{E}\left[1-\hat{g}(r')\,|r'\geq r\right].$$

Finally, if $\lim_{r\to 1}s_C(r)\in(0,1)$ (Case 3), $\frac{\rho_Y}{\rho_S}=\frac{\rho_Y}{\rho_C}=1$, and $\frac{1}{\zeta_Y^H}-s_s\frac{\partial\ln S(Y,r)}{\partial\ln Y}\left(\frac{\rho_Y}{\rho_S}\zeta_S-\zeta_C\left(\frac{\rho_Y}{\rho_C}-\frac{\zeta_{CY}}{\zeta_C}\right)\right)$ converges to $\zeta_Y-\zeta_{CY}+\zeta_C-s_C\zeta_{CY}$. In all three cases, equation (18) in FLT yields

$$\frac{\overline{\tau}_Y}{1-\overline{\tau}_Y}=\lim_{r\to 1}A(r)\,\mathbb{E}\left[1-\hat{g}(r')\,|r'\geq r\right]$$

where $A(r)=\frac{U_{Cr}}{U_C}-\frac{U_{Yr}}{U_Y}$ as defined above, and again the expression for the top labor wedge is equivalent to ours when $\lim_{r\to 1}\mathbb{E}[1-\hat{g}(r')\,|r'\geq r]=\lim_{r\to 1}B_C(r)$. $\square$

[25]Consider a perturbation $(\partial C,\partial Y,\partial S,\partial\tau_Y)$ along the households' FOCs for income $-\frac{\partial\tau_Y}{1-\tau_Y}=(\zeta_Y-\zeta_{CY})\frac{\partial Y}{Y}+(\zeta_C-s_C\zeta_{CY})\frac{\partial C}{C}$, and savings $\zeta_C\frac{\partial C}{C}-\zeta_{CY}\frac{\partial Y}{Y}=\zeta_S\frac{\partial S}{S}$ that keeps household utility unchanged: $U_C\partial C+U_Y\partial Y+\beta V'\partial S=0$, or $s_C\frac{\partial C}{C}+s_S\frac{\partial S}{S}=\frac{\partial Y}{Y}$. Solving these three equations for $-\frac{\partial Y}{Y}/\frac{\partial\tau_Y}{1-\tau_Y}$ yields ζ_Y^H.

B.2 Relationship to Scheuer and Slemrod (2021)

We now study a special case of the general environment of Section 5.2 that allows for heterogeneous initial capital holdings, and thus breaks the equality between the Pareto coefficients on income and wealth that the budget constraint imposes in our baseline model. The setting is the same as in our baseline model of Section 2, except that agents also receive an exogenous endowment $Z(r)$ that is strictly increasing in r. This framework nests that of Scheuer and Slemrod (2021), who assume that preferences satisfy the restrictions of Atkinson and Stiglitz (1976), that is, separable between consumption and income and homogeneous across consumers.

We show that if endowments have a strictly thinner tail than consumption and income, then the top income and savings taxes are the same as in our baseline model. Intuitively, endowments simply do not matter at the top of the distribution. When instead endowments have a thicker upper tail than income, inequality is mostly driven by inherited wealth and labor income becomes a negligible fraction of top earner's incomes. If, as in Scheuer and Slemrod (2021), endowments and consumption have an equal tail and preferences are separable, the solution for both labor and savings taxes is interior. However, this result is "knife-edge": As soon as consumption and income are complementary, it is optimal to impose arbitrarily large labor wedges on top earners. In the empirically plausible case where $\rho_Z = \rho_S < \rho_Y < \rho_C$, optimal taxes are just as stark: since the labor income and consumption of top earners are negligible, redistribution from the top is implemented through savings taxes that become arbitrarily large, and are accompanied by arbitrarily large earnings subsidies.

To summarize, the model with endowments substantially changes implications for optimal labor and savings taxes by shifting the burden of redistributive taxation from income to savings taxes when endowments become the main source of income for the top income earners.

Proof. Consider the same setting as in our baseline model, but suppose in addition that agents also receive an exogenous rank-specific endowment $Z(r)$. Since income and savings are taxed and hence observable, consumption is assumed to be unobserved. An agent with rank r then consumes $C(r, r') = C(r') + Z(r) - Z(r')$ when announcing type r'. Define the indirect utility function $W(r) = U(C(r), Y(r); r) + V(S(r))$, where we assume for simplicity that the second-period utility function is homogeneous across consumers. The planner's problem is stated as follows:

$$K(v_0) = \min_{\{C(r),Y(r),S(r)\}} \int_0^1 (C(r) - Y(r) + S(r))\, dr$$

subject to

$$\int_0^1 \omega(r) W(r) \, dr \geq v_0$$

$$W(r) = U(C(r), Y(r); r) + V(S(r))$$

$$W'(r) = U_C(C(r), Y(r); r) Z'(r) + U_r(C(r), Y(r); r).$$

Following analogous steps as in our baseline setting to solve this problem, we obtain the same characterization of optimal labor and savings wedges as in Section 4.4, except that we must adjust the definition of the incentive-adjustments and the marginal benefits of redistributing income, consumption, and savings B_Y, B_C, and B_S as follows:

$$\begin{aligned}
B_C(r) &= \mathbb{E}\left[\frac{M_C(r')}{M_C(r)} | r' \geq r\right] - \frac{\mathbb{E}\left[\frac{M_C(r')}{M_C(r)}\right] \mathbb{E}\left[\omega(r') U_C(r') \frac{M_C(r')}{M_C(r)} | r' \geq r\right]}{\mathbb{E}\left[\omega(r') U_C(r') \frac{M_C(r')}{M_C(r)}\right]} \\
B_Y(r) &= \mathbb{E}\left[\frac{M_Y(r')}{M_Y(r)} | r' \geq r\right] - \frac{\mathbb{E}\left[\frac{M_Y(r')}{M_Y(r)}\right] \mathbb{E}\left[\omega(r') (-U_Y(r')) \frac{M_Y(r')}{M_Y(r)} | r' \geq r\right]}{\mathbb{E}\left[\omega(r') (-U_Y(r')) \frac{M_Y(r')}{M_Y(r)}\right]} \\
B_S(r) &= \mathbb{E}\left[\frac{M_S(r')}{M_S(r)} | r' \geq r\right] - \mathbb{E}\left[\frac{M_S(r')}{M_S(r)}\right] \mathbb{E}\left[\omega(r') | r' \geq r\right]
\end{aligned}$$

with

$$\begin{aligned}
M_C(r) &= \frac{1}{U_C(r)} \exp\left[-\int_r^1 \left(\frac{U_{Cr}(r')}{U_C(r')} + \frac{U_{CC}(r')}{U_C(r')} Z'(r')\right) dr'\right] \\
M_Y(r) &= \frac{1}{-U_Y(r)} \exp\left[-\int_r^1 \left(\frac{U_{Yr}(r')}{U_Y(r')} + \frac{U_{CY}(r')}{U_Y(r')} Z'(r')\right) dr'\right] \\
M_S(r) &= \frac{1}{V'(S(r))}.
\end{aligned}$$

These marginal benefits converge to

$$\begin{aligned}
\lim_{r \to 1} B_C(r) &= \left[1 - (1 - s_Z) \frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y}\right]^{-1} = \frac{\bar{B}_C}{1 + \bar{B}_C s_Z \zeta_C / \rho_C} \\
\lim_{r \to 1} B_Y(r) &= \left[1 + \frac{\zeta_Y}{\rho_Y} - (1 - s_Z) s_C \frac{\zeta_{CY}}{\rho_C}\right]^{-1} = \frac{\bar{B}_Y}{1 + \bar{B}_Y s_Z s_C \zeta_{CY} / \rho_C} \\
\lim_{r \to 1} B_S(r) &= \left[1 - \frac{\zeta_S}{\rho_S}\right]^{-1} = \bar{B}_S,
\end{aligned}$$

where $s_Z = \lim_{r \to 1} \frac{Z'(r)}{C'(r)} = \frac{\rho_C}{\rho_Z} \lim_{r \to 1} \frac{Z(r)}{C(r)}$ represents the marginal increase in consumption scaled by the marginal increase in endowment at the top of the income (and endowment) distribution, and where $\bar{B}_C$, $\bar{B}_Y$, and $\bar{B}_S$ correspond to the marginal benefits of redistributing consumption, income, and savings in the baseline model without endowments.

The budget constraint implies that $\min\{\rho_Y, \rho_Z\} = \min\{\rho_C, \rho_S\}$, which allows us to distinguish different scenarios: *1.* Endowments have a thinner Pareto tail than income ($\rho_Y < \rho_Z$ and $s_Z s_C = 0$) and/or preferences are separable ($\zeta_{CY} = 0$); *2.* Endowments and income have equal Pareto tails ($\rho_Y = \rho_Z$), and consumption and income are complementary ($\zeta_{CY} > 0$); *3.* Endowments have a thicker Pareto tail than income ($\rho_Y > \rho_Z$), and consumption and income are complementary ($\zeta_{CY} > 0$).

In *Case* 1, $\lim_{r\to 1} B_Y(r)$ remains the same as in our baseline model, and hence endowments only affect the combined wedge $\frac{1-\overline{\tau}_Y}{1+\overline{\tau}_S} = \frac{1-\zeta_S/\rho_S}{1+\zeta_Y/\rho_Y}$ through their effect on the Pareto tail of savings. The thickness of the Pareto tail of consumption and endowments then governs the limit of $B_C(r)$: Specifically, if endowments have a thinner tail than consumption ($\rho_C < \rho_Z$), then $s_Z = 0$ and the top income and savings taxes are the same as in our baseline model. Intuitively, if endowments have a strictly thinner tail than consumption and income, then they simply do not matter at the top of the distribution: Top earners' endowments are negligible compared to their consumption and labor income. If instead endowments have the same tail as consumption ($\rho_C = \rho_Z$), then $s_Z > 0$, resulting in a shift from income to savings taxes. This shift can go so far as to make it optimal to subsidize income, and if endowments have a strictly thicker tail than consumption ($\rho_C > \rho_Z$), then $B_C(r) \to 0$ and earnings subsidies, along with savings taxes, become arbitrarily large for top income earners.

In *Case 2.*, $0 < s_Z s_C < \infty$ and the combined wedge is strictly lower than in the baseline model. If consumption has the same Pareto coefficient as income and endowments ($\rho_C = \rho_Y = \rho_Z$), then s_Z and s_C are both finite, so that the wedges $\overline{\tau}_Y$ and $\overline{\tau}_S$ are finite. The introduction of endowments reduces both B_Y and B_C, resulting in a strictly higher savings wedge and a lower combined wedge than in the baseline model; the labor wedge is reduced whenever $s_C \frac{\zeta_{CY}}{\rho_C} \frac{\bar{B}_Y}{\bar{B}_C} < 1$. If instead consumption has a strictly thinner tail ($\rho_Z = \rho_Y < \rho_C$) then $s_C \to 0$, $s_Z \to \infty$ and $B_C(r) \to 0$, resulting as before in arbitrarily large earnings subsidies and savings taxes at the top.

In *Case 3.*, $s_Z s_C = \infty$ and $B_Y(r) \to 0$, so that the combined wedge converges to 1. If consumption and endowments have equal tail coefficients ($\rho_C = \rho_Z$), then $0 < s_Z < \infty$ and $\overline{\tau}_S$ is finite and strictly larger than in our baseline economy, while the labor wedge becomes arbitrarily large ($\overline{\tau}_Y \to 1$). If $\rho_Z < \rho_C < \rho_Y$, we have both $s_Z \to \infty$ and $s_C \to \infty$ implying both arbitrarily large savings wedges (because $\rho_Z < \rho_C$) and arbitrarily large labor wedges (because $\rho_C < \rho_Y$). If $\rho_C = \rho_Y$, the savings wedge remains unbounded but the labor wedge is finite and given by $1 - \overline{\tau}_Y = \frac{\zeta_C}{s_C \zeta_{CY}}$. If $\rho_C > \rho_Y$, we obtain $\overline{\tau}_Y = -\infty$, making it optimal to have arbitrarily large savings taxes and earnings subsidies (but the combined wedge is always dominated by the savings wedge).

Intuitively, when endowments have a thicker upper tail than income, the planner's main tool for redistribution becomes the savings tax. Moreover, if consumption has a thinner tail than endowments (and savings), then a savings tax becomes non-distortionary at the top, and can therefore be arbitrarily large. The optimal labor wedge can then be understood by considering the spillover of labor income on savings: An increase in income allows households to both increase their spending on consumption and savings, and it induces them to substitute towards more consumption relative to savings. When s_C is high, the substitution effect dominates, which implies that an increase in income reduces savings, and hence the scope for redistribution through savings taxes. The planner then finds it optimal to tax income to reduce spill-overs to savings. In constrast, when s_C is low, the wealth effect of income on savings dominates, which makes it optimal to subsidize income. In the limit when $s_C \to 0$, and *a fortiori* when $\rho_C > \rho_Y$, the implied savings subsidy becomes arbitrarily large at the top. $\square$