

A Fair Day's Pay for a Fair Day's Work: Optimal Tax Design as Redistributional Arbitrage

Christian Hellwig and Nicolas Werquin

REVISED
January 20, 2023

WP 2022-03

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



FEDERAL RESERVE BANK *of* CHICAGO

*Working papers are not edited, and all opinions are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

A Fair Day's Pay for a Fair Day's Work: Optimal Tax Design as Redistributive Arbitrage*

Christian Hellwig

Toulouse School of Economics and CEPR

Nicolas Werquin

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

January 20, 2023

Abstract

We study optimal tax design based on the idea that policy-makers face trade-offs between multiple margins of redistribution. Within a Mirrleesian economy with labor income, consumption, and retirement savings, we derive a novel formula for optimal non-linear income and savings distortions based on redistributive arbitrage. We establish a sufficient statistics representation of the labor income and capital tax rates on top earners, which relies on comparing the Pareto tails of income and consumption. 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. We extend our representation of tax distortions based on redistributive arbitrage to economies with general preferences over an arbitrary number of periods and commodities, and we allow for return heterogeneity, age-contingent taxes, and stochastic evolution of types.

*We thank Arpad Abraham, Mark Aguiar, Gadi Barlevy, Marco Bassetto, Charlie Brendon, Steve Coate, 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, Mathieu Taschereau-Dumouchel, 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).

“Our Nation ... should be able to devise ways and means of insuring to all our able-bodied working men and women a fair day’s pay for a fair day’s work.”

Franklin D. Roosevelt, Message to Congress on Establishing Minimum Wages and Maximum Hours, 1937

1 Introduction

Originating with Mirrlees (1971), the problem of optimally designing taxes and social insurance programs is formalized as a trade-off between the social benefits of redistributing financial resources from richer to poorer households, and the efficiency costs of allocative distortions that such redistribution entails when these agents’ productivity types or inclination to work are not directly observable. One of the most celebrated achievements of this literature has been the derivation of the optimal tax rate on top income earners by Saez (2001) in terms of three observable statistics that give empirical meaning to this trade-off between incentives and redistribution: the elasticities of labor supply with respect to marginal tax rates and lump-sum transfers (substitution and income effects), and the Pareto coefficient of the tail of the income distribution, which measures the degree of top income inequality.

Despite its undisputed success in guiding tax policy design, the static Mirrleesian framework remains silent about a number of important policy questions. First, by focusing on a single consumption-labor supply margin, the model abstracts from the optimal design of policies that trade off between multiple policy tools. In practice, tax policies address concerns for redistribution along many dimensions: income, savings or consumption taxes, public social insurance programs for unemployment, healthcare or disability, subsidized provision of goods that are perceived to be essential necessities like housing, food, transportation, energy, education and even mass entertainment, or excess taxation of goods perceived to be luxuries. Moreover, the static Mirrleesian model implicitly assumes that the government is the sole channel of income redistribution. 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.

Second, abstracting from savings implies that we can always use the income distribution to proxy for consumption, or vice versa. However, this stark assumption is clearly rejected by empirical evidence which shows consumption to be substantially more evenly distributed than income (Toda and Walsh (2015)). The distinction between income and consumption inequality matters for quantitative conclusions of optimal tax policies: Applying Saez (2001)’s sufficient statistic representation, the optimal top income tax drops from 80% to 50% in our preferred calibration if we use

consumption- rather than income-based measures of inequality. In other words, the static representation of top optimal income taxes is based on an economic model that is inconsistent with the discrepancy between consumption and income inequality and provides no guidance about which measure is the most appropriate for estimating optimal income taxes. More generally, focusing exclusively on measures of income inequality may paint an incomplete picture of the link from allocations to welfare, which should be the key concern for optimal policy design.¹

In this paper, we develop a complementary perspective on optimal tax design, based on the premise that policy makers trade off between multiple dimensions of worker welfare and have potentially many policy tools at their disposal. Formally, in our baseline framework, we extend the canonical Mirrleesian tax design problem to allow for two separate consumption goods, which we interpret as “consumption” and “savings”. We consider a policy maker with a redistributive objective who designs income and savings taxes, while taking into account the households’ incentives to work, consume and save.

As our central result, we show that the optimal policy design obeys a simple principle of *redistributive arbitrage*. The policy maker has three means of extracting resources from the richest households: reducing their consumption, reducing their leisure (i.e., incentivizing them to work more), or reducing their wealth (taxing their savings). The optimal tax on labor income equalizes the resources that the policy maker can raise by asking the rich to work more—reducing their leisure—to the marginal resource gains from reducing their consumption. Similarly the optimal savings tax equalizes the marginal resource gains from reducing the richest households’ consumption to the marginal resource gains from reducing their savings. The same principle can be extended to any number of redistributive policy margins and thus serves as a guiding principle to design optimal redistributive policies along many different dimensions: The optimal policy equalizes the marginal resource gains from additional redistribution across different goods, since otherwise the tax designer would have an “arbitrage opportunity” by increasing redistribution along one margin and reducing it along a different one. Importantly, these redistributive arbitrages are constrained by the need to preserve the households’ incentives to work, consume, or save as intended by the policy maker.

Following Saez (2001), we express these marginal resource gains of redistributing consumption,

¹Consumption data provides an independent empirical test (and rejection) of the model underlying the representation of optimal taxes in the static model. This is an important caveat to the sufficient statistics approach: Its implications rely on the empirical validity of the underlying economic model. The empirical literature on risk-sharing emphasizes the importance of consumption, along with income data, for testing efficiency of risk-sharing arrangements 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.

leisure and savings—and hence the optimal income and savings taxes—in terms of observables, namely: the cross-sectional distribution (in particular, the Pareto tail coefficients) of each good, along with standard elasticity parameters that govern income and substitution effects. Abstracting from net complements or substitutes, the marginal gains from redistributing consumption are governed by the local Pareto coefficient of the consumption distribution and a risk-aversion parameter; the marginal gains from redistributing income or leisure are governed by the income distribution and labor supply elasticities; and the marginal gains from redistributing savings are governed by the wealth distribution and a risk aversion parameter over savings or second-period consumption. These representations clarify the respective roles of consumption, income and wealth inequality in determining optimal income and savings taxes.

The empirical evidence suggests that consumption has a thinner Pareto tail than income and savings. This implies that the consumption share of income converges to zero for top income earners, whose behavior thus reduces to a trade-off between leisure and savings. The static optimal tax formula of Saez (2001) then determines the combined wedge on labor income and savings. However, that does not answer how the combined wedge should be broken up into an income and a savings wedge. While the savings wedge can, in principle, be positive or negative, the fact that savings or wealth is substantially more unequally distributed than consumption implies that, for plausible levels of risk aversion, it is optimal to shift a significant share of the tax burden on top earners from income to savings. The static optimal tax formula overstates the marginal gains from redistribution and hence the optimal income taxes, because it fails to account for the fact that consumption is less unequally distributed than after-tax incomes and savings in the data.

Our calibration suggests that top savings taxes could be as high as 40%-50% of the level of savings, with a corresponding reduction in top income taxes from a static optimum of 80% at our baseline calibration towards 60%—almost doubling the top earners’ take-home pay. In a life-cycle context with a 30-year gap between the working period and retirement and a 5% annual return on savings, a savings tax of 40% corresponds to a 1.8% annual tax on accumulated wealth, or a 35% capital income tax. These estimates are thus in the same ballpark as existing proposals of annual wealth taxes in the range of 1% to 2% (Saez and Zucman (2019)). This shift from income towards savings taxes is a fairly robust feature of our quantitative results, and is driven by a combination of thinner consumption tails at the top of the income distribution and low consumption elasticities (risk-aversion and complementarity with labor effort). These features of the data imply that the marginal benefit of redistributing consumption is small compared to the marginal benefit of redistributing savings, making it optimal to shift part of the tax distortion

towards savings. They also suggest that capital income should still be taxed at a significantly lower rate than labor income.

We show that our baseline setting allows us to study two important rationales for taxing the capital of top earners: rate-of-return heterogeneity and the inverse Euler equation. In particular, our sufficient-statistic representation is such that the source (scale- vs. type-dependence) and extent of return heterogeneity does not affect our formulas and calibration. We then extend our results to a framework with one-dimensional preference types, but with general preferences over an arbitrary number of periods and commodities. We obtain a characterization of the optimal relative price distortions, or commodity taxes, as arbitraging between redistribution through one commodity vs. another. As an application of this generalized framework, we characterize the optimal income and capital taxes over the life cycle in terms of the age-dependent Pareto coefficients on income and consumption. We show that the accumulation of consumption inequality over the life cycle offers a new rationale for taxing the savings of working households, which is different from the rationale for taxing retirement savings in our baseline model.

While we are not aware of prior discussions or formalizations of redistributive arbitrage or related ideas in the economics literature on optimal tax design, the observation that redistributive policies act on many margins simultaneously is certainly not new to policy makers. For example, the labor movement's 19th century slogan "*A Fair Day's Pay for a Fair Day's Work*" epitomizes a joint concern for wages along with working hours or leisure of the working classes that permeated policy discussions over labor regulation and the concurrent emergence of the welfare state. The slogan was picked up by Roosevelt in a speech that led to the Fair Labor Standards Act (1938), which simultaneously introduced a minimum wage and regulations on total working hours. More recently, Aguiar and Hurst (2007) document a large increase in leisure inequality from the top to the bottom of the distribution since the 1960s in the U.S., mirroring the concurrent, well-documented and widely discussed rise in income inequality. Contemporary concerns for "work-life balance" suggest that high income earners today value leisure much like their working class peers in the 1930s or the 19th century, and employers acknowledge these concerns when granting workers leisure-related perks or non-pecuniary benefits, work-time flexibility or time-saving benefits like child-care services to working parents.²

²According to Cambridge online dictionary, work-life balance represents "*the amount of time you spend doing your job compared with the amount of time you spend with your family and doing things you enjoy.*" A 2011 report by the Council of Economic Advisors (Romer (2011)) reviews evidence suggesting that both employers and employees benefit from improved work-life balance: "*A study of more than 1,500 U.S. workers reported that nearly a third considered work-life balance and flexibility to be the most important factor in considering job offers. In another survey of two hundred human resource managers, two-thirds cited family-supportive policies and flexible hours as the single most important factor in attracting and retaining employees.*" The report itself is evidence that the joint importance of

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.³ By viewing tax policies as an arbitrage between different margins of redistribution, we generalize the representation of optimal income taxes obtained by Saez (2001) to a dynamic, or multiple-good, environment and derive a companion formula for optimal savings taxes. 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. In linking our characterization of optimal taxes to its empirical counterparts, we show that optimal top taxes rely not only on labor income data, as in the canonical Saez (2001) framework, but also on consumption data. We rely on the analyses of Toda and Walsh (2015), Blundell, Pistaferri, and Saporta-Eksten (2016), Straub (2019), and Buda et al. (2022) to argue that the Pareto tail of the distribution of consumption is significantly thinner than that of the income distribution.⁴

Gerritsen, Jacobs, Rusu, and Spiritus (2020), Schulz (2021), and Scheuer and Slemrod (2021), and especially Ferey, Lockwood, and Taubinsky (2021), are closest to our work. These papers characterize optimal savings taxes in models that are similar to ours, but use a different approach and obtain different results than us. First, our optimal tax formulas rely on a distinct set of perturbations and lead to redistributive arbitrage expressions that offer a unified perspective on the optimal design of taxes on multiple goods and bear little resemblance to the “ABC” expressions derived in these papers. Second, and most importantly, our representation maps to a different set of empirically observable sufficient statistics. Specifically, we show that the relative values of the

income and leisure for employee welfare is recognized at the highest levels of economic policy. The ongoing pandemic provides further evidence of the importance of leisure time for workers’ wellbeing: while the time savings and flexibility gains associated with remote work are greeted as a significant improvement in work-life balance, lack of access to child care and home schooling due to school closures are viewed as adding stress to working parents’ lives. Schieman et al. (2021) provide evidence from a sample of about 2000 Canadian households that reported work-life balance improved for most workers, excepted for those with children under the age of 12 who reported no change. Their cross-sectional controls further highlight that reported work-life balance appears to be as much affected by working hours and flexibility as it is by financial stress, but unrelated to income after controlling for other job characteristics.

³Several papers, such as 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 capital taxes to being linear. We abstract from several other extensions of the Atkinson-Stiglitz framework, such as multidimensional heterogeneity (Cremer, Pestieau, and Rochet (2003), Diamond and Spinnewijn (2011), Piketty and Saez (2013), and Saez and Stantcheva (2018)) or uncertainty (Diamond and Mirrlees (1978), Golosov, Kocherlakota, and Tsyvinski (2003), Farhi and Werning (2010), Shourideh (2012), Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2016), and Hellwig (2021)).

⁴This finding is consistent with Meyer and Sullivan (2017) who show that consumption inequality has seen a much more modest rise than income inequality since 2000.

Pareto tail coefficients on income and consumption, along with standard elasticity parameters, identify the underlying structure of preferences that pins down optimal income and capital taxes. While the alternative representation of Ferey, Lockwood, and Taubinsky (2021) offers additional insight into the identification of the preference elasticities along the bulk of the tax schedule, we show in Section 3.4 that their identification breaks down at the top of the income distribution. Thus, both papers are complementary, in the sense that ours offers prescriptions for top income and savings taxes, which is precisely where their sufficient statistics lose their identifying power. 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 case. On the other hand, these papers explore various microfoundations of return heterogeneity that are beyond the scope of our analysis. Finally, Scheuer and Slemrod (2021) derive a characterization of the capital tax rates on top earners when agents have exogenous endowments in addition to labor income. In contrast to our analysis, they take the labor income tax as given and restrict preferences to be separable between consumption and income, while the non-separability plays a critical role in our analysis. We discuss the relationship between our results and theirs in Section 5.

Outline of the Paper. We introduce our baseline model and derive theoretical formulas for optimal taxes in Section 2. In Section 3, we provide a sufficient-statistic representation of the optimal taxes. We calibrate the model and explore its quantitative implications in Section 4. Finally, Section 5 extends our results to a general framework.

2 Theory of Redistributive Arbitrage

2.1 Baseline 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 .⁵ They are represented as

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

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

⁵While 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.

approach 0 or ∞ . We interpret U as the first-period utility function, and V as the second-period utility function. The inter-temporal separability is inconsequential—we generalize our analysis to arbitrary preferences and commodities in Section 5.3. We discuss the interpretation of this baseline preference specification below.

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

$$\frac{\partial \ln(-U_Y/U_C)}{\partial r} \equiv \frac{U_{Yr}}{U_Y} - \frac{U_{Cr}}{U_C} < 0. \quad (1)$$

Furthermore, the marginal disutility of effort is decreasing in r , $U_{Yr}/U_Y < 0$. The MRS between consumption and savings $V_S(S; r)/U_C(C, Y; r)$ is monotonic in r for all (C, Y, S) , i.e.,

$$\frac{\partial \ln(V_S/U_C)}{\partial r} \equiv \frac{V_{Sr}}{V_S} - \frac{U_{Cr}}{U_C} \leq 0 \quad (2)$$

is either non-positive or non-negative everywhere.

The single-crossing condition (1) is standard (Mirrlees, 1971). It introduces a ranking of agents according to their preferences over leisure and consumption: On the margin, agents with higher rank r are more willing to work for a given consumption gain. The restriction $U_{Yr}/U_Y < 0$ implies that higher ranks r find it less costly to attain a given income level Y . This gives rise to a motive for redistributing effort from less to more productive agents, or equivalently leisure towards less productive agents; that is, redistribution “*from each according to his ability*”. The agent’s rank r may also directly enter the marginal utility of consumption when $U_{Cr} \neq 0$. This results in a second motive for redistribution—of consumption towards those agents who have the highest marginal utilities or “*consumption needs*”; that is, redistribution “*to each according to his needs*”. If $U_{Cr}/U_C \leq 0$, both redistribution motives favor lower ranks; if instead $U_{Cr}/U_C \geq 0$, consumption needs are higher for higher ranks, in which case the two redistribution motives are not aligned. Nevertheless, the single-crossing condition (1) guarantees that it is optimal to redistribute from higher to lower ranks.

The second part of Assumption 1 imposes that the inter-temporal MRS is monotonic. If it is increasing, so that (2) is positive, then higher ranks have a stronger taste for saving (relative to current consumption) than lower ranks. In other words, those who are the most inclined to work—the higher ranks—are also those who are the most inclined to save. If instead (2) 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. In addition, if second-period preferences are homogeneous,

so that $V(S; r) \equiv V(S)$ for all r , the sign of (2) boils down to that of U_{Cr} . More generally, the sign of V_{Sr} leads to a third motive for redistribution—of future consumption towards those who value it the most. For instance, if workers are heterogeneous in their discount factor, so that $V(S; r) \equiv \beta(r) V(S)$, then $V_{Sr}/V_S > 0$ whenever higher ranks are more patient than lower ranks.

The crucial point of this setup is that we are agnostic about the underlying preferences of individuals beyond Assumption 1. This is in contrast to most of the papers in the optimal taxation literature, which posit specific functional forms for the utility function—e.g., quasilinear, separable, GHH, etc. Such functional form assumptions are problematic since they carry strong implications for the optimal taxes on labor and capital. For instance, as is well known since Atkinson and Stiglitz (1976), preferences of the form $u(C) - v(Y, r) + \beta V(S)$ imply that the optimal tax rate on capital is equal to zero. More generally, as we show below, how the marginal utilities of each good vary with rank or inclination to work—that is, the values of $U_{Cr}(r)$, $U_{Yr}(r)$, $V_{Sr}(r)$ —are the key determinants of optimal tax rates. These variables are not directly observable empirically. A key contribution of our paper is to show that, rather than postulating arbitrary *a priori* restrictions on preferences to discipline these parameters, one can identify them from simple observable sufficient statistics—namely, standard elasticities and Pareto tails. That is, we “let the data speak” and inform us about the underlying structure of preferences (and, therefore, the optimal tax system) that is consistent with empirical evidence.

Social Planner’s Problem. Consumption, income, and savings are assumed to be observable, but an individual’s preference rank r is their private information. In our baseline model, we assume for simplicity that the social planner is Rawlsian and maximizes the lowest rank’s utility subject to incentive compatibility and break-even constraints.⁶ Taking the dual to the Rawlsian problem, the optimal allocation $\{C(r), Y(r), S(r)\}$ maximizes the net present value of tax revenue

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

subject to the incentive compatibility constraint:

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

⁶We generalize our analysis to arbitrary Bergson-Samuelson social welfare objectives in Section 5. Note that the formulas for optimal tax rates on top earners that we derive in Section 3 remain valid for any social welfare function.

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

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

We solve this problem using a Myersonian approach, replacing full incentive-compatibility by local incentive-compatibility. Define the indirect utility function $W(r) \equiv U(C(r), Y(r); r) + V(S(r); r)$.⁷ Then an allocation is locally incentive-compatible, if it satisfies

$$W'(r) = U_r(C(r), Y(r); r) + V_r(S(r); r). \quad (3)$$

We refer to $W'(r)$ as the *marginal information rent* of type r . The lower bound constraint can be re-stated as $W(0) \geq W_0$. The solution to this relaxed problem is obtained using optimal control techniques and is fully described in the Appendix.

2.2 Optimal Taxes

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. The following theorem, which is the first main result of this paper, provides a full characterization of the optimal taxes in our setting:

Theorem 1 (Redistributional Arbitrage). *The optimal labor wedge τ_Y satisfies*

$$1 - \tau_Y(r) = \frac{B_Y(r)}{B_C(r)} \equiv \frac{\mathbb{E} \left[\frac{U_Y(r)}{U_Y(r')} \exp \left(\int_r^{r'} \frac{U_{Yr}(r'')}{U_Y(r'')} dr'' \right) \mid r' \geq r \right]}{\mathbb{E} \left[\frac{U_C(r)}{U_C(r')} \exp \left(\int_r^{r'} \frac{U_{Cr}(r'')}{U_C(r'')} dr'' \right) \mid r' \geq r \right]}, \quad (4)$$

and the optimal savings wedge τ_S satisfies

$$1 + \tau_S(r) = \frac{B_S(r)}{B_C(r)} \equiv \frac{\mathbb{E} \left[\frac{V_S(r)}{V_S(r')} \exp \left(\int_r^{r'} \frac{V_{Sr}(r'')}{V_S(r'')} dr'' \right) \mid r' \geq r \right]}{\mathbb{E} \left[\frac{U_C(r)}{U_C(r')} \exp \left(\int_r^{r'} \frac{U_{Cr}(r'')}{U_C(r'')} dr'' \right) \mid r' \geq r \right]}. \quad (5)$$

Theorem 1 summarizes the principle of *redistributional arbitrage*. It formalizes the idea that, at the optimal allocation, the planner is indifferent between redistributing slightly less along one

⁷To ease notation, we further 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 .

margin of inequality—consumption, leisure, or wealth—and slightly more along another. Formally, 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. This interpretation stems from a simple set of perturbation arguments that we describe in Section 2.3. 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 consume less. Similarly, the ratio B_S/B_C describes the trade-off between redistributing consumption or savings. Comparing equations (4) and (5) 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).

Interpretation of the Model. One 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”, and let C and Y stand for life-time income and consumption. In this case our results would reinterpret the savings tax as a tax on bequests.

As we discuss formally in Section 5.1, letting the function V depend on rank r allows us to nest the case of heterogeneous rates of return on savings. Furthermore, we argue in Section 5.2 that our specification of second-period preferences can capture the individual’s expected utility of future consumption and earnings in a setting with stochastically evolving types r_t . Thus, our characterization of optimal wedges naturally extends to a dynamic Mirrleesian economy.

Finally, we could also interpret C as “basic necessities” and S as “luxury goods” in a static interpretation of our model. In this case the savings tax represents a relative price distortion between the two, possibly in the form of subsidies on basic necessities. More broadly, we show in Section 5.3 that our analysis can be straightforwardly extended to a framework with fully general preferences over an arbitrary number of periods and commodities, and we discuss various applications of this generalized framework.

2.3 Perturbation-Based Interpretation of Theorem 1

In this section, we formalize the interpretation of Theorem 1 as an arbitrage between various margins of redistribution—consumption, leisure, or savings. Fix a given rank $r > 0$ and consider the following perturbation: We simultaneously raise the consumption of ranks $r' \geq r$ by $\Delta C(r') > 0$ and raise their income—i.e., reduce their leisure—by $\Delta Y(r') > 0$, while preserving local incentive compatibility (3). Moreover, we design this joint perturbation such that the utility of agent r remains unchanged, thus ensuring that the incentives of agents with ranks $r' < r$ are preserved; that is, $\Delta C(r) = \frac{-U_Y(r)}{U_C(r)} \Delta Y(r)$. We show below that the first part of this perturbation—providing agents $r' \geq r$ with higher consumption—lowers the planner’s resources by $-B_C(r) \Delta C(r)$, while the second part—raising their output—increases resources by $B_Y(r) \Delta Y(r)$. At the optimum allocation, this joint perturbation must neither raise nor lower resources, so that $\frac{B_Y(r)}{B_C(r)} = \frac{\Delta C(r)}{\Delta Y(r)} = \frac{-U_Y(r)}{U_C(r)}$. Formula (4) follows immediately. The optimum savings wedge (5) is obtained analogously as a no-arbitrage condition between redistributing via consumption and savings.

Marginal Cost of Raising Consumption: Case $U_{Cr} = 0$. Consider first the resource cost of raising the *consumption* of ranks $r' \geq r$. If preferences satisfy $U_{Cr} = 0$, this perturbation preserves local incentive compatibility for all $r' > r$ if and only if it induces a uniform increase in utility above rank r . To see this formally, notice that for any r' , an increase in the consumption of rank r' by $\Delta C(r')$ does not affect the marginal information rent at r' , since $\Delta U_r(r') = U_{Cr}(r') \Delta C(r') = 0$, and hence does not require any further change in utility above r' . Now, this uniform increase in utility above rank r implies that the consumption of agents $r' > r$ must increase in proportion to their inverse marginal utility $\frac{1}{U_C(r')}$. As a result, the perturbation lowers the planner’s resources by

$$-\mathbb{E} \left[\frac{1}{U_C(r')} \mid r' \geq r \right] \Delta W(r) = -B_C(r) \Delta C(r),$$

where $\Delta W(r) = U_C(r) \Delta C(r)$ represents the increase in utility for rank r associated with the perturbation of consumption. Therefore, $B_C(r)$ represents the marginal resource cost of raising the consumption of ranks $r' > r$ in an incentive-compatible manner.

Marginal Cost of Raising Consumption: General Case. With general non-separable preferences $U_{Cr} \neq 0$, a uniform increase in utility no longer preserves local incentive compatibility. Rather, the perturbation must now raise the utility of ranks $r' > r$ in proportion to $\mu_C(r, r') \equiv \exp\left(\int_r^{r'} \frac{U_{Cr}(r'')}{U_C(r'')} dr''\right)$, and consumption in proportion to $\frac{1}{U_C(r')} \mu_C(r, r')$, thus leading to the expression of the marginal benefits B_C in equations (4) and (5). This is because the perturbation $\Delta C(\cdot)$

changes utility levels for $r' > r$ by $\Delta W(r') = U_C(r') \Delta C(r')$ and marginal information rents by $\Delta U_r(r') = U_{Cr}(r') \Delta C(r')$. It therefore preserves local incentive compatibility if and only if

$$\Delta W'(r') = \Delta U_r(r') = \frac{U_{Cr}(r')}{U_C(r')} \Delta W(r').$$

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. Intuitively, suppose that higher ranks have lower consumption needs, i.e., $U_{Cr} < 0$. We then have $\mu_C < 1$, so that the utility of higher ranks does not need to increase by as much as that of lower ranks to maintain incentive compatibility. This is because the higher level of consumption at rank r' is not that attractive for higher ranks $r'' > r'$, who don't value consumption as highly; thus, a relatively small increase in utility at r'' is sufficient to deter them from mimicking lower ranks.

Marginal Benefit of Reducing Leisure or Savings. Consider now the second part of the perturbation, whereby the planner reduces the *leisure*, or raises the income, of ranks $r' \geq r$. Following analogous steps as in the previous case, we find that if preferences satisfied $U_{Yr} = 0$, the utility of ranks $r' \geq r$ would need to fall uniformly to preserve local incentive compatibility, so that their output would need to rise in proportion to $1/(-U_Y(r'))$. The non-separability $U_{Yr} < 0$ requires an incentive-adjustment $\mu_Y(r, r') = \exp\left(\int_r^{r'} \frac{U_{Yr}(r'')}{U_Y(r'')} dr''\right)$. As a result, this perturbation frees an amount of resources equal to $B_Y(r) \Delta Y(r)$, where B_Y is defined in equation (4). Similarly, a perturbation that lowers the utility of types $r' > r$ by reducing their savings, while preserving local incentive compatibility, raises resources in proportion to $B_S(r)$, defined in equation (5).

Welfare-Improving Perturbations and Independence of Taxes. The elementary perturbations described above can also be used to identify possible directions of welfare improvement to a sub-optimal tax schedule. If one of the marginal benefits of redistribution exceeds another, then the planner gains resources by increasing redistribution along one margin and reducing it along another. This argument immediately implies that optimal taxes can be set independently of one another: The arbitrage formula (4) characterizes the optimal labor income taxes regardless of the value (optimal or not) of the savings taxes. Similarly the arbitrage formula (5) characterizes the optimal savings taxes regardless of the level of labor income taxes.

2.4 Relationship to the “ABC” Optimal Tax Formulas

Our representation of the optimal tax system contrasts with the “ABC” expressions typically derived in the literature following Diamond (1998); see, e.g., Gerritsen et al. (2020), Schulz (2021), and Ferey, Lockwood, and Taubinsky (2021). The proof of Theorem 1 shows that the optimal income and savings wedges can also be expressed as the solution to the following three equations:

$$\frac{\tau_Y(r)}{1 - \tau_Y(r)} = A(r) B_C(r), \quad \tau_Y(r) = A(r) B_Y(r), \quad \frac{\tau_Y(r)}{1 - \tau_Y(r)} (1 + \tau_S(r)) = A(r) B_S(r), \quad (6)$$

where $A \equiv \frac{U_{Cr}}{U_C} - \frac{U_{Yr}}{U_Y}$.

The first equation in (6) (“consumption-ABC”) re-states and generalizes the familiar ABC formula from Theorem 1 in Saez (2001) to the present environment.⁸ It equates the marginal efficiency cost of increasing the labor wedge at rank r , $\frac{\tau_Y}{1 - \tau_Y} \frac{1}{A \cdot U_C}$, to the additional resources the planner can raise by reducing the consumption of infra-marginal ranks $r' > r$, B_C/U_C . To see this, consider a perturbation $(\Delta C(r), \Delta Y(r))$ that keeps rank r indifferent by marginally reducing both their consumption and their output, so that $\Delta Y = (-U_C/U_Y) \Delta C$. The resource cost of this perturbation is given by $\Delta Y - \Delta C = \frac{\tau_Y}{1 - \tau_Y} \Delta C$. At the same time, the perturbation reduces the marginal information rent at rank r by $\Delta U_r = U_{Cr} \Delta C + U_{Yr} \Delta Y = A \cdot U_C \Delta C$ and thereby makes it strictly less attractive for ranks $r' > r$ to mimic rank r . This allows the planner to reduce the consumption of ranks $r' > r$, with a resource gain (per our earlier analysis) equal to $(B_C/U_C) \Delta U_r = A \cdot B_C \Delta C$.

Analogously, the second equation (“leisure-ABC”), which is novel, equates the marginal cost of the tax distortion at r ,⁹ to the marginal resource gains of reducing the leisure of agents $r' > r$, $B_Y/(-U_Y)$. The third equation (“savings-ABC”) equates the marginal cost of the tax distortion at r , to the marginal benefit of reducing the savings of agents $r' > r$, B_S/V_S . Our arbitrage representations (4) and (5) are then obtained by eliminating the marginal cost of tax distortions $A(r)$ from these ABC formulas.¹⁰

Importantly, because leisure, consumption and savings are linked through the incentive compatibility and budget constraints, the three formulas that characterize the optimal labor income taxes

⁸Note in particular that, if the utility function takes the form $u(C, Y/\theta(r))$, where $\theta(r)$ represents worker r ’s productivity and is distributed according to a distribution F , then $A = \frac{1 + \zeta_Y^M}{\zeta_Y^H} \cdot \frac{1 - F(\theta)}{\theta f(\theta)}$, where ζ_Y^M and ζ_Y^H denote respectively the Marshallian (uncompensated) and Hicksian (compensated) elasticities of labor supply.

⁹Note that the marginal cost can be expressed as: $\frac{\tau_Y}{1 - \tau_Y} \frac{1}{A \cdot U_C} = \tau_Y \frac{1}{A \cdot (-U_Y)} = \frac{\tau_Y}{1 - \tau_Y} \frac{1 + \tau_S}{A \cdot V^r}$.

¹⁰Note moreover that the ABC formulas imply $A(r) = 1/B_Y(r) - 1/B_C(r)$. Thus, our arbitrage representation provides the decomposition of this term—which drives optimal taxes—into the consumption- and the leisure-based motive for redistribution.

(consumption-ABC, leisure-ABC, and redistributive arbitrage) are all *equivalent* to each other. However, as we shall see below, they differ in terms of the observable statistics that they emphasize, and therefore the calibration of optimal income taxes. Furthermore, comparing formulas (4), (5) and (6) highlights that the principle of redistributive arbitrage, in contrast to the ABC representations, offers a unified perspective on optimal income and savings taxes. This representation also clarifies that optimal savings taxes are independent of income taxes, which has direct implications for the set of parameters and observables that determine the optimal savings wedge: It depends on the parameters that enter B_S and B_C directly, but is independent of the parameters that only affect B_Y or A .

2.5 When Should Savings Be Taxed?

Our savings wedge representation (5) nests the uniform commodity taxation theorem of Atkinson and Stiglitz (1976) as a special case. Specifically, the optimal savings wedge is equal to zero for all types—i.e., redistribution should be achieved only through income taxes—if the marginal rate of substitution between consumption and savings is homogeneous across ranks r . The following corollary also shows that the converse statement is true:

Corollary 1 (Atkinson-Stiglitz Theorem). *The optimal allocation satisfies $B_S(r) \geq B_C(r)$ and the optimal savings wedge is $\tau_S(r) \geq 0$ for all r , if and only if $\frac{V_{Sr}(r)}{V_S(r)} - \frac{U_{Cr}(r)}{U_C(r)} \geq 0$ for all r .*

In other words, the optimal savings tax inherits the sign of $V_{Sr}/V_S - U_{Cr}/U_C$. This insight is already present in Mirrlees (1976). If the intertemporal MRS is increasing (resp., decreasing) with r , so that higher ranks are more inclined to save (resp., consume) their current income, then it is optimal to tax (resp., subsidize) savings at the top of the income distribution. When $V_{Sr}/V_S = U_{Cr}/U_C$, the optimal allocation equalizes the marginal benefit of redistributing savings to the marginal benefit of redistributing consumption for all r , and there is no reason to tax savings differently than consumption.¹¹ When $V_{Sr}/V_S > U_{Cr}/U_C$, the planner can screen the more productive ranks—i.e., deter them from mimicking lower ranks—via positive savings taxes on lower ranks by exploiting the fact that their taste for savings (relative to current consumption) is stronger than that of lower ranks. Formally, consider a perturbation that increases consumption for rank r by $\Delta C(r)$, and reduces their savings by $\Delta S(r)$ so as to keep their utility unchanged—that is, $U_C(r) \Delta C(r) + V_S(r) \Delta S(r) = 0$. This perturbation changes their information rent by $U_{Cr}(r) \Delta C(r) + V_{Sr}(r) \Delta S(r)$. Thus, $U_{Cr}/U_C - V_{Sr}/V_S$ measures the change in information rents

¹¹It is straightforward to check from the definitions of the marginal benefits B_S, B_C that, when $V_{Sr}/V_S = U_{Cr}/U_C$, $B_S(r) = B_C(r)$ for all r if and only if $1/V_S(r) = 1/U_C(r)$, or $\tau_S(r) = 0$, for all r .

that comes with an increase in consumption and a reduction in savings that leave individuals of rank r indifferent. If such a perturbation reduces information rents, i.e. $U_{Cr}/U_C - V_{Sr}/V_S < 0$, then it allows the planner to increase the static redistribution from higher towards lower ranks, thus leading to a rationale for taxing savings.¹²

3 Sufficient Statistics Representation of Optimal Top Tax Rates

In this section, we express the marginal benefits of redistribution B_C , B_Y , and B_S , and hence the optimal income and savings taxes, in terms of sufficient statistics that can be observed empirically. Theorem 1 and Corollary 1 imply that the needs-based, ability-based, and savings-based complementarity variables U_{Cr}/U_C , U_{Yr}/U_Y , and V_{Sr}/V_S play a critical role. We first derive an identification result (Lemma 1) that shows that these variables can be identified from the distribution of income and consumption, along with standard behavioral elasticities. We then apply this result to obtain our sufficient-statistic expressions for optimal taxes (Theorem 2).

3.1 Identification Lemma

We start by introducing the relevant Pareto coefficients and elasticities, before deriving the identification of the three complementarity variables in terms of these parameters.

Sufficient Statistics. We denote by $s_C(r)$ the share of consumption in retained income at rank r , and $\rho_C(r)$, $\rho_Y(r)$, $\rho_S(r)$ the local Pareto coefficients of the distributions of consumption, labor income, and savings, respectively:

$$s_C(r) \equiv \frac{C(r)}{(1 - \tau_Y(r))Y(r)}$$

and

$$\frac{1}{\rho_X(r)} \equiv -\frac{d \ln X(r)}{d \ln(1-r)} = \left[\frac{d \ln(1 - F_X(X(r)))}{d \ln X(r)} \right]^{-1}$$

for any $X \in \{C, Y, S\}$, where F_X and f_X denote the c.d.f. and p.d.f. of the distribution of X . In addition, we define four elasticity variables $\zeta_C(r)$, $\zeta_Y(r)$, $\zeta_{CY}(r)$, $\zeta_S(r)$ as follows. Let

$$\zeta_C(r) \equiv -\frac{\partial \ln U_C(C, Y; r)}{\partial \ln C} \Big|_{C=C(r), Y=Y(r)} = -\frac{C(r)U_{CC}(r)}{U_C(r)}$$

¹²As we show in Section 5, the intuition and the result generalize to preferences of the form $U(C, S, Y; r)$, allowing for interaction between S and r along the same lines as C and r . Uniform commodity taxation then holds ($\tau_S = 0$ for all r) if and only if $\frac{U_{Cr}}{U_C} = \frac{U_{Sr}}{U_S}$ for all r , in which case the incentive-adjustments are the same: $\mu_C(r, r') = \mu_S(r, r')$.

and

$$\zeta_S(r) \equiv - \left. \frac{\partial \ln V(S; r)}{\partial \ln S} \right|_{S=S(r)} = - \frac{S(r) V_{SS}(r)}{V_S(r)}$$

denote the coefficients of relative risk aversion in periods 0 and 1, respectively. Let also

$$\zeta_Y(r) \equiv \left. \frac{\partial \ln (-U_Y(C, Y; r))}{\partial \ln Y} \right|_{C=C(r), Y=Y(r)} = \frac{Y(r) U_{YY}(r)}{U_Y(r)}$$

denote an inverse elasticity of labor supply; if the utility function is separable, so that $\zeta_{CY} = 0$, then ζ_Y is the inverse of the Frisch elasticity.¹³ Finally, let

$$\zeta_{CY}(r) \equiv \left. \frac{\partial \ln U_C(C, Y, r)}{\partial \ln Y} \right|_{C=C(r), Y=Y(r)} = \frac{Y(r) U_{CY}(r)}{U_C(r)}$$

denote the coefficient of complementarity between consumption and labor supply. These four elasticity parameters all have direct empirical counterparts (see Section 4.1).

Identification. We now show that the complementarity variables U_{Cr}/U_C , U_{Yr}/U_Y , V_{Sr}/V_S can be expressed in terms of the above sufficient statistics and the tax schedule. More specifically, we show that they are identified up to one degree of freedom, which we take to be $\rho_{U_C}^{-1}(r) \equiv -d \ln U_C(r) / d \ln(1-r)$, the inverse of the local Pareto tail coefficient on inverse marginal utilities of consumption. This degree of freedom stems from the fact that the solution to our optimal taxation problem is invariant to any monotone transformation of the agents' indirect utility; that is, any monotone function of $U(C, Y, r) + V(S, r)$ leaves agents' incentive and lower bound constraints unchanged but results in a shift of the value of $\rho_{U_C}^{-1}(r)$. Nevertheless, the observable parameters introduced in the previous paragraph are sufficient to fully identify the differences $U_{Yr}/U_Y - U_{Cr}/U_C$ and $V_{Sr}/V_S - U_{Cr}/U_C$, which govern how intra- and inter-temporal marginal rates of substitution vary across ranks, as well as the incentive-adjusted inverse marginal utilities $(U_X(r)/U_X(r')) \exp \int_r^{r'} (U_{Xr}/U_X) dr''$ for $X \in \{C, Y\}$ and $(V_S(r)/V_S(r')) \exp \int_r^{r'} (V_{Sr}/V_S) dr''$ that determine optimal taxes via formulas (4) and (5). Therefore, the Rawlsian optimal tax schedule and, as long as the Inada conditions hold, the optimal top tax rates under any social welfare objective, are invariant to the value judgement embedded in the choice of $\rho_{U_C}^{-1}(r)$ and depend only on the empirically observable sufficient statistics introduced in the previous paragraph.

¹³More generally, the inverse Frisch elasticity is equal to $\zeta_Y - s_C \zeta_C (\zeta_{CY}/\zeta_C)^2$. The empirical evidence suggests that $0 \leq \zeta_{CY}/\zeta_C < 0.15$ and $\lim_{r \rightarrow 1} s_C(r) = 0$ (see Section 4.1). Thus, ζ_Y^{-1} is quantitatively very close to the Frisch elasticity and converges to the latter for top income earners.

Lemma 1 (Identification). *The variables U_{Cr}/U_C , U_{Yr}/U_Y , and V_{Sr}/V_S can be expressed as:*

$$(1-r) \frac{U_{Cr}(r)}{U_C(r)} = \frac{\zeta_C(r)}{\rho_C(r)} - \frac{\zeta_{CY}(r)}{\rho_Y(r)} + \frac{1}{\rho_{U_C}(r)}, \quad (7)$$

and

$$(1-r) \frac{U_{Yr}(r)}{U_Y(r)} = -\frac{\zeta_Y(r)}{\rho_Y(r)} + \frac{s_C(r)\zeta_{CY}(r)}{\rho_C(r)} - \frac{d \ln(1-\tau_Y(r))}{d \ln(1-r)} + \frac{1}{\rho_{U_C}(r)}, \quad (8)$$

and

$$(1-r) \frac{V_{Sr}(r)}{V_S(r)} = \frac{\zeta_S(r)}{\rho_S(r)} - \frac{d \ln(1+\tau_S(r))}{d \ln(1-r)} + \frac{1}{\rho_{U_C}(r)}. \quad (9)$$

Moreover, for $X \in \{C, Y, S\}$ the incentive-adjusted inverse marginal utilities have inverse local Pareto tail coefficients equal to $\frac{\zeta_C(r)}{\rho_C(r)} - \frac{\zeta_{CY}(r)}{\rho_Y(r)}$, $-\frac{\zeta_Y(r)}{\rho_Y(r)} + \frac{s_C(r)\zeta_{CY}(r)}{\rho_C(r)}$, and $\frac{\zeta_S(r)}{\rho_S(r)}$, respectively.

Lemma 1 is a generalization of Lemma 1 in Saez (2001) to our economy. Equations (7), (8), and (9) show that empirically observable parameters—standard elasticities, Pareto coefficients, and measures of tax progressivity—together pin down the three key complementarity parameters, up to one degree of freedom captured by the Pareto coefficient $\rho_{U_C}^{-1}(r)$.¹⁴ These expressions are obtained by totally differentiating $U_C(r)$, $U_Y(r)$, and $V_S(r)$, which allows us to decompose their respective (inverse) local Pareto tail coefficients into a component dependent on U_{Cr}/U_C , U_{Yr}/U_Y , and V_{Sr}/V_S that captures the rank dependence of marginal utilities for a given allocation, and a component that captures the variation of allocations at a given rank. The latter is fully identified from preference elasticities and the local Pareto tail coefficients on allocations, which are observable.

Furthermore, these observable sufficient statistics fully identify how the marginal rates of substitution vary with rank (and hence, by Corollary 1, the sign of the optimal capital tax rate), and the incentive-adjusted inverse marginal utilities (and hence, by Theorem 1, the optimal income and capital tax schedules more generally), since the latter net the variation in marginal utilities that is due to rank-dependent preferences out of the inverse marginal utilities and only retain the part that varies with allocations. Crucially, this identification does not rely on any specific functional form assumption for preferences: The “data” implicitly inform us about the underlying correlation structure between ranks and marginal utilities that matters for optimal (Rawlsian) taxes.¹⁵

¹⁴The proof of Lemma 1 shows that $\lim_{r \rightarrow 1} \rho_{U_C}^{-1}(r) < 1$ (imposing a lower bound on the Pareto tail coefficient of inverse marginal utilities) whenever $\lim_{r \rightarrow 1} U_C(r) = 0$, and $\lim_{r \rightarrow 1} \rho_{U_C}^{-1}(r) = 0$ (implying that inverse marginal utilities are thin-tailed) whenever $\lim_{r \rightarrow 1} U_C(r) > 0$.

¹⁵By contrast, as we already discussed above, many papers in the literature impose strong *a priori* assumptions on the utility function to derive optimal taxes in terms of elasticity parameters and Pareto coefficients, before resorting to empirical estimates of these parameters to evaluate the formulas quantitatively. As emphasized by Chetty (2009), a potential pitfall of this “sufficient statistic” approach is that these empirical estimates may not be compatible with the structural restrictions imposed by the underlying model that led to the formula. For instance, suppose that

To understand the key insight of Lemma 1, focus on top earners ($r \rightarrow 1$), for whom the Pareto coefficients ρ_C, ρ_Y, ρ_S and marginal tax rates τ_Y, τ_S converge to constants. Suppose moreover that the risk-aversion parameters over consumption and savings are equal (to one, say), $\zeta_C = \zeta_S = 1$, and that the complementarity coefficient ζ_{CY} is small relative to risk aversion, as is the case empirically. Equations (7) and (9) then imply that $(1 - r) [V_{Sr}/V_S - U_{Cr}/U_C] = 1/\rho_S - 1/\rho_C$. Thus, the sign of $V_{Sr}/V_S - U_{Cr}/U_C$ is determined by the relative thickness of the Pareto tails ρ_C vs. ρ_S . Specifically, it is positive—so that capital should be taxed—if and only if $\rho_C > \rho_S$, i.e., iff consumption is strictly more evenly distributed than savings at the top. Intuitively, the relative thickness of the tails of consumption and savings (or, more generally the ratios of elasticities and Pareto coefficients ζ_C/ρ_C vs. ζ_S/ρ_S) reflect how the taste for current consumption relative to savings varies along the ability distribution. In particular, observing that $\rho_C > \rho_S$ indicates that the consumption share s_C converges to 0 as $r \rightarrow 1$; that is, top earners spend a vanishing fraction of their labor income on current consumption, which in turn implies that V_S/U_C must be increasing along the ability distribution for given C, Y and S . More generally, equations (7) to (9) show that these elasticities and Pareto coefficients determine not only the signs, but also the values of $U_{Yr}/U_Y - U_{Cr}/U_C$ and $V_{Sr}/V_S - U_{Cr}/U_C$, as well as those of the incentive-adjusted inverse marginal utilities that appear in Theorem 1. They are therefore natural and transparent sufficient statistics for optimal labor and capital taxes.

3.2 Optimal Top Tax Rates

We now express the optimal labor income and savings wedges at the top of the income distribution in terms of the sufficient statistics introduced in Section 3.1.

Assumption 2. *The optimal allocation $\{C(\cdot), Y(\cdot), S(\cdot)\}$ is co-monotonic, and the distributions of income, consumption, savings, and rates of return have unbounded support and upper Pareto tails with coefficients $\rho_Y, \rho_C, \rho_S, \rho_R$, respectively. In addition, the elasticities $\zeta_C, \zeta_S, \zeta_Y, \zeta_{CY}$ and the parameter s_C converge to finite limits as $r \rightarrow 1$.*

Lemma 1, along with Assumption 2, allows us to derive empirical counterparts for the marginal

the values of the calibrated parameters imply that the value of $V_{Sr}/V_S - U_{Cr}/U_C$ implied by equations (7) and (9) is strictly negative, as will most often be the case in our quantitative exercises of Section 4. This overidentifying restriction is inconsistent with, e.g., separable preferences with a marginal utility of consumption that is independent of r . To take an even more striking example, suppose that optimal taxes were derived under the assumption that preferences are GHH, $U = u(g(C) - v(Y/\theta(r)))$ for some concave constant-elasticity functions u and g and convex function v . While this utility function implies $U_{Cr} \leq 0$, we can show that this functional form must either violate the restrictions of Lemma 1, or impose that $\rho_C = \rho_Y$, which as we discuss below is not consistent with empirical evidence.

benefits terms B_C, B_Y, B_S that appear in the optimal tax formulas of Theorem 1. We find¹⁶

$$\lim_{r \rightarrow 1} B_C(r) = \left[1 - \frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y} \right]^{-1} \quad (10)$$

and

$$\lim_{r \rightarrow 1} B_Y(r) = \left[1 + \frac{\zeta_Y}{\rho_Y} - \frac{s_C \zeta_{CY}}{\rho_C} \right]^{-1} \quad (11)$$

and

$$\lim_{r \rightarrow 1} B_S(r) = \left[1 - \frac{\zeta_S}{\rho_S} \right]^{-1}. \quad (12)$$

Abstracting for now from complementarities, these expressions show that there is a natural mapping between consumption (resp., income, savings) data and the marginal benefits of redistributing consumption (leisure, wealth). The marginal benefits of redistributing consumption B_C (resp., savings B_S) are increasing in the level of consumption (resp., savings) inequality, as measured by the respective inverse Pareto coefficients $1/\rho_C$ and $1/\rho_S$. The marginal benefits of redistributing leisure, B_Y , are increasing in the level of leisure inequality, or decreasing in the level of income inequality $1/\rho_Y$; intuitively, high income inequality indicates that top earners are hard-working and have relatively little leisure. Finally, the complementarity between consumption and income ζ_{CY} lowers (resp., raises) the marginal benefits of redistributing consumption (resp., leisure).

Expressions (10), (11) and (12) immediately lead to the following theorem, which is the second main result of this paper:

Theorem 2 (Sufficient-Statistic Representation). *Suppose that the optimal allocation satisfies Assumption 2. Then the optimal labor wedge on top income earners $\bar{\tau}_Y \equiv \lim_{r \rightarrow 1} \tau_Y(r)$ satisfies*

$$1 - \bar{\tau}_Y = \frac{1 - \zeta_C/\rho_C + \zeta_{CY}/\rho_Y}{1 + \zeta_Y/\rho_Y - s_C \zeta_{CY}/\rho_C} \quad (13)$$

and the optimal savings wedge on top income earners $\bar{\tau}_S \equiv \lim_{r \rightarrow 1} \tau_S(r)$ satisfies

$$1 + \bar{\tau}_S = \frac{1 - \zeta_C/\rho_C + \zeta_{CY}/\rho_Y}{1 - \zeta_S/\rho_S}, \quad (14)$$

where $\frac{\zeta_C}{\rho_C} < 1 + \frac{\zeta_{CY}}{\rho_Y}$ and $\frac{\zeta_S}{\rho_S} < 1$.

¹⁶As long as leisure is a normal good, B_Y is finite and bounded above by 1. On the other hand, the representation of B_C requires that $\frac{\zeta_C}{\rho_C} < 1 + \frac{\zeta_{CY}}{\rho_Y}$; if this condition is violated then the marginal benefits of redistributing consumption B_C are infinite, and thus the allocation cannot be optimal. Similarly, the representation of B_S requires that $\frac{\zeta_S}{\rho_S} < 1$; otherwise B_S is infinite. These restrictions are imposed jointly on the primitive preference parameters and on the Pareto tails of the income, consumption, and savings distributions. They are, in principle, testable.

Equation (13) provides a very simple generalization of the standard top income tax rate formula of Saez (2001) to a dynamic environment, and equation (14) provides an analogous sufficient statistics formula for savings taxes. *Ceteris paribus*, high income and consumption inequality 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. A higher degree of complementarity unambiguously lowers the optimal top income tax rate, and raises the optimal top savings tax rate. This is a familiar result: When preferences are non-separable, it is optimal to tax less heavily the goods that are complementary to labor (Corlett and Hague (1953)).

Importantly, the optimal income tax rate (13) depends explicitly on the Pareto tail coefficient of *consumption* in addition to that of labor income. This dependence arises naturally from the marginal benefits of redistributing consumption B_C and intuitively captures the notion that the marginal gains of further redistribution are linked to the tail of the consumption distribution, that is, to how much the tax system—as well as, potentially, all of the additional private insurance mechanisms to which individuals have access—already manages to redistribute. Thus, the central take-away is that, in dynamic economies, the optimal design of taxes should rely not only on income, but also on consumption data. Our redistributive arbitrage representation gives a transparent interpretation of this result.

By the same reasoning, in the static framework, the optimal income tax rate should also depend implicitly on both consumption and income inequality. However, in the static model, consumption is equal to after-tax income, so that the Pareto coefficients ρ_Y and ρ_C coincide—an over-identifying restriction that can be tested and is generally rejected by the data. Because of this equivalence, the existing literature systematically expresses the optimal static tax formula in terms of ρ_Y only, and uses income data to estimate it. But there is no compelling conceptual reason to do so: One could alternatively express the static optimum formula in terms of ρ_C and estimate it using consumption data. Breaking the equivalence between consumption and after-tax income by adding a consumption-savings margin to the model clarifies that both coefficients ρ_Y and ρ_C matter independently for the level of optimal labor income taxes.

3.3 A Tale of Three Tails

The budget constraint in our model imposes that income is split between consumption and savings. This in turn leads to $\rho_Y = \min\{\rho_C, \rho_S\}$, that is, consumption and savings are both at least as evenly distributed as labor income.¹⁷ In particular, this restriction implies that one cannot choose

¹⁷If $\rho_C < \rho_Y$ (resp., $\rho_S < \rho_Y$), then the consumption (resp., savings) shares of after-tax income must grow arbitrarily large, which violates that these shares are both bounded between 0 and 1. If $\min\{\rho_C, \rho_S\} > \rho_Y$, then the

all three Pareto coefficients freely from the data. This is the analogue of the condition $\rho_Y = \rho_C$ in the static setting. Our model is thus consistent with the following three scenarii:

1. $\rho_Y = \rho_C < \rho_S$, so that savings are strictly more evenly distributed than income and consumption. Equivalently, the budget share of consumption s_C converges to 1 for top earners.¹⁸
2. $\rho_Y = \rho_S < \rho_C$, so that consumption is strictly more evenly distributed than income and savings. Equivalently, the budget share of consumption s_C converges to 0 for top earners.
3. $\rho_Y = \rho_C = \rho_S$, so that income, consumption, and savings are all as evenly distributed. Equivalently, the budget share of consumption s_C takes on any value between 0 and 1.

Previewing our quantitative results, we present empirical evidence in Section 4 that $\rho_C > \rho_Y$, which in turn requires that $\rho_S = \rho_Y$ (Case 2). In the sequel, we denote by $\bar{\tau}_Y^{Saez}$ the static optimum derived by Saez (2001, equation (8)). It is expressed in terms of the Hicksian (compensated) and Marshallian (uncompensated) elasticities of labor supply ζ_Y^H, ζ_Y^M as:

$$\bar{\tau}_Y^{Saez} = \frac{1}{1 - \zeta_Y^I + \rho_Y \zeta_Y^H}, \quad (15)$$

where $\zeta_Y^I \equiv \zeta_Y^H - \zeta_Y^M$ is the income effect parameter. We derive analytically the map between ζ_Y^H, ζ_Y^I and our elasticities $\zeta_C, \zeta_Y, \zeta_S, \zeta_{CY}$ in the Appendix.

Case 1. Savings have a Thinner Tail than Income and Consumption. Suppose first that savings have a thinner tail than income and consumption, so that $\rho_Y = \rho_C < \rho_S$ and $s_C = 1$. In this case, the Hicksian and Marshallian elasticities ζ_Y^H, ζ_Y^M identify ζ_Y and ζ_C , and it is straightforward to show that formula (13) reduces to the static optimum (15).¹⁹ Thus, the static analysis of Saez (2001) delivers the correct optimal tax rate on labor income, and data on consumption (or savings) is not required to evaluate it. Intuitively, when $s_C = 1$ the dynamic model is equivalent to a static model at the top, since the savings share of income converges to zero: Top earners spend most of their income on current consumption. Unfortunately, as we argue below, this case is not the empirically relevant one.

consumption and savings shares must both converge to 0, which violates the inter-temporal budget constraint. Thus, $\min\{\rho_C, \rho_S\} = \rho_Y$.

¹⁸Differentiating the inter-temporal budget constraint with respect to r and taking limits implies $\frac{\rho_Y}{\rho_C} s_C + \frac{\rho_Y}{\rho_S} s_S = 1$ with $s_C + s_S = 1$, which pins down s_C in Cases 1 and 2.

¹⁹In Case 1, we have $\tilde{\zeta}_Y = (1 - \zeta_Y^I)/\zeta_Y^H$ and $\tilde{\zeta}_C = \zeta_Y^I/\zeta_Y^H$ where $\tilde{\zeta}_Y \equiv \zeta_Y - \zeta_{CY}$ and $\tilde{\zeta}_C \equiv \zeta_C - \zeta_{CY}$. Conversely, $\zeta_Y^H = 1/(\tilde{\zeta}_Y + \tilde{\zeta}_C)$ and $\zeta_Y^I = \tilde{\zeta}_C/(\tilde{\zeta}_Y + \tilde{\zeta}_C)$. Hence, $1 - \bar{\tau}_Y^{Saez} = \frac{1 - \tilde{\zeta}_C/\rho_Y}{1 + \tilde{\zeta}_Y/\rho_Y}$.

Case 2. Consumption has a Thinner Tail than Income and Savings. Suppose next that consumption has a thinner tail than income and savings, so that $\rho_Y = \rho_S < \rho_C$ and $s_C = 0$. In this case, ζ_Y^H, ζ_Y^M identify ζ_Y and ζ_S , and the static optimum $\bar{\tau}_Y^{Saez}$ given by equation (15) reduces to the combined wedge on income and savings:²⁰

$$1 - \bar{\tau}_Y^{Saez} = \frac{1 - \bar{\tau}_Y}{1 + \bar{\tau}_S} = \frac{1 - \zeta_S/\rho_S}{1 + \zeta_Y/\rho_Y}. \quad (16)$$

Intuitively, when $s_C = 0$ at the top, so that top earners save most of their income, the optimal allocation for top earners is determined by a static trade-off between income and savings. Equation (16) shows that the static optimum $\bar{\tau}_Y^{Saez}$ now characterizes the optimal wedge between income and savings, which is the combination of the labor and savings wedges $\bar{\tau}_Y$ and $\bar{\tau}_S$. Hence, the optimal top labor income tax rate $\bar{\tau}_Y$ no longer coincides with the static optimum $\bar{\tau}_Y^{Saez}$ given by equation (15), unless $\zeta_S/\rho_Y = \zeta_C/\rho_C - \zeta_{CY}/\rho_Y$, that is, unless the Atkinson-Stiglitz theorem applies, so that the optimum savings tax rate $\bar{\tau}_S$ is equal to zero. Furthermore, by Corollary 1, the static optimum $\bar{\tau}_Y^{Saez}$ overstates the correct optimum $\bar{\tau}_Y$ whenever the optimal savings tax rate $\bar{\tau}_S$ is strictly positive, i.e., if preferences are such that $U_{Cr} < 0$, and it underestimates the optimum top labor income tax rate if it is optimal to subsidize savings. Theorem 2 gives the optimal breakdown of the combined wedge (16) into labor income and capital taxes.

Case 3. Income, Consumption, and Savings have Identical Tails. Suppose finally that the distributions of income, consumption, and savings all have the same tail coefficient, so that $\rho_Y = \rho_C = \rho_S$ and $s_C \in (0, 1)$. In this case, the optimal top income tax rate (13) generally differs from the static optimum (15). The dynamic adjustments can only be neglected when the first-period utility is quasilinear in consumption, so that $U_{CC} = U_{CY} = 0$.²¹ However, whenever the utility of consumption is strictly concave, even if preferences are GHH, the response of savings to labor income taxes modifies the optimal top income tax rate, and the standard formula of Saez (2001) ceases to apply.

3.4 Alternative Representations: 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) emphasizes different sufficient statistics, namely the cross-sectional

²⁰In Case 2, we have $\zeta_Y = (1 - \zeta_Y^I)/\zeta_Y^H$ and $\zeta_S = \zeta_Y^I/\zeta_Y^H$, or conversely, $\zeta_Y^H = 1/(\zeta_S + \zeta_Y)$ and $\zeta_Y^I = \zeta_S/(\zeta_S + \zeta_Y)$.

²¹Indeed, we then have $\zeta_C = \zeta_{CY} = \zeta_Y^I = 0$ and $\zeta_Y = 1/\zeta_Y^H$, so that the optimal labor income tax rate is equal to $1/(1 + \rho_Y/\zeta_Y)$ both in the static and the dynamic settings.

variation of savings with income net of the causal effect of income on savings (“ s'_{het} ”), to estimate optimal savings taxes. Intuitively, 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 the Atkinson-Stiglitz result. FLT’s representation of optimal savings taxes is an ABC formula scaled by the variable s'_{het} .

In the Appendix we derive the precise relationship between our optimal tax formulas and this alternative representation. We argue that both representations are equivalent *provided that* $s_C(r) > 0$, i.e., consumption takes up a non-negligible fraction of after-tax income. In particular, the sufficient statistic highlighted in FLT offers an additional moment condition to infer the ratio of risk-aversion parameters ζ_S/ζ_C , along with the Hicksian and Marshallian labor supply elasticities. However, if—as we argue is empirically plausible—consumption has a strictly thinner tail than savings, then $\lim_{r \rightarrow 1} s_C(r) = 0$, and the identification of FLT breaks down for top earners; that is, their additional sufficient statistics lose their informational content.

Intuitively, $\lim_{r \rightarrow 1} s_C(r) = 0$ implies that all the cross-sectional variation in savings is driven by labor income, while the impact of cross-sectional variation in preferences vanishes, so that $s'_{het} = 0$. Nevertheless, this does *not* imply that the optimal savings tax goes to zero. Indeed, FLT’s optimum formula scales the cross-sectional variation in preferences s'_{het} by a compensated elasticity of savings to savings taxes (holding income constant). This compensated elasticity also vanishes in the top as $\lim_{r \rightarrow 1} s_C(r) = 0$, since the substitution effect from an increase in the savings tax becomes negligible relative to the income effect—savings become inelastic. Yet the ratio between s'_{het} and the compensated elasticity of savings to savings taxes converges to a finite limit, that we show can be represented in terms of the Pareto tail coefficients of income, consumption and savings, as well as preference elasticities.

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 $\rho_C = \rho_Y$. By contrast, our result based on the Pareto tails of consumption and savings offers an alternative that identifies top income taxes even in the empirically relevant case where $\rho_C > \rho_Y$. This discussion shows that both papers are complementary, in the sense that we are able to offer prescriptions for top income and savings taxes, on which their sufficient statistics are unable to shed light.

4 Quantitative Implications

In this section, we calibrate our model in Case 2, which is likely to be the relevant case empirically. For completeness, we propose an alternative calibration for Case 3 in the Appendix.

4.1 Calibration

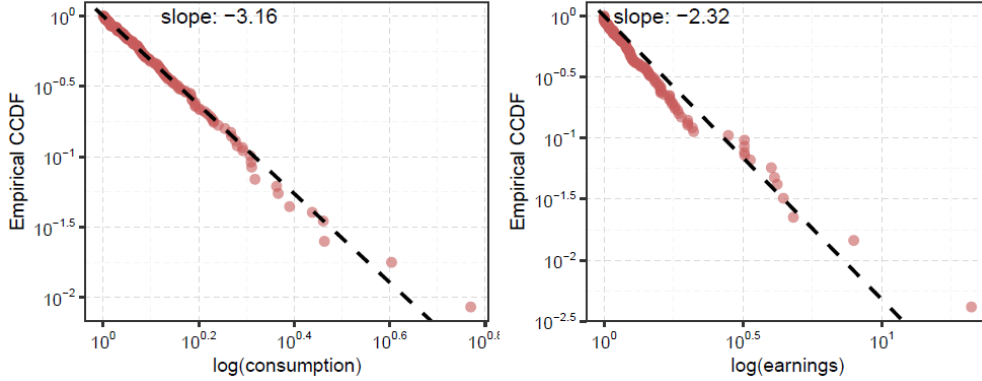
Pareto Tails of Income and Savings: ρ_Y, ρ_S . The fact that the income distribution has a Pareto tail is well documented. In the U.S., the Pareto coefficient on income is approximately equal to 1.5 (Diamond and Saez (2011)). Moreover, our model imposes $\rho_Y = \rho_S$ in Case 2. As we discuss below, our model allows for heterogeneous rates of return—which imply a strictly thicker tail for wealth than for income or savings—but our sufficient statistics allow us to remain agnostic about the source and extent of such return heterogeneity.

Before we proceed, note that this calibration follows the Mirrleesian literature by using annual income data to evaluate the Pareto coefficient. However, the relevant parameter in our model should rather be a measure of lifetime—or working-life—income inequality. The permanent income hypothesis suggests that the corresponding tail could be much thinner than that estimated from annual data. In fact, Karahan, Ozkan, and Song (2022) estimate a Pareto coefficient for lifetime earnings equal to $\rho_Y = 2.13$. (As we show next, this value is still far smaller than all of the measures of the Pareto coefficient of consumption that we could find.)²²

We do not use this lifetime value for ρ_Y , however, for two reasons. First, we want our calibration to follow as closely as possible those of the literature, which typically uses annual data to calibrate for this parameter. Second, and most importantly, the calibration should also ideally use lifetime, rather than annual, measures of consumption inequality, as well as estimates of the income and substitution effects on lifetime labor supply. Since the literature does not provide reliable estimates of these parameters, we chose for transparency to be consistent and use annual data for all of the relevant variables of our analysis.

²²The 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 Introduction 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 Coefficients of Consumption and Total Income



Pareto Tail of Consumption: ρ_C . Turning to measures of consumption inequality,²³ Toda and Walsh (2015) argue using CEX data that consumption is also Pareto distributed at the top, and they estimate an upper tail coefficient of $\rho_C = 3.38$, so that $\rho_Y/\rho_C = 0.44$. Straub (2019) finds that the income elasticity of consumption is equal to 0.7, which pins down the ratio of Pareto coefficients of income and consumption, $\rho_Y/\rho_C = \frac{C'/C}{Y'/Y} = 0.7$ or $\rho_C = 2.14$. These estimates suggest that consumption has a substantially thinner tail than income, so that $s_C \rightarrow 0$ as $r \rightarrow 1$: That is, top earners save most of their income.

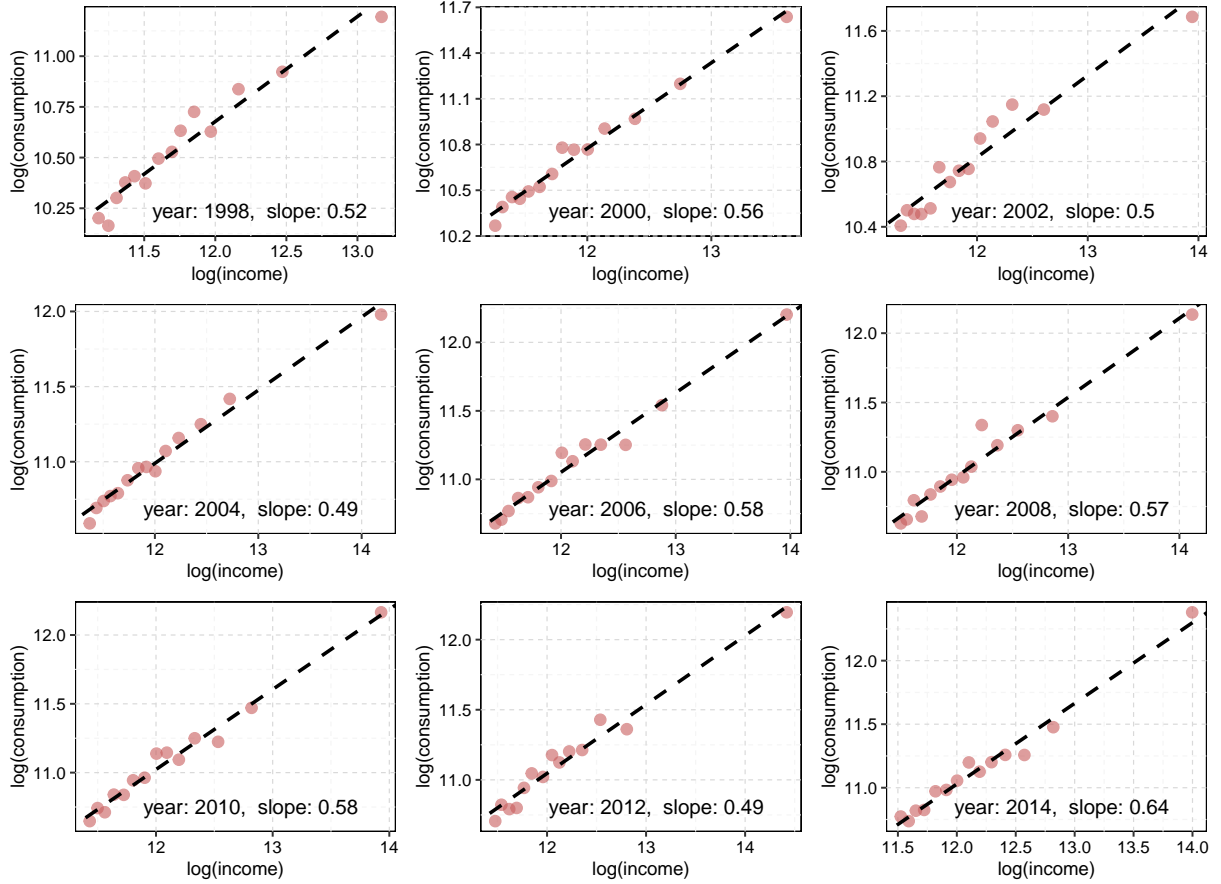
We can also impute the ratio of Pareto coefficients ρ_Y/ρ_C based on our own computations of the consumption and income shares of top earners, using the data from Blundell, Pistaferri, and Saporta-Eksten (2016) which are based on the PSID from 1998 to 2014. Since the PSID is top-coded, these estimates should be taken as suggestive. However, they allow us to represent graphically the tails of the income and consumption distributions. Figure 1 plots the log of the survival c.d.f. $1 - F(X)$ against $\log X$, where the variable X represents either consumption (left panel) or total income (right panel) % in 2014; similar figures for every year between 1998 and 2014.²⁴ We use a threshold of the top 90% for consumption, and the top 95% for income. If X was exactly Pareto distributed at the tail with coefficient ρ_X , the resulting graph would be a straight line with slope $-\rho_X$. The figure highlights that assuming a Pareto tail for consumption with a significantly larger coefficient than for income is indeed reasonable. In particular, the ratio of the slope coefficients is $2.32/3.16 = 0.73$, very close to the value found by Straub (2019).

However, notice that the x-axes of these two graphs are different. Our theoretical model further imposes that there is perfect co-monotonicity between income and consumption, which of course is

²³Note that consumption inequality should be less affected by the concern that annual measures may differ significantly from lifetime measures.

²⁴We are grateful to Alexandre Gaillard for computing these statistics for us.

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



not the case in the data. To be consistent with our theoretical analysis, in Figure 2 we plot the mean log-consumption of workers within each income quantile: by averaging, we remove consumption variation conditional on income rank. Each graph represents one year between 1998 and 2014. We use the quantiles between 0.80 to 0.94 in increments of 0.02, every percentile between 0.95 and 0.99, and 0.995.) Since income is Pareto distributed, the fact that the data points align along a straight line confirms that consumption is also Pareto-distributed at the top. Moreover, the slope of the relationship gives estimates of the ratio of Pareto coefficients ρ_Y/ρ_C between 0.49 and 0.64, which are intermediate between the values obtained by Toda and Walsh (2015) and Straub (2019).

Furthermore, the fact that consumption has a significantly thinner Pareto tail than that of income can be verified in other countries that have much better consumption data. In particular, Buda et al. (2022) use a large representative panel of consumption expenditures in Spain that contains transaction-level data from all the retail accounts of one of the World’s largest banks, BBVA—amounting to 3 billion individual transactions by 1.8 million bank customers. They construct distributional national accounts that capture 100% of aggregate consumption, allowing them

to compute consumption at each quantile of the distribution. They show that consumption inequality is substantially smaller than its income counterpart: for instance, 22.4% (resp., 4.1%, 0.8%) of 2017 aggregate consumption accrued to the top 10% (resp, 1%, 0.1%) consumption-richest adults, while the World Inequality Database shows that 31% (resp., 11%, 4.2%) of total national post-tax income accrues to the top 10% (resp, 1%, 0.1%) income earners. Moreover, they find that the power law parameterization of the tail of the consumption distribution provides a statistically significant better fit when compared to lognormal or exponential alternatives. They estimate a power-law shape parameter of $\rho_C = 3.91$ at the tail, slightly larger than the estimate of Toda and Walsh (2015) for the U.S. By contrast, the Pareto coefficient for income in Spain is approximately equal to $\rho_Y = 2$ (Blanchet et al, 2018). Thus, the ratio ρ_Y/ρ_C is equal to 0.51, a value that is close to our own estimate for the U.S.

As a result, we evaluate our optimal tax formulas below for $\rho_Y = \rho_S = 1.5$ and $\rho_Y/\rho_C \in \{0.45, 0.6, 0.75\}$.

Labor Supply Elasticities: ζ_Y, ζ_S . Recall that in Case 2, there is a one-to-one map between the Hicksian and Marshallian elasticities of labor supply ζ_Y^H, ζ_Y^M , on the one hand, and the elasticity parameters ζ_Y, ζ_S , on the other hand. There is a vast literature that estimates the elasticities of labor 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 $\zeta_Y^H = 1/3$. For top income earners, Gruber and Saez (2002) estimate a value of $\zeta_Y^H = 1/2$. Empirical evidence about the size of the income effects $\zeta_Y^I = \zeta_Y^H - \zeta_Y^M$ is mixed; see, e.g., Keane (2011). Gruber and Saez (2002) find small income effects, while Golosov, Graber, Mogstad, and Novgorodsky (2021) estimate that \$1 of additional unearned income reduces the pre-tax income in the highest income quartile by 67 cents, which for a top marginal tax rate of 50 percent translates into an income effect of 1/3. For our baseline calibration, we choose $\zeta_Y^H = 1/3$ for the Hicksian elasticity and the intermediate value $\zeta_Y^I = 1/4$ for the income effect. These values imply $\zeta_Y^{-1} = \zeta_Y^H/(1-\zeta_Y^I) = 4/9$ and $\zeta_S = \zeta_Y^I/\zeta_Y^H = 0.75$, reasonable values for the Frisch elasticity and the relative risk-aversion of top earners. We then evaluate the robustness of our quantitative results to the alternative parameter values $\zeta_Y^H = 1/2$ (so that $\zeta_S = 0.5$ and $\zeta_Y^{-1} = 2/3$) and $\zeta_Y^I = 1/3$ (so that $\zeta_S = 1$ and $\zeta_Y^{-1} = 0.5$).

Risk-Aversion and Complementarity: ζ_C, ζ_{CY} . Because the combined wedge on income and savings is equal to the static wedge (equation (16)), the values of the labor supply elasticity ζ_Y^H and the income effect parameter ζ_Y^I are sufficient to evaluate the ratio $\frac{B_Y}{B_S} = \frac{1-\bar{\tau}_Y}{1+\bar{\tau}_S}$. Information about consumption, i.e., the remaining two elasticities ζ_C and ζ_{CY} , are only required to quantify

the breakdown of the combined wedge into income and savings taxes. In our baseline calibration, we choose a first-period risk-aversion coefficient for top earners of $\zeta_C = \zeta_S = 0.75$, and we evaluate the robustness of our results to the alternative value $\zeta_C = 1.25$. To calibrate the complementarity between consumption and labor ζ_{CY} , we follow Chetty (2006) who shows that this parameter can be bounded as a function of the coefficient of relative risk aversion by $\zeta_{CY} \leq \frac{\Delta \ln C}{\Delta \ln Y} \cdot \zeta_C$, where $\frac{\Delta \ln C}{\Delta \ln Y}$ is the change in consumption that results from an exogenous variation in labor supply (e.g., due to job loss or disability). He then estimates the latter parameter in the data and finds an upper bound $\frac{\Delta \ln C}{\Delta \ln Y} < 0.15$. We use $\zeta_{CY} = 0$ as our baseline value (separable utility function) and evaluate the robustness of our results to the upper bound $\frac{\zeta_{CY}}{\zeta_C} = 0.15$.

4.2 Quantitative Results

Table 1 below summarizes our quantitative results for the optimal top tax rates on labor income and savings. The first row reports the results for our baseline calibration $(\rho_Y, \zeta_Y^H, \zeta_Y^L, \zeta_C, \zeta_{CY}) = (\frac{3}{2}, \frac{1}{3}, \frac{1}{4}, \frac{3}{4}, 0)$ and three values of the Pareto coefficient on consumption $\rho_C \in \{0.45, 0.6, 0.75\}$. We also report the static optimum $\bar{\tau}_Y^{Saez} = 1 - \frac{1 - \zeta_S / \rho_S}{1 + \zeta_Y / \rho_Y}$. The remaining rows of the table vary one parameter at a time. Note that while $\bar{\tau}_Y$ represents a marginal labor income tax on gross income, $\bar{\tau}_S$ 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 $\frac{\bar{\tau}_S}{1 + \bar{\tau}_S}$, which is the variable we report in the table. 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 $\frac{\bar{\tau}_S}{1 + \bar{\tau}_S} = 40\%$, say, corresponds to a 1.8% (resp., 1.7%) annual tax on accumulated wealth, or a 35% (resp., 58%) capital income tax. Alternatively, if we interpret our model as one of retirement saving, a wedge of 40% means that top income earners can only expect to receive a present value of 0.71 dollars of additional pension payments for each additional dollar in social security contributions.

Note that we do not restrict the utility function *a priori*: Our calibration of the elasticities and Pareto tails implicitly determines the underlying structure of preferences (see Lemma 1). Some parameter values can only be generated by $U_{Cr} < 0$, so that savings should be taxed, while others are only consistent with $U_{Cr} > 0$, so that savings should be subsidized. Specifically, the breakdown of the combined wedge $\bar{\tau}_Y^{Saez}$ between savings and income taxes $\bar{\tau}_Y, \bar{\tau}_S$ is pinned down by the ratios of risk-aversion parameters and Pareto coefficients $\zeta_C / \rho_C, \zeta_Y / \rho_Y, \zeta_S / \rho_S$ that respectively drive the marginal benefits of redistributing consumption, leisure, and savings.

Table 1: Optimal Taxes in Case 2

	$\rho_Y/\rho_C = 0.45$		$\rho_Y/\rho_C = 0.6$		$\rho_Y/\rho_C = 0.75$		$\bar{\tau}_Y^{Saez}$
	$\bar{\tau}_Y$	$\frac{\bar{\tau}_S}{1+\bar{\tau}_S}$	$\bar{\tau}_Y$	$\frac{\bar{\tau}_S}{1+\bar{\tau}_S}$	$\bar{\tau}_Y$	$\frac{\bar{\tau}_S}{1+\bar{\tau}_S}$	
<i>Baseline</i>	69%	35%	72%	29%	75%	20%	80%
$\zeta_Y^H = 0.5$	61%	14%	65%	5%	69%	-7%	67%
$\zeta_Y^I = 1/3$	67%	57%	70%	52%	73%	47%	86%
$\zeta_C = 1.25$	75%	20%	80%	0%	85%	-33%	80%
$\zeta_{CY}/\zeta_C = 0.15$	66%	41%	69%	35%	72%	29%	80%

For low values of the first-period risk aversion or a very thin consumption tail ($\rho_Y/\rho_C = 0.45$), B_C is relatively low, so that the savings tax is high and the labor income tax rate is substantially lower than in the static framework. If the consumption and savings elasticities are the same, then the fact that consumption appears to have a thinner tail than savings, or that top income earners save most of their income, suggests that the marginal benefits of redistribution are higher for savings than for consumption ($B_S > B_C$), and thus that it is optimal to load tax distortions into savings rather than consumption, resulting in a lower income and a higher savings tax. Which of these marginal benefits dominates is then a matter of the elasticity estimates on consumption vs. savings, along with the tail coefficients of the consumption and savings distributions.

For higher values of the first-period risk aversion or more unequal distributions of consumption, the savings tax is lower and the labor income tax closer to the static optimum. The marginal gains of redistributing consumption eventually exceed those of redistributing savings ($B_C > B_S$), in which case the optimum income tax $\bar{\tau}_Y$ exceeds $\bar{\tau}_Y^{Saez}$ and savings are subsidized, $\bar{\tau}_S < 0$. Analogously, higher values of the second-period risk-aversion ζ_S , driven either by a higher income effect parameter ζ_Y^I or a lower Hicksian elasticity ζ_Y^H , reduce (resp., raise) the optimal labor (savings) tax. With $\zeta_{CY} = 0$, our model also provides a lower bound on optimal income taxes and an upper bound on savings wedges that depends only on the Pareto coefficients ρ_Y and ρ_S . Since $B_C \geq 1$, we have $\bar{\tau}_Y \geq 1 - B_Y = \frac{1}{1+\rho_Y/\zeta_Y} = 60\%$ and $\bar{\tau}_S \leq B_S - 1 = \frac{1}{\rho_S/\zeta_S - 1}$ so $\frac{\bar{\tau}_S}{1+\bar{\tau}_S} \leq 52\%$ in our baseline calibration.

Next, the complementarity between consumption and labor income $\zeta_{CY} > 0$ leaves the combined labor and savings wedge unchanged but shifts the wedge from labor to savings taxes. As we discussed above, when income and first-period consumption are complements, the Corlett-Hague rule implies that the planner should reduce the tax rate on labor income and raise the tax rate on savings. Quantitatively, the complementarity correction has a significant impact on the optimal tax rates for reasonable empirical values of ζ_{CY} . Formulas (13) and (14) imply that the correction for complementarity ζ_{CY}/ρ_Y is equivalent to adjusting the Pareto tail coefficient on consumption

upwards to $\tilde{\rho}_C$ defined by $\rho_Y/\tilde{\rho}_C = \rho_Y/\rho_C - \zeta_{CY}/\zeta_C$. It thus amounts to increasing the effective gap between income and consumption inequality. In our baseline calibration, the adjustment reduces the ratio of tail coefficients from $\rho_Y/\rho_C = 0.45$ to $\rho_Y/\tilde{\rho}_C = 0.30$. For $\zeta_C = 0.75$, this lowers the marginal benefit of redistributing consumption B_C from 1.25 to 1.14, equivalent to a 9.6% increase in after-tax labor income and a corresponding increase in the savings wedge.

Savings should be taxed if and only if $\zeta_S/\zeta_C > \rho_S/\tilde{\rho}_C$ where $\tilde{\rho}_C$ is the adjusted Pareto tail coefficient. Without the complementarity correction, the values $\zeta_S = 0.75$ and $\rho_S/\rho_C = 0.45$ (resp., 0.75) imply that savings should be taxed unless the first-period risk-aversion coefficient for top earners ζ_C is larger than $\frac{\rho_C}{\rho_S}\zeta_S = 1.67$ (resp., 1). With the complementarity correction, we have $\rho_S/\tilde{\rho}_C = 0.3$ (resp., 0.6), so risk aversion ζ_C would need to exceed 2.5 (resp., 1.25) to overturn the conclusion that savings should be taxed. To sum up, already without complementarity the marginal benefit of redistributing savings appear to be high relative to the marginal benefit of redistributing consumption, as consumption has a much thinner upper tail than income and savings. The complementarity between consumption and effort only reinforces this conclusion. So unless ζ_C is very large, the marginal benefits of redistributing consumption remain substantially smaller than the marginal benefits of redistributing savings, resulting in a significant shift from income to savings taxes at the optimal allocation.

5 Extensions and Applications

In this last section, we first show that our baseline setting encompasses two important rationales for taxing the capital of top earners: rate-of-return heterogeneity and the inverse Euler equation. Next, we extend our analysis of redistributive arbitrage and the sufficient-statistic representations of optimal taxes to an environment with general preferences over an arbitrary number of periods and set of commodities, and study an application of this general framework to age-dependent taxation over the life-cycle.

5.1 Return Heterogeneity

Recent empirical evidence suggests that heterogeneous rates of return, whereby wealthier agents earn higher 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) and Fagereng, Guiso, Malacrino, and Pistaferri (2020). 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). While several recent papers derived ABC representations of optimal taxes in settings with return heterogeneity (Gerritsen, Jacobs, Rusu, and Spiritus, 2020; Schulz, 2021), the same caveats as those of Section 3.4 apply to these contributions.

In this section, we show that the generic utility function $V(S; r)$ introduced in our baseline framework of Section 2 nests the case of heterogeneous returns, thus allowing us to immediately apply our analysis to this case. To see this, interpret $V(\cdot; r)$ as an indirect utility function over *initial* savings, rather than over second-period consumption. Thus, the function V incorporates the return on savings, which are allowed to be type-dependent via the argument r . 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 , and $R(S, r) S \equiv C_2(r)$ denotes the second-period consumption. Note that this expression also allows for heterogeneity in discount rates. 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 3.3, 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.²⁵

One important advantage of the calibration in Section 4 of top income and savings taxes is that 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 Hicksian and Marshallian elasticities ζ_Y^H, ζ_Y^M , 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. Instead, return heterogeneity enters 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 ζ_{C_2} is the second-period risk aversion and $\eta \equiv SR_S(S, r)/R(S, r)$ is the scale-dependence parameter. Hence, scale dependence of returns affects the savings elasticity ζ_S through the parameter η whenever $\zeta_{C_2} \neq 1$. Specifically, increasing returns to savings ($\eta > 0$) lower ζ_S and thus optimal savings taxes when $\zeta_{C_2} < 1$, and increase

²⁵In the Appendix, we plot the tail distributions of the rates of return calculated by Gaillard and Wangner (2021) and the tail distribution of wealth. Unfortunately, the relationship between log-returns and log-income is very noisy and unstable; some of the graphs suggest an estimate of $\rho_S/\rho_R = 0.05$, which combined with our calibrated value $\rho_S = 1.5$ implies a Pareto coefficient for wealth in our framework equal to $\rho_{C_2} = \rho_S/[1 + \rho_S/\rho_R] = 1.43$, close to that observed in the data (1.4).

ζ_S and optimal savings taxes when $\zeta_{C_2} > 1$. The opposite result holds if savings have decreasing returns ($\eta < 0$). Finally, note that in our framework, type-dependence of returns does not affect optimal taxes: intuitively, this is because it does not generate any behavioral responses.

5.2 Inverse Euler Equation

Our second interpretation of V shows how our analysis can be linked to the “Inverse Euler Equation” emerging in dynamic Mirrleesian economies (Goloso, Kocherlakota, and Tsyvinski (2003), Farhi and Werning (2013), and Goloso, Troshkin, and Tsyvinski (2016)) in which types evolve stochastically over time. In such economies, an alternative motive for savings taxes arises from the need to preserve incentives over the entire working life, as savings or wealth have adverse effects on incentives. However, much of the dynamic Mirrleesian 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 (or commodity taxation more broadly) in our setting.

Specifically, 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 $R > 0$ entering the second period. 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

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

and incentive-compatibility constraints

$$v(C_2(r_2), Y_2(r_2), r_2) \geq v(C_2(r'_2), Y_2(r'_2), r_2)$$

for all $r_2, r'_2 \in [0, 1]$.

We can then characterize the optimal labor distortion in period t by equalizing the marginal benefits of redistributing second-period consumption (B_{C_2}) and second-period income (B_{Y_2}), with a similar characterization of top labor income taxes as in Sections 2 and 3.²⁶ In addition, the

²⁶The only difference is that here we are working with a utilitarian welfare criterion, rather than a Rawlsian one, but as we will show in the next subsection, this distinction does not affect the characterization of top income taxes.

resulting solution implies that

$$V_S(S) = \beta R \left\{ \mathbb{E} \left[\frac{1}{v_{C_2}(r_2)} \frac{\mu_{C_2}(0, r_2)}{\mathbb{E}[\mu_{C_2}(0, r_2)]} \right] \right\}^{-1}, \quad (17)$$

where $\mu_{C_2}(0, r_2) \equiv \exp\left(\int_0^{r_2} \frac{v_{C_2r}(r')}{v_{C_2}(r')} dr'\right)$. In other words, adjusting for discounting β and returns R , the inverse marginal utility of savings $1/V_S(S)$ is equal to an expected inverse marginal utility of second-period consumption, weighted by an adjustment factor $m(r_2)$ that is analogous to the first-period incentive adjustments described in Section 2.

This adjustment factor follows from a simple perturbation argument along the same lines as in Section 2. Suppose first that second-period preferences are separable, or $v_{C_2r}/v_{C_2} = 0$. Then, in order to preserve incentive compatibility in the second period, returns to savings must be distributed so as to raise consumption utility uniformly for all ranks r_2 , or returns must be proportional to $1/v_{C_2}(r_2)$. In this case, $\mathbb{E}[1/v_{C_2}(r_2)]$ represents the marginal resource cost of increasing agents' expected utility while preserving incentive compatibility, and $\beta R\{\mathbb{E}[1/v_{C_2}(r_2)]\}^{-1}$ is the agent's marginal utility of extra savings at the end of the first period. When preferences are non-separable ($v_{C_2r}/v_{C_2} \leq 0$), the same arguments as in Section 2 then imply that returns to savings must raise the utility of different ranks in proportion to $\mu_{C_2}(0, r_2)$ in order to preserve incentive compatibility. Thus, the expectation in the right-hand side of equation (17) represents the marginal resource cost of increasing agent's expected utility, so the above expression for $V_S(S)$ represents, again, the agent's marginal value of extra savings at the end of the first period.

Combining this expression for $V_S(S)$ with our characterization of the first-period savings wedge then yields the following generalization of the Inverse Euler Equation:

$$(1 + \tau_S(r)) U_C(r) = V_S(r) = \beta R \left\{ \mathbb{E} \left[\frac{1}{v_{C_2}(r_2)} \frac{\mu_{C_2}(0, r_2)}{\mathbb{E}[\mu_{C_2}(0, r_2)]} \right] \right\}^{-1},$$

where $1 + \tau_S(r) = \frac{B_S(r)}{B_C(r)}$ was characterized in Theorem 2. 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: First, it incorporates the optimal savings wedge $\tau_S(r)$ that accounts for heterogeneity in inter-temporal marginal rates of substitution and the extent to which savings reduce first-period information rents. Second, extending the logic of the inverse Euler equation to non-separable preferences, it accounts for the adverse effect of savings on future incentives by characterizing the marginal value of savings as a harmonic expectation of second-period marginal utilities. Furthermore, with non-separable preferences these marginal utilities are

further reweighted to account for the additional incentive adjustment required to preserve incentive compatibility in the second period.

The present discussion was kept deliberately simple by assuming that ranks were i.i.d. across time and across agents. This assumption implies that private information is short-lived, and the indirect utility of savings $V(S)$ depends on the first-period rank only through the choice of savings $S(r)$. Hellwig (2021) analyzes a dynamic Mirrleesian economy with arbitrary Markovian shock processes that integrates motives for savings taxes due to preference heterogeneity and complementarities—as in the present analysis—with wealth effects on incentives from the dynamic Mirrleesian setting. The analysis applies the above characterization of redistributive consumption and income perturbations to both intra- and inter-temporal tradeoffs to generalize both the Inverse Euler Equation and the static sufficient statistics formulas for income and savings taxes on top earners in Theorem 2. The key observation for the latter result is that the top income and savings taxes remain based on a Rawlsian logic of maximum revenue extraction, even if at other points of the distribution there are strong motives for linking labor and savings taxes intertemporally based on tax-smoothing motives. One key difference in the dynamic Mirrleesian economy is that the sufficient statistics required to compute optimal taxes are now based on the distributions of income, consumption and savings *conditional on* the entire prior sequence of types, or equivalently the entire earnings history, since the latter determines the within-period trade-off between incentives and redistribution that describes the optimal tax system. Just as age-dependence will alter the level of Pareto coefficients in Section 5.4 below, 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.

5.3 General Preferences and Multiple Commodities

In our baseline model of Section 2, we assumed that preferences were additively separable, so that the benefits of “savings” were independent of “consumption” and “income”. As we discuss formally below, it is straightforward to extend Theorem 1 to general preferences of the form $U(C, S, Y; r)$. Moreover, our analysis can be directly extended to an arbitrary set of consumption goods, leading to a characterization of optimal relative price distortions as arbitraging between redistribution through one commodity vs. another.

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 2: 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. With unrestricted preferences, the analysis will require estimates for two additional preference elasticities to account for the complementarity of consumption and income with savings.

Formally, suppose that agents' preferences are defined as $U(\mathbf{X}; r)$, where \mathbf{X} is an N -dimensional commodity vector and $r \in [0, 1]$. Let $\frac{\partial U}{\partial x_n} = U_n$ and $\frac{\partial U_r}{\partial x_n} = U_{nr}$ and assume that $\frac{U_{nr}}{U_n}$ is increasing in n . Hence, $\frac{U_m}{U_n}$ is increasing in r whenever $m > n$. The planner's cost of providing an aggregate commodity vector \mathbf{X} is $C(\mathbf{X})$, and we let $p_n = \frac{\partial C}{\partial x_n}$ denote the "price" of good n . The planner's problem reads

$$\max_{\mathbf{X}(\cdot)} \int_0^1 \omega(r) G(U(\mathbf{X}(r); r)) dr - C\left(\int_0^1 \mathbf{X}(r) dr\right)$$

subject to the agents' incentive compatibility constraints. In this formulation, $\omega(\cdot)$ represents rank-dependent Pareto weights, and the concave function $G(\cdot)$ represents the planner's aversion to inequality.

Let $\hat{\omega}(r) \equiv \omega(r) G'(U(r))$ represent the marginal welfare weight on rank r and $\mu_k(r, r') \equiv \exp\left(\int_r^{r'} \frac{U_{kr}}{U_k} dr''\right)$ denote the incentive-adjustment specific to commodity k . The optimal wedge between any pair of goods then takes the form

$$\frac{U_m(r) p_n}{U_n(r) p_m} \equiv 1 - \tau_{m,n}(r) = \frac{B_m(r)}{B_n(r)}, \quad (18)$$

where, for any $k \in \{n, m\}$,

$$B_k(r) = \mathbb{E} \left[\frac{U_k(r)}{U_k(r')} \mu_k(r, r') \mid r' \geq r \right] \left(1 - \frac{\mathbb{E} [\hat{\omega}(r') \mu_k(r, r') \mid r' \geq r]}{p_k \mathbb{E} [(U_k(r'))^{-1} \mu_k(r, r') \mid r' \geq r]} \right) \quad (19)$$

represents the marginal benefits of reducing the consumption of commodity k for ranks above r while preserving incentive-compatibility for $r' \geq r$. This representation multiplies the Rawlsian marginal benefit of redistribution $\mathbb{E} \left[\frac{U_k(r)}{U_k(r')} \mu_k(r, r') \mid r' \geq r \right]$ by an adjustment that factors in the effective Pareto weight on types $r' \geq r$. Note that the Inada conditions ensure that this adjustment factor converges to 1 at the top of the type distribution: If $\lim_{r \rightarrow 1} \hat{\omega}(r) U_k(r) = 0$, we recover the Rawlsian representation of $B_k(r)$ of Theorem 1.

In the proof of Corollary 1, 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 , or equivalently iff the incentive adjustments $\mu_m(r, r')$ and $\mu_n(r, r')$ coincide. 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 $\mu_n(r, r') > \mu_m(r, r')$ for all r

and $r' > r$.

This representation (18)-(19) generalizes the redistributive arbitrage argument of Theorem 1 to an arbitrary number of goods and arbitrary individual and social preferences. Fix $r \in (0, 1)$ and consider a perturbation such that: (i) the consumption of good n increases for all $r' \geq r$; (ii) the consumption of good m decreases for all $r' \geq r$; (iii) the utility of rank r remains unchanged; (iv) incentive-compatibility is preserved for all $r' \geq r$. The unique perturbation $\{\delta x_n(r'), \delta x_m(r')\}$ that satisfies these four requirements is given by $\delta x_k(r') = \frac{1}{U_k(r')} \mu_k(r, r') \Delta$, for $k \in \{n, m\}$ and small positive Δ . This perturbation around the optimal allocation must keep the planner's objective function unchanged, or in other words, the resource gains from reducing consumption of good m must exactly offset the resource cost of increasing consumption of good n for $r' \geq r$, for otherwise the perturbation or its negative would lead to a strict welfare improvement. This redistributive arbitrage yields condition (18), where the incentive-adjusted marginal benefits of redistribution are characterized by (19).

Furthermore, we can also generalize Lemma 1 and thus represent $\lim_{r \rightarrow 1} B_n(r)$ in terms of observables. Taking derivatives of $M_n(r') \equiv \frac{1}{U_n(r')} \mu_n(r, r')$ with respect to r' yields

$$\frac{M'_n(r')}{M_n(r')} = \frac{U_{nr}}{U_n} - \frac{d \log U_n}{dr} = - \sum_{k=1}^N \frac{U_{nk}(r')}{U_n(r')} x_k(r') \cdot \frac{x'_k(r')}{x_k(r')}.$$

If the preference elasticities $\zeta_{nk}(r) \equiv -\frac{U_{nk}(r)}{U_n(r)} x_k(r)$ and local tail coefficients $\rho_k(r) \equiv -\frac{\partial \ln x_k(r)}{\partial \ln(1-r)}$ converge to constants ζ_{nk} and ρ_k as $r \rightarrow 1$, it then follows that $M_n(r') \sim \prod_{k=1}^N x_k(r')^{\zeta_{nk}}$ as $r \rightarrow 1$, and

$$\lim_{r \rightarrow 1} B_k(r) = \left[1 - \sum_{k=1}^N \frac{\zeta_{nk}}{\rho_k} \right]^{-1}. \quad (20)$$

Equation (20) shows that the optimal wedge at the top between any pair of commodities can be represented as a function of: (i) the distributions of consumption of all N commodities (or more specifically their Pareto tail coefficients ρ_k); and (ii) the full matrix of income and substitution effects of all commodities which is summarized by $\{\zeta_{nk}\}_{1 \leq n, k \leq N}$.

As we discussed in the context of Corollary 1, our model reveals a potential redistributive rationale for non-uniform commodity taxation, which our baseline model of Section 2 displayed through savings taxes. This rationale arises whenever two different commodities yield different incentive-adjustments $\mu_n(r, r')$. Potential departures from uniform commodity taxation are then linked to these incentive-adjustments which can in turn be mapped to observables. Our analysis thus develops a template for future empirical work that seeks to identify optimal commodity taxes

and subsidies by identifying the required marginal benefits of redistribution for any commodity, using observed distributions of consumption and estimated demand elasticities. Subsidies for basic necessities, such as subsidized rent, food stamps, public transportation, education or health services play a central role in increasing the welfare of low-income households. On the other hand, governments may also find it opportune to tax certain consumption goods favored by higher income households. One key application of this framework may be to housing which is an important budget component of most households, thus displaying important wealth effects, and which benefits from a whole array of redistributive interventions, from subsidized public housing or rent subsidies at the low end of the income distribution to mortgage interest deductions at the upper end. Our analysis may offer an efficiency rationale for implementing such policies, as well as practical guidance on how they should be structured to achieve the government's redistributive objective.

5.4 Income and Savings Taxes over the Life Cycle

As an application, we can illustrate the power of redistributive arbitrage in the generalized N -good economy, studied in Section 5.3, by exploring 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 $t = 1, \dots, T$. Their initial preference rank is drawn prior to date $t = 1$, and is private information. The households' preferences are given by

$$\mathcal{U}(\{C_t, Y_t\}; r) \equiv \sum_{t=1}^T \beta^t U(C_t, Y_t; r, t)$$

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 taxes on top earners are then given by the static trade-off between redistributing income and redistributing consumption at date t , while the age-dependent savings taxes are given by the trade-off between redistributing consumption at date t vs. consumption at date $t + 1$:

$$1 - \bar{\tau}_Y(t) = \lim_{r \rightarrow \infty} \frac{B_{Y_t}(r)}{B_{C_t}(r)} = \frac{1 - \zeta_{C_t}/\rho_{C_t} + \zeta_{C_t Y_t}/\rho_{Y_t}}{1 + \zeta_{Y_t}/\rho_{Y_t} - s_{C_t} \zeta_{C_t Y_t}/\rho_{C_t}}$$

and

$$1 + \bar{\tau}_S(t) = \lim_{r \rightarrow \infty} \frac{B_{C_{t+1}}(r)}{B_{C_t}(r)} = \frac{1 - \zeta_{C_t}/\rho_{C_t} + \zeta_{C_t Y_t}/\rho_{Y_t}}{1 - \zeta_{C_{t+1}}/\rho_{C_{t+1}} + \zeta_{C_{t+1} Y_{t+1}}/\rho_{Y_{t+1}}},$$

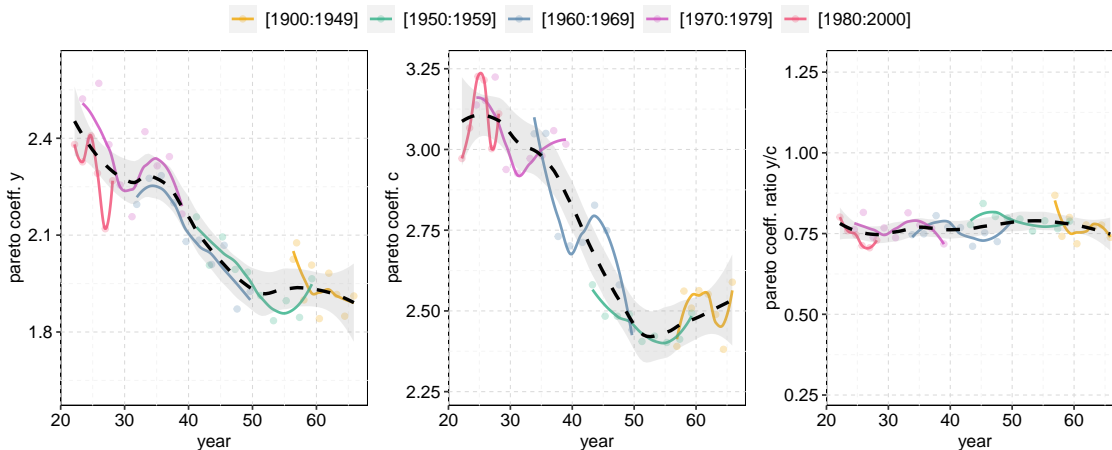
where the marginal benefits of redistribution are computed as before, but are now based on age-specific rather than unconditional preference elasticities and Pareto tail coefficients.

Following the same procedure as described in Section 4.1, we can use the data of Blundell, Pistaferri, and Saporta-Eksten (2016) to impute age-specific Pareto coefficients from top earners' consumption and income shares. This imputation gives us ball-park estimates of the variation in consumption and income inequality with age. In Figure 3, we compute the Pareto coefficients for consumption and income, as well as their ratio, by birth cohort from different PSID waves, and then plot them against age. We observe that the Pareto coefficient on income declines from about 2.4 around age 20 to about 1.8 for age 50. The Pareto coefficient for consumption displays a similar pattern but at a strictly higher level, starting from about 3 at age 20 to stabilize around 2.4 at age 50 and slightly rising again towards retirement. These figures illustrate well the growth of income and consumption inequality over the first half of the life cycle. The ratio of Pareto coefficients is remarkably stable across ages, with values between 0.75 and 0.8.

Note that our estimates of the Pareto coefficients for income by age are consistent with those found by Karahan, Ozkan, and Song (2022) using a confidential employer-employee matched panel of the earnings histories of male workers between 1978 and 2013 from the U.S. Social Security Administration. They show that lifetime earnings inequality—measured by the P90/P10 ratio—is roughly half the cross-sectional earnings inequality. They confirm that the top end of the lifetime earnings distribution follows a power law with the top 0.1% (resp., 1%) accounting for around 29% of total lifetime earnings among the top 1% (resp., 10%) of the population. This corresponds to a Pareto coefficient for lifetime earnings equal to $\rho_Y = 2.13$. Furthermore, this power law also holds in the cross-sectional distribution of earnings conditional on age. Earnings concentration at the top—measured as the relative earnings share of the top 0.1% to the top 1%—increases sharply over the life cycle from 0.23 at age 25 to 0.38 at age 55. This corresponds to Pareto coefficients at age 25 (resp., 31, 37, 43, 49, 55) equal to 2.78 (resp., 2.61, 2.21, 1.85, 1.67, 1.58).

What do these age-specific Pareto coefficients imply for the evolution of income 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 result in uniformly lower income taxes. Using $\zeta_{Y_t}^{-1} = 4/9$ and $\zeta_{C_t} = 0.75$ as in our baseline calibration along with $\rho_{C_t}/\rho_{Y_t} = 0.75$ 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 (vs. 75% in our baseline model) if there are no complementarities ($\zeta_{C_t Y_t} = 0$). With complementarities ($\zeta_{C_t Y_t}/\zeta_{C_t} = 0.15$), top optimal

Figure 3: Pareto Coefficients conditional on Age



income taxes increase from 58% at age 20 to 67% at age 50 and beyond (vs. 72% in our baseline model).

For savings taxes, the gradual increase in consumption inequality suggests that the marginal benefits of redistribution increase with age. This in turn introduces a rationale for back-loading redistribution, or taxing savings. With a consumption elasticity of 0.75 (as in our baseline model) and a ratio of Pareto tail coefficients equal to $\rho_{C_t}/\rho_{Y_t} = 0.75$, comparing the marginal benefits of redistributing consumption at age 20 vs. age 50 implies a cumulative savings tax over 30 years of 7.7% (with preference complementarity) to 10% (without preference complementarity), or equivalently to about 0.26% to 0.36% per annum, before dropping to zero beyond age 50. These estimates are smaller than the ones 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 or wealth inequality in the cross-section.

Of course these numbers should be taken to be at best suggestive, since the model abstracts—importantly—from life-cycle uncertainty and income shocks that accumulate and contribute to income inequality with age. They also assume that preferences are age-independent, which is of course a strong assumption: For example, it would seem reasonable to assume that labor supply may be more elastic for younger or older workers who have some margin of control over when to transition from education to full-time employment, and from full-time employment to retirement. Nevertheless, the results highlight how thinking about optimal redistribution as an arbitrage between different policy margins has the potential to yield novel insights about the optimal design of tax policies.

5.5 Further Extensions

Heterogeneous Initial Capital. In the Appendix, we study a special case of the general environment of Section 5.3 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.

Multi-Dimensional Types. We conclude by briefly discussing another important extension that is outside the scope of the present paper. The assumption of a one-dimensional type (“rank”) space becomes more difficult to justify as one moves beyond a single consumption good, since there is no reason why individual ability should be perfectly aligned with tastes for different commodities, for example. In line with this assumption, our derivation of sufficient statistics made use of the fact that consumption, income, and savings were perfectly co-monotonic at the optimal solution. Such perfect co-monotonicity seems implausible from an empirical point of view, even with a simple commodity space with three goods, like ours. Another natural extension is therefore to extend the present analysis to multi-dimensional type spaces. While multi-dimensional screening is notoriously

challenging, due to the lack of conclusive results about the validity of the first-order approach to optimal screening, Kleven, Kreiner, and Saez (2009, Online Appendix) suggest that the first-order approach can be applied in specific tax settings.²⁷ Assuming that the first-order approach is valid, preliminary results in Hellwig (2022) show that core ideas from the present analysis generalize to multi-dimensional screening problems, in particular the representation of incentive-preserving perturbations and the characterization of optimal relative price distortions by a generalization of the redistributive arbitrage formula presented in equation (18). These preliminary results suggest that there is scope to generalize the core idea of redistributive arbitrage to multi-dimensional type spaces.

6 Conclusion

We develop a new perspective on optimal tax design, based on the idea that optimal allocations trade off not only between efficiency and redistribution, but also between the margins along which redistribution takes place. The optimal tax system then equalizes the marginal benefit of redistribution from higher to lower ranks for all goods, around any given rank r , a property that we call redistributive arbitrage. As our main result, we derived a simple new formula for optimal tax distortions based on redistributive arbitrage. We show how to infer the respective marginal benefits of redistribution from income and consumption data and key preference elasticities, thus giving empirical content to this new perspective on optimal tax design.

As our main policy implication, our calibration results suggest that there may be significant gains from taxing and redistributing savings at the top of the income distribution. Our model suggests that it may be optimal to tax savings (wealth) by up to 2% per year, while lowering top income taxes substantially relative to existing sufficient statistics calibrations. These results are consistent with the empirical observation that savings, like income, appear to be far more unequally distributed than consumption, suggesting potential welfare gains from shifting redistribution from consumption towards savings.

The importance of multiple dimensions of worker welfare—e.g., leisure and consumption—is both historically and contemporaneously well documented. This generates trade-offs between different margins of redistributing welfare. Redistributive arbitrage formalizes how these trade-offs are resolved by optimal tax policies. In practice, many policy makers probably develop an intuitive understanding for redistributive arbitrage, when determining what policies are popular

²⁷See also the recent work by Golosov and Krasikov (2022). Both papers show that the first-order approach can be valid absent participation constraints.

with their voters and matter for their welfare. In fact, the Roman emperors are perhaps the first rulers on record to perform redistributive arbitrage, since they already knew that the most cost-effective way to keep their working population happy was to provide them with a combination of *panem et circenses*, or bread and entertainment!²⁸

References

- Aguiar, Mark and Erik Hurst (2007). “Measuring trends in leisure: The allocation of time over five decades”. In: *The Quarterly Journal of Economics* 122.3, pp. 969–1006.
- 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.
- Auclert, Adrien (2019). “Monetary policy and the redistribution channel”. In: *American Economic Review* 109.6, pp. 2333–67.
- 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.
- 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.

²⁸To be fair, the Roman poet Juvenal coined the phrase *panem et circenses* in the early 2nd century to mock the high levels of political corruption, motives that are outside the tradeoffs considered by our benevolent social planner. But what worked for a corrupt Roman politician also works for a benevolent Mirrleesian planner, as long as the working population’s welfare depends on being provided the right mix of bread and entertainment.

- Cremer, Helmuth, Pierre Pestieau, and Jean-Charles Rochet (2003). “Capital income taxation when inherited wealth is not observable”. In: *Journal of Public Economics* 87.11, pp. 2475–2490.
- 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 James Mirrlees (1978). “A model of social insurance with variable retirement”. In: *Journal of Public Economics* 10.3, pp. 295–336.
- 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.
- Diamond, Peter and Johannes Spinnewijn (2011). “Capital income taxes with heterogeneous discount rates”. In: *American Economic Journal: Economic Policy* 3.4, pp. 52–76.
- 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 (2010). “Progressive estate taxation”. In: *The Quarterly Journal of Economics* 125.2, pp. 635–673.
- (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.
- Gaillard, Alexandre and Philipp Wangner (2021). “Wealth, Returns, and Taxation: A Tale of Two Dependencies”. In: *Available at SSRN 3966130*.
- Gauthier, Stéphane and Fanny Henriët (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 and Ilia Krasikov (2022). *Multidimensional Screening in Public Finance: The Optimal Taxation of Couples*. Working Paper. University of Chicago.

- 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.
- Hellwig, Christian (2021). *Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach*. TSE Working Paper.
- (2022). *Multi-dimensional Screening: a First-Order Approach*. Work in progress.
- Jacobs, Bas and Robin Boadway (2014). “Optimal linear commodity taxation under optimal non-linear income taxation”. In: *Journal of Public Economics* 117, pp. 201–210.
- 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.
- Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez (2009). “The optimal income taxation of couples”. In: *Econometrica* 77.2, pp. 537–560.
- 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.
- Meyer, Bruce D and James X Sullivan (2017). *Consumption and Income Inequality in the US Since the 1960s*. Tech. rep. National Bureau of Economic Research.
- 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.
- Piketty, Thomas and Emmanuel Saez (2013). “A theory of optimal inheritance taxation”. In: *Econometrica* 81.5, pp. 1851–1886.
- Romer, Christina (2011). *Work-life balance and the economics of workplace flexibility*. DIANE publishing.
- Saez, Emmanuel (2001). “Using elasticities to derive optimal income tax rates”. In: *The review of economic studies* 68.1, pp. 205–229.

- Saez, Emmanuel (2002). “The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes”. In: *Journal of Public Economics* 83.2, pp. 217–230.
- Saez, Emmanuel and Stefanie Stantcheva (2018). “A simpler theory of optimal capital taxation”. In: *Journal of Public Economics* 162, pp. 120–142.
- Saez, Emmanuel and Gabriel Zucman (2019). “Progressive wealth taxation”. In: *Brookings Papers on Economic Activity* 2019.2, pp. 437–533.
- Scheuer, Florian and Joel Slemrod (2021). “Taxing our wealth”. In: *Journal of Economic Perspectives* 35.1, pp. 207–30.
- Schieman, Scott, Philip J Badawy, Melissa A. Milkie, and Alex Bierman (2021). “Work-life conflict during the COVID-19 pandemic”. In: *Socius* 7.
- Schulz, Karl (2021). *Redistribution of Return Inequality*. CESifo Working Paper.
- Shourideh, Ali (2012). “Optimal taxation of wealthy individuals”. In: *Work. pap. U. of Pennsylvania*.
- Straub, Ludwig (2019). “Consumption, savings, and the distribution of permanent income”. In: *Unpublished manuscript, Harvard University*.
- 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 “A Fair Day’s Pay for a Fair Day’s Work”

Christian Hellwig, Nicolas Werquin

A Proofs and Derivations

Proof of Theorem 1. 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 $r > 0$.

We solve it as an optimal control problem using $W(\cdot)$ 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.

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 $\phi'(r) = -\frac{\partial \mathcal{H}}{\partial W}$, 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 $\phi(0) = \phi(1) = 0$. Define $\frac{U_{Cr}(r)}{U_C(r)} = \frac{m'_C(r)}{m_C(r)}$, or $m_C(r) = \exp\left(-\int_r^1 \frac{U_{Cr}(r')}{U_C(r')} dr'\right)$. Substituting into the above 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

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

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

$$\begin{aligned} \frac{\phi(r)}{1-r} &= \mathbb{E} \left[\frac{1}{U_C(r')} \frac{m_C(r')}{m_C(r)} \Big| r' \geq r \right] - \frac{\mathbb{E} \left[\frac{1}{U_C(r')} \frac{m_C(r')}{m_C(r)} \right] \mathbb{E} \left[\omega(r') \frac{m_C(r')}{m_C(r)} \Big| r' \geq r \right]}{\mathbb{E} \left[\omega(r') \frac{m_C(r')}{m_C(r)} \right]} \\ &\equiv \frac{1}{U_C(r)} B_C(r). \end{aligned}$$

Notice that $\frac{m_C(r')}{m_C(r)} = \mu_C(r, r')$ defined in the text. 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,

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

Let $m_Y(r) = \exp\left(-\int_r^1 \frac{U_{Yr}(r')}{U_Y(r')} dr'\right)$ and apply the same steps as above to get

$$\begin{aligned} \frac{\phi(r)}{1-r} &= \mathbb{E} \left[\frac{1}{-U_Y(r')} \frac{m_Y(r')}{m_Y(r)} \Big| r' \geq r \right] - \frac{\mathbb{E} \left[\frac{1}{-U_Y(r')} \frac{m_Y(r')}{m_Y(r)} \right] \mathbb{E} \left[\omega(r') \frac{m_Y(r')}{m_Y(r)} \Big| r' \geq r \right]}{\mathbb{E} \left[\omega(r') \frac{m_Y(r')}{m_Y(r)} \right]} \\ &\equiv \frac{1}{-U_Y(r)} B_Y(r), \end{aligned}$$

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

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

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

Let $m_S(r) = \exp\left(-\int_r^1 \frac{V_{Sr}(r')}{V_S(r')} dr'\right)$. The previous ODE can be integrated and solved along the same lines as above to find

$$\begin{aligned} \frac{\phi(r)}{1-r} &= \mathbb{E} \left[\frac{1}{V_S(r')} \frac{m_S(r')}{m_S(r)} \Big| r' \geq r \right] - \frac{\mathbb{E} \left[\frac{1}{V_S(r')} \frac{m_S(r')}{m_S(r)} \right] \mathbb{E} \left[\omega(r') \frac{m_S(r')}{m_S(r)} \Big| r' \geq r \right]}{\mathbb{E} \left[\omega(r') \frac{m_S(r')}{m_S(r)} \right]} \\ &= \frac{1}{V_S(r)} B_S(r), \end{aligned}$$

with $\lambda = \frac{\mathbb{E}[m_S/V_S]}{\mathbb{E}[m_S\omega]}$. Equating this last expression to $\frac{1}{U_C(r)}B_C(r)$ then yields the expression for the savings wedge:

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

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

Proof of Corollary 1. We saw in the proof of Theorem 1 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 $\frac{U_{Cr}(r)}{U_C(r)} - \frac{V_{Sr}(r)}{V_S(r)}$ has a constant sign, we get $U_C(r) \leq V_S(r)$, or $\tau_S(r) \geq 0$ for all r , if and only if $\frac{U_{Cr}(r)}{U_C(r)} - \frac{V_{Sr}(r)}{V_S(r)} \leq 0$ for all r .

More generally, consider the general framework of Section 5.3. For any two goods $m < n$, suppose that the marginal rate of substitution $\frac{U_m(r)}{U_n(r)}$ is weakly increasing in r , so that $\frac{U_n(r)}{U_n(r')} \geq \frac{U_m(r)}{U_m(r')}$ for all $r' > r$. Equivalently, $\frac{U_{mr}(r)}{U_m(r)} \geq \frac{U_{nr}(r)}{U_n(r)}$ for all r , or $\mu_m(r, r') \geq \mu_n(r, r')$ for all r, 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 two incentive adjustments $\mu_m(r, r')$ and $\mu_n(r, r')$ coincide, or equivalently iff the MRS $U_m(r)/U_n(r)$ is uniform across types. More generally, we have $\frac{U_m(r)}{U_n(r)} \frac{p_n}{p_m} < 1$, so that $\tau_{m,n}(r) > 0$, iff $\frac{U_{nr}(r)}{U_n(r)} > \frac{U_{mr}(r)}{U_m(r)}$, or equivalently $\mu_n(r, r') > \mu_m(r, r')$. \square

Proof of Lemma 1. 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 elasticities and Pareto coefficients $\rho_C(r)$, $\rho_Y(r)$, $\rho_S(r)$ introduced in Section 3.1, the two first-order conditions $-\frac{U_Y}{U_C} = 1 - \tau_Y$ and $\frac{V_S}{U_C} = 1 + \tau_S$, and noting that $\frac{CU_{CY}}{-U_Y} = \frac{C}{(1-\tau_Y)Y} \frac{YU_{CY}}{U_C} = s_C \zeta_{CY}$ implies that these three equations can be rewritten as

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

Using the definition of $\rho_{U_C}(r)$ and rearranging terms leads to equations (7), (8), and (9). It follows immediately that

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

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

$$\begin{aligned} M_C(r) &= \frac{1}{U_C(r)} e^{-\int_r^1 \frac{d}{dr'} \frac{U_C(r')}{U_C(r')} dr'} e^{\int_r^1 \left\{ -\zeta_C(r') \frac{C'(r')}{C(r')} + \zeta_{CY}(r') \frac{Y'(r')}{Y(r')} \right\} dr'} \\ &= e^{\int_r^1 \left\{ -\frac{\zeta_C(r')}{\rho_C(r')} + \frac{\zeta_{CY}(r')}{\rho_Y(r')} \right\} \frac{dr'}{1-r'}}, \end{aligned}$$

and similarly

$$\begin{aligned} M_Y(r) &= \frac{1}{-U_Y(r)} e^{-\int_r^1 \frac{d}{dr'} \frac{(-U_Y(r'))}{-U_Y(r')} dr'} e^{\int_r^1 \left\{ \zeta_Y(r') \frac{Y'(r')}{Y(r')} - s_C(r') \zeta_{CY}(r') \frac{C'(r')}{C(r')} \right\} dr'} \\ &= e^{\int_r^1 \left\{ \frac{\zeta_Y(r')}{\rho_Y(r')} - \frac{s_C(r') \zeta_{CY}(r')}{\rho_C(r')} \right\} \frac{dr'}{1-r'}}, \end{aligned}$$

and

$$M_S(r) = \frac{1}{V_S(r)} e^{-\int_r^1 \frac{d}{dr'} \frac{V_S(r')}{V_S(r')} dr'} e^{-\int_r^1 \zeta_S(r') \frac{S'(r')}{S(r')} dr'} = e^{-\int_r^1 \frac{\zeta_S(r')}{\rho_S(r')} \frac{dr'}{1-r'}}.$$

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

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

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

Proof of Theorem 2. It follows from Assumption 2 and the previous proof that

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

For the numerator, define $X(r) \equiv (Y(r))^{-\zeta_Y} (C(r))^{s_C \zeta_{CY}}$. We wish to compute $\mathbb{E} \left[\frac{X(r')}{X(r)} \mid r' \geq r \right]$, 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 $\left(-\frac{\zeta_Y}{\rho_Y} + \frac{s_C \zeta_{CY}}{\rho_C} \right)^{-1}$. This implies

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

Along the same lines,

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

and therefore

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

At the optimal allocation, $B_C(r)$ must be finite, and therefore $\frac{\zeta_C}{\rho_C} < 1 + \frac{\zeta_{CY}}{\rho_Y}$. It then follows automatically that $\lim_{r \rightarrow 1} \tau_Y(r) < 1$. To prove the second part of Theorem 2, follow analogous steps as above to get

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

for $\zeta_S/\rho_S < 1$. Combining this result with $\lim_{r \rightarrow 1} B_C(r) = \left[1 - \frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y} \right]^{-1}$, we get

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

This concludes the proof. \square

Relationship with Ferey, Lockwood, and Taubinsky (2021). 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:²⁹

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

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 $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$ is strictly increasing in $\frac{\zeta_S(r)}{\zeta_C(r)}$, for given preferences, spending shares, and Pareto tails. However, if $\lim_{r \rightarrow 1} s_C(r) = 0 = 1 - \lim_{r \rightarrow 1} s_S(r)$, then $\lim_{r \rightarrow 1} \frac{\partial \ln S(Y,r)}{\partial \ln Y} = 1$ and $\lim_{r \rightarrow 1} \left(-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}\right) = 0$, regardless of the other parameters, which confirms that the identifying power of $\frac{\partial \ln S(Y,r)}{\partial \ln Y}$ vanishes when $\lim_{r \rightarrow 1} s_C(r) = 0$ 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:

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

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 that³⁰

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

where $s_S(r) \zeta_C(r) + s_C(r) \zeta_S(r)$ represents the inverse of the inter-temporal elasticity of substitu-

²⁹Consider 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 1 to substitute out $\frac{\partial \tau_S}{1 + \tau_S} + \frac{U_{Cr}}{U_C}$ yields the expression for $-\frac{\partial \ln S(Y,r)}{\partial \ln(1-r)}$.

³⁰Consider 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 $-\frac{\partial \ln S(Y,r)}{\partial \ln(1 + \tau_S)} \Big|_{Y,T(Y,S) \text{ constant}}$.

tion. 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 $\lim_{r \rightarrow 1} s_C(r)$, 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 $r \rightarrow 1$.

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, $\mathbb{E}[1 - \hat{g}(r') | r' \geq r]$, converges to $B_C(r)$. 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 $\lim_{r \rightarrow 1} \mathbb{E}[1 - \hat{g}(r') | r' \geq r] = \lim_{r \rightarrow 1} B_C(r)$.

In addition, we can rewrite equation (18) in FLT as

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

where the compensated income elasticity ζ_Y^c satisfies³¹

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

Substituting $s_s \frac{\partial \ln S(Y,r)}{\partial \ln Y} = \frac{s_S \zeta_C (1 - s_C \frac{\zeta_{CY}}{\zeta_C})}{s_S \zeta_C + s_C \zeta_S}$ then allows us to evaluate the above expression in the limit as $r \rightarrow 1$: If $\lim_{r \rightarrow 1} s_C = 1$ (Case 1), it follows that $\frac{1}{\zeta_Y^c} = \zeta_Y - \zeta_{CY} + \zeta_C - \zeta_{CY}$ and $s_s \frac{\partial \ln S(Y,r)}{\partial \ln Y} \rightarrow 0$ so $\frac{\tau_Y}{1 - \tau_Y} = \{\zeta_Y - \zeta_{CY} + \zeta_C - \zeta_{CY}\} \frac{1}{\rho_Y} \mathbb{E}[1 - \hat{g}(r') | r' \geq r]$. If $\lim_{r \rightarrow 1} s_C = 0$ (Case 2), it follows that $\frac{1}{\zeta_Y^c} = \zeta_Y + \zeta_S$ and $s_s \frac{\partial \ln S(Y,r)}{\partial \ln Y} \rightarrow 1$ so $\frac{\tau_Y}{1 - \tau_Y} = \left\{ \zeta_Y + \zeta_C \left(\frac{\rho_Y}{\rho_C} - \frac{\zeta_{CY}}{\zeta_C} \right) \right\} \frac{1}{\rho_Y} \mathbb{E}[1 - \hat{g}(r') | r' \geq r]$. Finally if $\lim_{r \rightarrow 1} s_C(r) \in (0, 1)$ (Case 3), $\frac{\rho_Y}{\rho_S} = \frac{\rho_Y}{\rho_C} = 1$, and $\frac{1}{\zeta_Y^c} - s_s \frac{\partial \ln S(Y,r)}{\partial \ln Y} \left(\frac{\rho_Y}{\rho_S} \zeta_S - \zeta_C \left(\frac{\rho_Y}{\rho_C} - \frac{\zeta_{CY}}{\zeta_C} \right) \right)$ converges to $\zeta_Y - \zeta_{CY} + \zeta_C - s_C \zeta_{CY}$. In all three cases, equation (18) in FLT yields

$$\frac{\bar{\tau}_Y}{1 - \bar{\tau}_Y} = \lim_{r \rightarrow 1} A(r) \mathbb{E}[1 - \hat{g}(r') | r' \geq r]$$

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 $\lim_{r \rightarrow 1} \mathbb{E}[1 - \hat{g}(r') | r' \geq r] = \lim_{r \rightarrow 1} B_C(r)$. \square

Income and Substitution Effects: Hicksian and Marshallian Elasticities. Consider a labor income tax schedule $T_Y(Y)$ and a savings tax schedule $T_S(S)$. For ease of notation, assume

³¹Consider 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^c .

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

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

and

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

with

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

We obtain the responses of income, consumption and savings by taking first-order Taylor expansions of the two perturbed FOCs as $\delta \rightarrow 0$:

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

and

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

where $\tilde{\zeta}_C \equiv \zeta_C - s_C \zeta_{CY}$, $\tilde{\zeta}_Y \equiv \zeta_Y - \zeta_{CY}$, $\tilde{\zeta}_S \equiv \zeta_S + s_S \zeta_{CY}$. Note that as $r \rightarrow 1$, so that $Y, S \rightarrow \infty$ and T_Y', T_S' converge to constants, we have $s_C + s_S \rightarrow 1$. Solving this system leads to

$$\frac{\partial Y}{Y} = -\zeta_Y^H \frac{\partial T_Y'}{1 - T_Y'} + \zeta_Y^I \frac{\partial T_Y}{(1 - T_Y') Y},$$

with

$$\zeta_Y^H = \frac{1}{\tilde{\zeta}_Y + \frac{\tilde{\zeta}_C \tilde{\zeta}_S}{s_S \tilde{\zeta}_C + s_C \tilde{\zeta}_S}}, \quad \text{and} \quad \zeta_Y^I = \frac{\frac{\tilde{\zeta}_C \tilde{\zeta}_S}{s_S \tilde{\zeta}_C + s_C \tilde{\zeta}_S}}{\tilde{\zeta}_Y + \frac{\tilde{\zeta}_C \tilde{\zeta}_S}{s_S \tilde{\zeta}_C + s_C \tilde{\zeta}_S}}.$$

In particular, when $s_C \rightarrow 1$ and $s_S \rightarrow 0$ (Case 1), we have $\zeta_Y^H = \frac{1}{\tilde{\zeta}_Y + \tilde{\zeta}_C}$ and $\zeta_Y^I = \frac{\tilde{\zeta}_C}{\tilde{\zeta}_Y + \tilde{\zeta}_C}$. When $s_C \rightarrow 0$ and $s_S \rightarrow 1$ (Case 2), we have $\zeta_Y^H = \frac{1}{\tilde{\zeta}_Y + \tilde{\zeta}_S}$ and $\zeta_Y^I = \frac{\tilde{\zeta}_S}{\tilde{\zeta}_Y + \tilde{\zeta}_S}$. \square

Calibration for Case 3. In case 3, the Pareto coefficients of consumption, income, and savings must coincide: $\rho_Y = \rho_C = \rho_S$. We set this parameter to 1.5, the value we used for income and savings in the calibration of Case 2. To calibrate the elasticities, we take $\zeta_Y^H = 1/3$, $\zeta_Y^I = 1/4$. Using the expressions derived above and imposing that the risk aversion parameters are the same in both periods, so that $\zeta_C = \zeta_S$, we obtain $\zeta_C = \frac{\zeta_Y^I}{\zeta_Y^H} + s_C \zeta_{CY}$ and $\zeta_Y = \frac{1}{\zeta_Y^H} - \frac{\zeta_Y^I}{\zeta_Y^H} + s_C \zeta_{CY} \left(1 - \frac{s_S \zeta_{CY}}{\zeta_Y^I / \zeta_Y^H + s_C \zeta_{CY}}\right)$.

In our benchmark calibration, we take $\zeta_{CY} = 0$ and get $\zeta_C = \zeta_S = 3/4$ and $\zeta_Y = 9/4 = 2.25$. We finally need to calibrate the consumption share s_C . To do so, note first that, by the above derivations, we can express the consumption response to a lump-sum tax transfer, or marginal propensity to consume (MPC), as

$$\frac{\partial C}{-\partial T_Y} = s_C \frac{\tilde{\zeta}_Y}{\tilde{\zeta}_C} \zeta_Y^I.$$

We match an MPC of top income earners of 0.2 (see Figure 2 in Auclert (2019)). This implies $s_C = \frac{4}{3} MPC = 0.27$.

In this benchmark calibration with $\zeta_C = \zeta_S$ and $\zeta_{CY} = 0$, we obtain an optimal savings wedge $\bar{\tau}_S = 0$ and an optimal labor wedge $\bar{\tau}_Y = \bar{\tau}_Y^{Saez} = 80\%$. This is a consequence of the Atkinson-Stiglitz theorem, or Corollary 1. Indeed, preferences are then separable and the utility of consumption is homogeneous across consumers. This implies that the benefits of redistributing via consumption and savings are then identical: $B_C = 1/(1 - \zeta_C/\rho_C)$ and $B_S = 1/(1 - \zeta_S/\rho_S)$.

Now, when preferences are non-separable (or when $\zeta_C \neq \zeta_S$), it becomes optimal to distort savings. We take $\zeta_{CY}/\zeta_C = 0.15$ (the upper bound in Chetty (2006)) and $MPC = 0.2$. Solving the non-linear system of three equations in three unknowns ζ_C, ζ_Y, s_C derived above, leads to $\zeta_C = \zeta_S = 0.79$, $\zeta_Y = 2.29$, and $s_C = 0.35$. As in Case 2, the complementarity between consumption and income raises the optimal savings wedge and lowers the labor wedge: We get $\bar{\tau}_Y = 78\%$ and $\bar{\tau}_S = 17\%$. \square

Extension to a Model with Heterogeneous Endowments. 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 Theorem 1, 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:

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

with

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

Under Assumption 2, these marginal benefits converge to

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

where $s_Z = \lim_{r \rightarrow 1} \frac{Z'(r)}{C'(r)} = \frac{\rho_C}{\rho_Z} \lim_{r \rightarrow 1} \frac{Z(r)}{C(r)}$ 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 are given by equations (10)-(12) and 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., $\lim_{r \rightarrow 1} B_Y(r)$ remains the same as in our baseline model, and hence endowments only affect the combined wedge $\frac{1-\bar{\tau}_Y}{1+\bar{\tau}_S} = \frac{1-\zeta_S/\rho_S}{1+\zeta_Y/\rho_Y}$ 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 $B_C(r) \rightarrow 0$ 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 $\bar{\tau}_Y$ and $\bar{\tau}_S$ 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 $s_C \frac{\zeta_{CY}}{\rho_C} \frac{\bar{B}_Y}{B_C} < 1$. If instead consumption has a strictly thinner tail ($\rho_Z = \rho_Y < \rho_C$) then $s_C \rightarrow 0$, $s_Z \rightarrow \infty$ and $B_C(r) \rightarrow 0$, resulting as before in arbitrarily large earnings subsidies and savings taxes at the top.

In Case 3., $s_Z s_C = \infty$ and $B_Y(r) \rightarrow 0$, 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 $\bar{\tau}_S$ is finite and strictly larger than in our baseline economy, while the labor wedge becomes arbitrarily large ($\bar{\tau}_Y \rightarrow 1$). If $\rho_Z < \rho_C < \rho_Y$, we have both $s_Z \rightarrow \infty$ and $s_C \rightarrow \infty$ 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 $1 - \bar{\tau}_Y = \frac{\zeta_C}{s_C \zeta_{CY}}$. If $\rho_C > \rho_Y$, we obtain $\bar{\tau}_Y = -\infty$, 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 contrast, when s_C is low, the wealth effect of income on savings dominates, which makes it optimal to subsidize income. In the limit when $s_C \rightarrow 0$, and *a fortiori* when $\rho_C > \rho_Y$, the implied savings subsidy becomes arbitrarily large at the top. \square

Additional Graphs. Figure 4 reports the Pareto coefficients of the consumption distribution between 1998 and 2014, computed analogously to those of Figure 4 in the main text. Similarly, Figure 5 reports the Pareto coefficients of the income distribution between 1998 and 2014.

Figure 6 plots the tail distributions of the rates of return calculated from the SCF by Gaillard and Wangner (2021) between 1998 and 2014 using a threshold of the top 95%; we refer to this paper for the construction of the data. Returns are defined as (one plus) the ratio of income from investments to wealth. There seems to be systematic deviation from the linear relationship, which tends to indicate that rates of return are lognormally distributed rather than Pareto distributed at the tail. (Recall that our formulas hold regardless of the underlying distribution of returns.) To be consistent with our analysis that imposes a one-to-one map between returns and income, Figure 6 plots the distribution of log-returns against that of log-income, following a procedure analogous to that used to construct Figure 2, using quantiles between 0.80 to 0.90 in increments of 0.05, as well as 0.925, 0.95, 0.96, 0.97, 0.98, 0.99, and 0.995. The relationship is very noisy and unstable.

Finally, Figure 8 shows the Pareto tail of wealth computed using the SCF in 2014, augmented with the Forbes list data at the very top; it indicates a Pareto coefficient $\rho_{C_2} = 1.4$. \square

Figure 4: Pareto Coefficients of Consumption

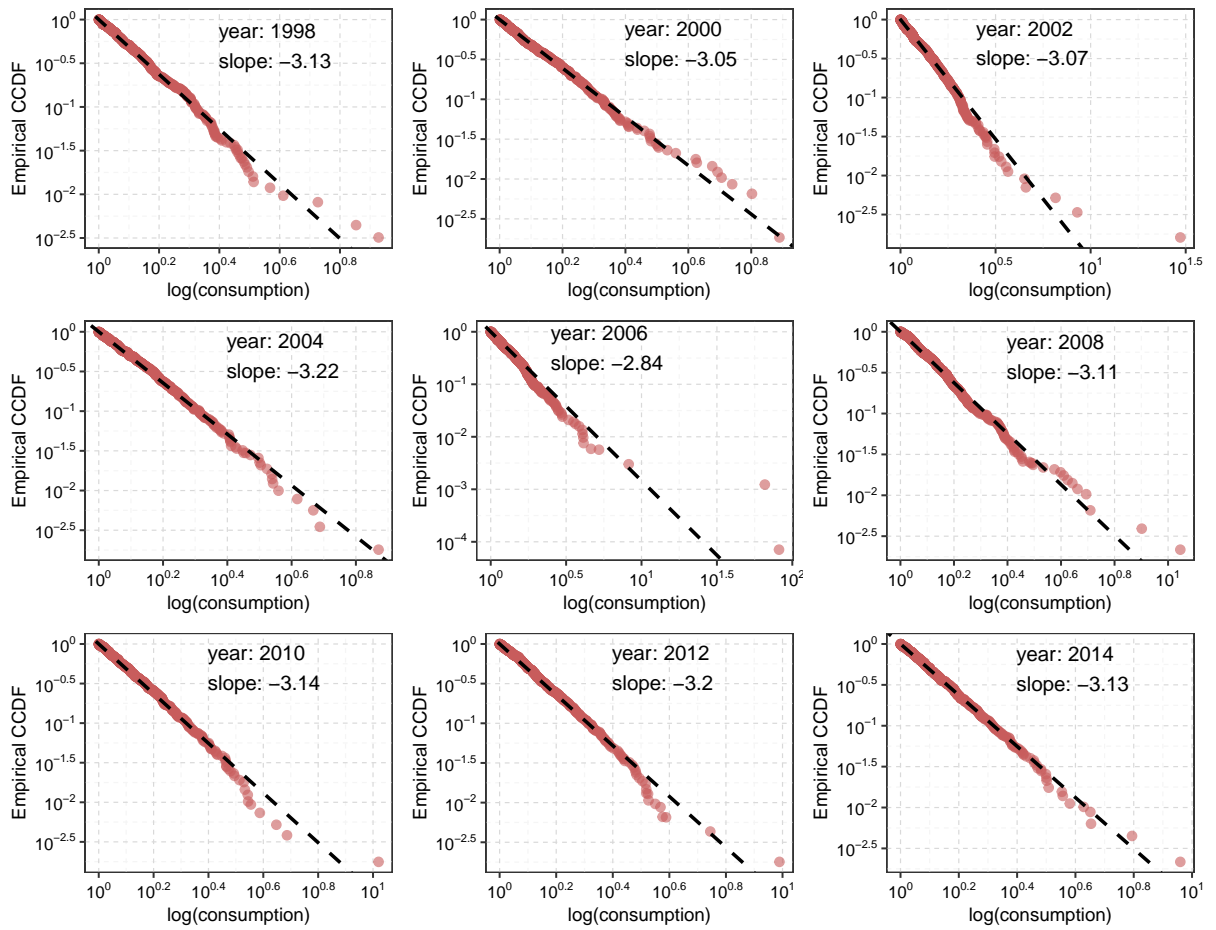


Figure 5: Pareto Coefficients of Income

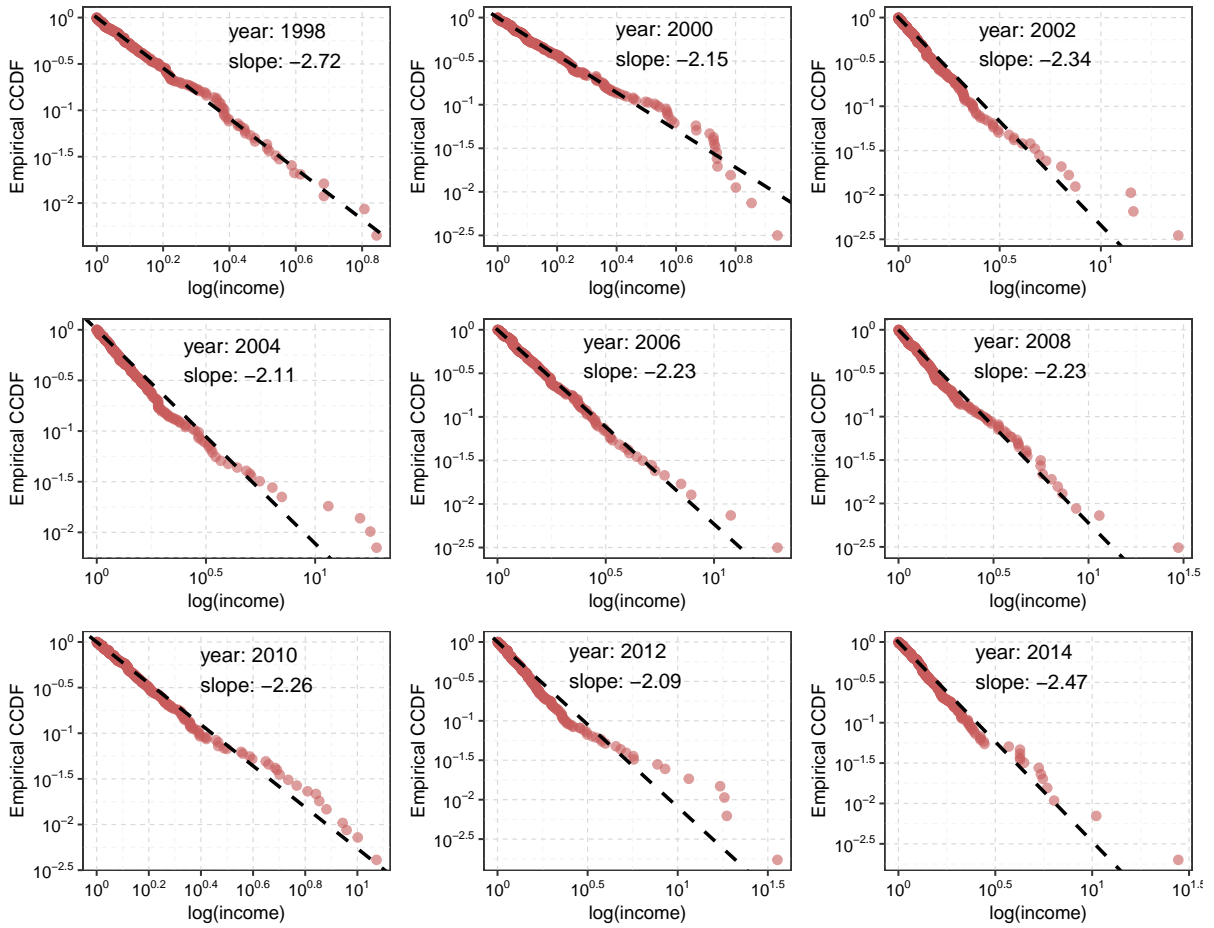


Figure 6: Pareto Coefficients of Rates of Return

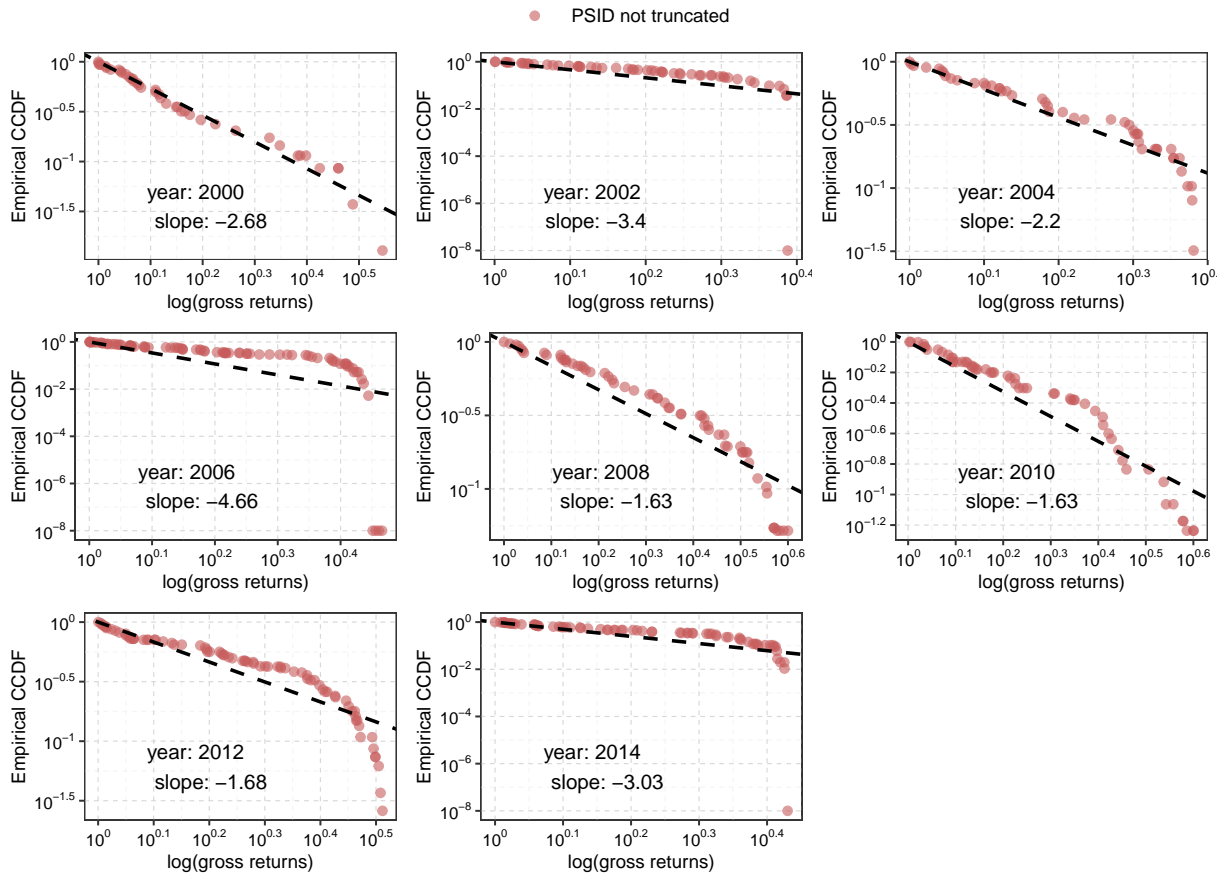


Figure 7: Ratio of Pareto Coefficients: Returns vs. Income

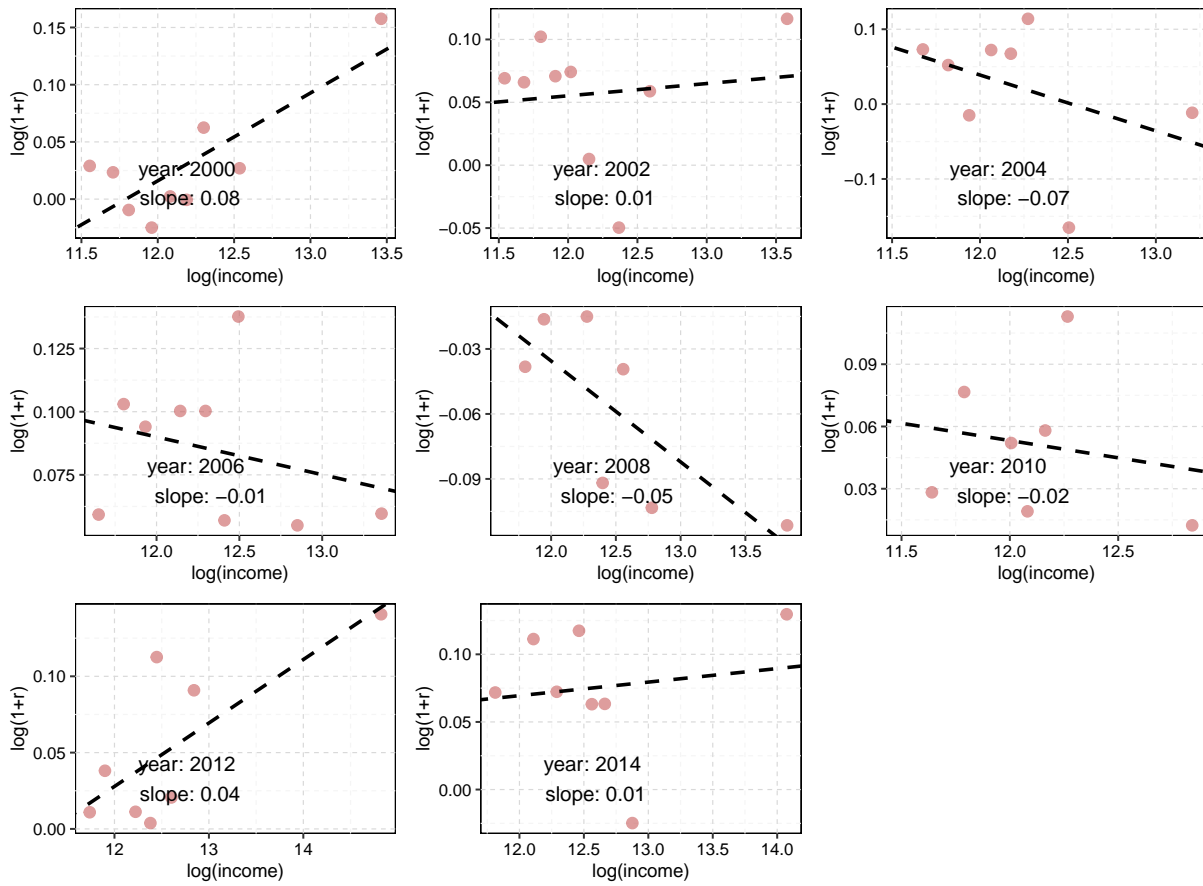


Figure 8: Pareto Coefficient of Wealth

