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

Christian Hellwig

Toulouse School of Economics and CEPR

Nicolas Werquin

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

July 19, 2022

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 earnings, consumption and retirement savings, we derive a novel formula for optimal non-linear income and savings distortions based on redistributional arbitrage. We establish a sufficient statistics representation of the labor income and capital tax rates on top income earners, which relies on the comparison between 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. Our results generalize to an economy with general preferences over an arbitrary set of commodities.

*We thank Arpad Abraham, Mark Aguiar, Gadi Barlevy, Marco Bassetto, Charlie Brendon, Ashley Craig, Joel David, Antoine Ferey, Lucie Gadenne, Alexandre Gaillard, Aart Gerritsen, Mike Golosov, Francois Gourio, Martin Hellwig, Jim Hines, Camille Landais, Alexander Ludwig, Yena Park, Dominik Sachs, Florian Scheuer, Karl Schulz, Emmanuel Saez, Stefanie Stantcheva, Aleh Tsyvinski, and Philipp Wangner for useful comments. Opinions expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Christian Hellwig acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future program (Investissements d'Avenir, grant ANR-17-EURE-0010).

"*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 tradeoff between the social benefits of redistributing financial resources from richer to poorer households, and the efficiency costs of allocative distortions that such redistribution necessarily 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 tradeoff between efficiency 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 earnings 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, education and mass entertainment, or excess taxation of goods perceived to be luxuries. Moreover, the static Mirrleesian model implicitly assumes that the government is the only provider of insurance against labor market risks. In practice, agents may smooth labor market risks through other means than the government, such as private insurance, precautionary savings, or intra-family transfers.

Second, equating consumption to after-tax income 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: The optimal top income tax drops from 80% to 50% in our preferred calibration if we instead use consumption- rather than income-based measures of inequality. In other words, Saez (2001)'s sufficient statistic 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", and consider a policy maker with a redistributional objective who designs income and savings taxes. Importantly, the policy maker is not completely free in designing redistributional policies, but these redistributional arbitrages must preserve the households' incentives to work, consume, or save as intended by the policy maker.

As our central result, we show that the optimal policy design obeys a simple principle of *redistributional arbitrage*. The policy maker has three means of extracting resources from the richest households: reducing their consumption, reducing their leisure (i.e., asking them to work more), or reducing their wealth (taxing their savings). The optimal tax on labor income equalizes the resources the policy maker can raise by asking the rich to work more, or reducing their leisure, to the marginal resource gain from reducing their consumption. Similarly the optimal savings tax equalizes the marginal resource gain from reducing the richest household's consumption to the marginal resource gain 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 redistributional policies along many different policy margins: 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.

Following Saez (2001), we express these marginal resource gains of redistributing consumption, leisure and savings—and hence the optimal income and savings taxes—in terms of observables, namely: the cross-sectional distribution (Pareto tail coefficients) of each good, along with standard

[1] Consumption 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.

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 earnings 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. This implies that the consumption share of income converges to zero for top income earners. Thus, these agents' incentive problem reduces to a static trade-off between leisure and savings. It follows that the formula of Saez (2001) 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, simple calibrations of our optimum formulas lend support to the conclusion that it is optimal to shift a significant share of the tax burden on top earners from income to savings, unless the consumption elasticity takes on an implausibly large value. Therefore, the formula obtained in the static model overstates the top optimal labor income tax rate, because it fails to account for the fact that consumption is less unequally distributed than after-tax incomes in the data.

Our calibration suggests that top savings taxes should 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 their 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 (2019a) and Saez and Zucman (2019b)). This shift from income towards savings taxes is a fairly robust feature of our quantitative results, and is driven by a combination of low consumption elasticities and/or thinner consumption tails at the top of the income distribution. 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.

While we are not aware of prior discussions or formalizations of redistributional arbitrage or

related ideas in the economics literature on optimal tax design, the observation that redistributional 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]

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). By viewing tax policies as an arbitrage between different margins of redistribution, we generalize Saez's representation of optimal income taxes to a dynamic, or multiple-good, environment and derive a companion formula for optimal savings taxes. In linking this 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 analysis of Toda and Walsh (2015) and Straub (2019), who show that the Pareto tail of the distribution of consumption is significantly thinner than that of

[2] According 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 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.

the income distribution.[3]

Our model is based on Atkinson and Stiglitz (1976). Because we allow for arbitrary preferences, their uniform commodity taxation theorem only applies as special case of our framework.[4] 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 or express the formulas in terms of empirically observable sufficient statistics. Gerritsen et al. (2020) and Schulz (2021) focus on a different departure from the Atkinson-Stiglitz benchmark, namely, a model with heterogeneous returns.[5] Saez and Stantcheva (2018) characterize optimal top labor and capital tax rates under restricted individual preferences. The papers that are closest to our work are Ferey, Lockwood, and Taubinsky (2021) and Scheuer and Slemrod (2021). They study a model similar to ours,[6] but using different approaches. First, our optimal tax formulas rely on a different set of perturbations and lead to redistributional arbitrage expressions that bear little resemblance to the "ABC" expressions derived in these papers by perturbing tax schedules rather than allocations. Conceptually, our novel representation offers a unified perspective on the optimal design of taxes on multiple goods (income and savings). Second, our representation maps to a different—and, we think, particularly direct and transparent—set of empirically observable sufficient statistics: specifically, the Pareto tail coefficients on income and consumption and standard elasticity parameters. Finally, Scheuer and Slemrod (2021), like us, also derive a characterization of the capital tax rates on top earners. In contrast to our analysis, however, they take as given the labor income tax. We discuss in more detail the relationship between our paper 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.

[3] This 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.

[4] Several 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)).

[5] Our general framework of Section 5 nests this case. On the other hand, these papers explore various microfoundations of return heterogeneity that are beyond the scope of our analysis.

[6] The model of Scheuer and Slemrod (2021) has heterogeneous endowments and separable preferences. In Section 5, we apply our methodology to an extension of our baseline model that has both heterogeneous endowments and arbitrary preferences.

2 Theory of Redistributional 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. There are two periods 0 and 1. The preferences of agents of rank r are defined over consumption C and labor income Y in period 0, and consumption S ("savings") in period 1.[7] They are represented as

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

where for any r, the functions U and V are twice continuously differentiable with $U_C > 0$, $U_{CC} < 0$, $U_Y < 0$, $U_{YY} < 0$, $V' > 0$, $V'' < 0$ and satisfy the usual Inada conditions as C, Y or S approach 0 or ∞.

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

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

Furthermore, the marginal disutility of effort is decreasing in r, $U_{Yr}/U_Y < 0$. The marginal utility of consumption is monotonic in r, i.e., U_{Cr}/U_C is either non-positive or non-negative everywhere.

Assumption 1 introduces a ranking of agents according to their preferences over consumption and leisure bundles: On the margin, agents with higher 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 easier 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 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: In other words, those who are the most inclined to work are also those who are the most inclined to save. If instead $U_{Cr}/U_C \geq 0$, consumption needs are higher for higher ranks, in which case the two redistribution

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

motives are not aligned: Those who are the most inclined to work are also those who are the most inclined to spend their incomes on current consumption.

Social Planner's Problem. Consumption, earnings, and savings are assumed to be observable, but an individual's preference rank r is their private information. Resources can be saved at a rate $R > 0$ from period 0 to 1. In our baseline model, we assume that the planner's objective is to maximize revenue; equivalently, the social objective is Rawlsian.[8] Thus, the optimal allocation $\{C(r), Y(r), S(r)\}$ maximizes the net present value of tax revenue:

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

subject to the incentive compatibility constraint:

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

for all types r and announcements r'.

We solve this problem using a Myersonian approach, replacing full incentive-compatibility by local incentive-compatibility. Define the indirect utility function $W(r) \equiv U(C(r), Y(r); r) + \beta V(S(r))$.[9] Then an allocation is locally incentive-compatible, if it satisfies

$$W'(r) = U_r(C(r), Y(r); r). \tag{2}$$

We refer to $U_r(r) \equiv U_r(C(r), Y(r); r)$ as the *marginal information rent* of type r. 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 earnings. Let $\tau_S(r) \equiv \beta R V'(r)/U_C(r) - 1$ denote the *savings wedge* at rank r, i.e., the inter-temporal distortion in the agent's first-order condition for savings.

[8] We generalize our analysis to arbitrary Bergson-Samuelson social welfare objectives in Section 5. In particular, the optimal top tax rate formulas of Section 3 remain valid for any social welfare function.

[9] To 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.

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*

$$1-\tau_Y(r) = \frac{B_Y(r)}{B_C(r)}, \tag{3}$$

and the optimal savings wedge τ_S *satisfies*

$$1+\tau_S(r) = \frac{B_S(r)}{B_C(r)}, \tag{4}$$

where for $X \in \{C, Y\}$,

$$B_X(r) \equiv \mathbb{E}\left[\frac{U_X(r)}{U_X(r')}\exp\left(\int_r^{r'}\frac{U_{Xr}(r'')}{U_X(r'')}dr''\right) \mid r' \geq r\right], \text{ and } B_S(r) \equiv \mathbb{E}\left[\frac{V'(S(r))}{V'(S(r'))} \mid r' \geq r\right]. \tag{5}$$

Theorem 1 summarizes the principle of *redistributional arbitrage*. It formalizes the idea that, at the optimal allocation, the planner is indifferent between redistributing slighltly less along one 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 earnings 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 (3) and (4) with the individual's first-order conditions $1-\tau_Y = -U_Y/U_C$ and $1+\tau_S = \beta R V'/U_C$ 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. 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 5.1 that our analyis can be straightforwardly extended to a framework with fully general preferences over an arbitrary number of commodities.

2.3 Perturbation-Based Interpretation of Theorem 1

To interpret the marginal benefits of redistribution B_C, B_Y, B_S, fix a given rank $r > 0$ and perturb (lower or raise) the utility of agents with rank $r' > r$. We can do so in three independent ways: by adjusting their consumption, their leisure, or their savings. The impact of any such perturbation on the social objective is given by the (positive or negative) amount of extra resources that it allows the planner to levy. In the sequel, we focus on the trade-off between redistributing via consumption and leisure, and show how an arbitrage argument between these two margins leads to formula (3). The optimum savings wedge (4) is obtained analogously as a no-arbitrage condition between redistributing via consumption and savings.

Consider the following perturbation: We simultaneously *raise* the consumption of ranks $r' \geq r$ by $\Delta C(r') > 0$ and *reduce* their leisure, or raise their output by $\Delta Y(r') > 0$, while preserving local incentive compatibility (2). 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 (3) follows immediately.

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 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

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

where $\Delta W\left(r\right) = U_C\left(r\right) \Delta C\left(r\right)$ represents the increase in utility for rank r associated with the perturbation of consumption. Therefore, $B_C\left(r\right)$ 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 $\mu_C\left(r, r'\right) \equiv \exp\left(\int_r^{r'} \frac{U_{Cr}(r'')}{U_C(r')} dr''\right)$, and consumption in proportion to $\frac{1}{U_C(r')}\mu_C\left(r, r'\right)$, thus leading to expression (5) for the marginal benefits B_C. This is because the perturbation $\Delta C\left(\cdot\right)$ changes utility levels for $r' > r$ by $\Delta W\left(r'\right) = U_C\left(r'\right) \Delta C\left(r'\right)$ and marginal information rents by $\Delta U_r\left(r'\right) = U_{Cr}\left(r'\right) \Delta C\left(r'\right)$. It therefore preserves local incentive compatibility if and only if

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

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. Consider now the second part of the perturbation, whereby the planner reduces the *leisure*, or raises the output, 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 $1/\left(-U_Y\left(r'\right)\right)$. The non-separability $U_{Yr} < 0$ then requires an incentive-adjustment $\mu_Y\left(r, r'\right) = \exp\left(\int_r^{r'} \frac{U_{Yr}(r'')}{U_Y(r'')} dr''\right)$. As a result, this perturbation frees an

amount of resources equal to $B_Y(r)\,\Delta Y(r)$.[10]

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 (3) characterizes the optimal labor income taxes regardless of the value (optimal or not) of the savings taxes. Similarly the arbitrage formula (4) 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 constrasts with the "ABC" expressions typically derived in the literature following Diamond (1998). To understand the link between these two characterizations, 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:

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

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.[11] 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 mimick rank r. This allows the planner to increase the redistribution of consumption from ranks $r' > r$, with a resource gain (per our

[10] 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 $B_S(r)$. Notice that the derivation of these marginal benefits is simplified by the fact that the marginal utility of savings is independent of r, so that no incentive-adjustment μ_S is necessary in this case.

[11] Note 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.

earlier analysis) equal to $(B_C/U_C)\,\Delta U_r$.

Analogously, the second equation ("leisure-ABC"), which is novel, equates the marginal cost of the tax distortion at r,[12] 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 $r' > r$, $B_S/(\beta R V')$. Our arbitrage representations (3) and (4) are then obtained by eliminating the marginal cost of tax distortions $A(r)$ from these ABC formulas.

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 (consumption-ABC, leisure-ABC, and redistributional 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 (3), (4) and (6) highlights that the principle of redistributional arbitrage, in contrast to the ABC representation, 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?

The uniform commodity taxation theorem of Atkinson and Stiglitz (1976) is nested as a special case of our savings wedge representation (4). Specifically, the optimal savings wedge must be equal to zero for all types—i.e., redistribution should only be achieved through income taxes—if the agents' marginal rate of substitution between consumption and savings is homogeneous across ranks r. Coupled with the fact that preferences over savings are independent of rank, this condition is equivalent to $U_{Cr} = 0$ for all r.[13] The following corollary also shows that the converse statement is true:

Corollary 1. *The optimal allocation satisfies* $B_S(r) \gtreqless B_C(r)$ *and the optimal savings wedge is* $\tau_S(r) \gtreqless 0$ *for all* r*, if and only if* $U_{Cr}(r) \lesseqgtr 0$ *for all* r*.*

In other words, the optimal savings tax inherits the sign of $-U_{Cr}$. This insight is already present

[12] Note that the marginal cost can be expressed as: $\frac{\tau_Y}{1-\tau_Y}\frac{1}{A\cdot U_C} = \tau_Y \frac{1}{A\cdot(-U_Y)} = \frac{\tau_Y}{1-\tau_Y}\frac{1+\tau_S}{A\cdot\beta R V'}$.

[13] This is a weaker restriction than the weak separability assumption imposed in Atkinson and Stiglitz (1976), which in our setting would require $U_{CY} = 0$. In particular, if the utility function takes the form $u(C,Y) - v(Y,r) + V(S)$, we have $U_{Cr} = 0$ but $U_{CY} \neq 0$.

in Mirrlees (1976). If the marginal utility is increasing (resp., decreasing) with r, so that higher ranks also have higher (lower) consumption needs, then it is optimal to subsidize (tax) savings at the top of the income distribution. When $U_{Cr} = 0$, 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.[14] When $U_{Cr} < 0$, the planner can screen the more productive ranks—i.e., deter them from mimicking lower ranks—via positive savings taxes by exploiting the fact that their taste for savings over current consumption is stronger than that of lower ranks. Formally, a perturbation that increases consumption, and hence reduces savings, for rank r by $\Delta C(r)$ raises their current utility by $U_C \Delta C(r)$ and changes their information rent by $U_{Cr} \Delta C(r)$. Thus, U_{Cr}/U_C measures the ratio of the change in information rents to the increase in utility that comes with a reduction in savings. If such a perturbation reduces information rents ($U_{Cr} < 0$), then it allows the planner to increase the static redistribution from higher towards lower ranks, thus leading to a rationale for taxing savings.[15] Hellwig (2021) establishes the same result in a general multi-period dynamic Mirrlees model and shows that the alignment of ability with consumption needs offers a new rationale for taxing or subsidizing savings.

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 τ_Y, τ_S in terms of sufficient statistics that can be observed empirically. In particular, Theorem 1 and Corollary 1 imply that the needs-based and ability-based complementarity variables U_{Cr}/U_C and U_{Yr}/U_Y play a critical role in our analysis. Our key identification result shows that the distributions of earnings and consumption, along with standard behavioral elasticities, are sufficient to separately identify these variables in the data.[16]

3.1 Identification Lemma

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, respec-

[14] It is straightforward to check from equation (5) that, when $U_{Cr} = 0$, $B_S(r) = B_C(r)$ for all r if and only if $1/(\beta R V'(S(r))) = 1/U_C(r)$, or $\tau_S(r) = 0$, for all r.

[15] As we show in Section 5, the intuition and the result generalizes 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 $\mu_C(r, r')$ and $\mu_S(r, r')$ are the same.

[16] Lemma 1 also suggests that wealth inequality plays a role in the identification. However, as we discuss in Section 3.3, the distribution of savings in our model is pinned down as a residual from those of earnings and consumption.

tively:

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

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 $\zeta_C(r), \zeta_Y(r), \zeta_S(r), \zeta_{CY}(r)$ by

$$\zeta_C(r) \equiv -\frac{\partial \ln U_C(r)}{\partial \ln C} = \frac{-C(r) U_{CC}(r)}{U_C(r)}, \quad \zeta_Y(r) \equiv \frac{\partial \ln(-U_Y(r))}{\partial \ln Y} = \frac{Y(r) U_{YY}(r)}{U_Y(r)},$$

and

$$\zeta_S(r) \equiv -\frac{\partial \ln V'(S(r))}{\partial \ln S} = -\frac{S(r) V''(S(r))}{V'(S(r))}, \quad \zeta_{CY}(r) \equiv \frac{\partial \ln U_C(r)}{\partial \ln Y} = \frac{Y(r) U_{CY}(r)}{U_C(r)}.$$

These four elasticity parameters all have direct empirical counterparts (see Section 4.1). The variables ζ_C and ζ_S are the coefficients of relative risk aversion in periods 0 and 1, respectively, and ζ_{CY} is a coefficient of complementarity between consumption and labor supply. The variable ζ_Y is an inverse elasticity of labor supply. Specifically, denote by ζ_Y^H and ζ_Y^M the Hicksian (compensated) and Marshallian (uncompensated) elasticities, and by $\zeta_Y^I \equiv \zeta_Y^H - \zeta_Y^M$ the income effect parameter. If the consumption share s_C converges to zero for top earners, which we argue below is empirically relevant, then we get $\zeta_Y = \left(1-\zeta_Y^I\right)/\zeta_Y^H$ and $\zeta_S = \zeta_Y^I/\zeta_Y^H$ as $r \to 1$.[17] More generally, we derive the map between ζ_Y^H, ζ_Y^I and $\zeta_C, \zeta_Y, \zeta_S, \zeta_{CY}$ in the Appendix.

Lemma 1 (**Identification**). *The variables U_{Cr}/U_C and U_{Yr}/U_Y can be expressed in terms of sufficient statistics as:*

$$\frac{U_{Cr}(r)}{U_C(r)} = \frac{\zeta_C(r)}{\rho_C(r)} - \frac{\zeta_S(r)}{\rho_S(r)} - \frac{\zeta_{CY}(r)}{\rho_Y(r)} - \frac{\tau_S'(r)}{1+\tau_S(r)} \tag{7}$$

and

$$\frac{U_{Yr}(r)}{U_Y(r)} = -\frac{\zeta_Y(r)}{\rho_Y(r)} - \frac{\zeta_S(r)}{\rho_S(r)} + \frac{s_C(r)\zeta_{CY}(r)}{\rho_C(r)} - \frac{\tau_Y'(r)}{1-\tau_Y(r)} - \frac{\tau_S'(r)}{1+\tau_S(r)}. \tag{8}$$

This result, which generalizes Lemma 1 in Saez (2001) to our dynamic economy, is obtained by differentiating the first-order conditions of the individual's problem $1-\tau_Y = -U_Y/U_C$ and $1+\tau_S = \beta R V'/U_C$ with respect to the rank r. It shows that empirically observable parameters—standard elasticities, Pareto coefficients, and measures of tax progressivity—together pin down the weights on the consumption- and needs-based redistribution motives. This result does not rely on

[17] Inverting these relationships leads to $\zeta_Y^H = 1/(\zeta_S + \zeta_Y)$ and $\zeta_Y^I = \zeta_S/(\zeta_S + \zeta_Y)$.

specific functional form assumptions for preferences: The "data" implicitly informs us about the underlying correlation structure between ranks and marginal utilities.[18]

To understand the key insight of Lemma 1, focus on top earners, i.e. $r \to 1$, for whom the Pareto coefficients ρ_C, ρ_Y, ρ_S converge to constant and the progressivity terms $\frac{\tau'_Y}{1-\tau_Y}$ and $\frac{\tau'_S}{1+\tau_S}$ converge to zero. Suppose moreover that the risk-aversion parameters in periods 1 and 2 are equal, $\zeta_C = \zeta_S$, and that the complementarity coefficient ζ_{CY} is small relative to risk aversion, as is the case empirically. Equation (7) then implies that the sign of U_{Cr}/U_C is determined by the relative thickness of the Pareto tails ρ_C vs. ρ_S. Specifically, U_{Cr} is negative—so that capital should be taxed—if and only if $\rho_C > \rho_S$, i.e., iff consumption is strictly more evenly distributed than wealth at the top. Intuitively, the relative thickness of the tails of consumption and wealth—or, more generally the ratios of elasticities and Pareto coefficients ζ_C/ρ_C, ζ_S/ρ_S—reflect how the taste for current consumption relative to savings varies along the ability distribution. In particular, observing that $\rho_C > \rho_S$ implies that the consumption share s_C converges to 0 as $r \to 1$; that is, top earners spend a vanishing fraction of their labor earnings on current consumption. We view these sufficient statistics for optimal labor and capital taxes, based on the direct comparison of two Pareto tails, as particularly direct and transparent.[19]

More generally, these elasticities and Pareto coefficients not only determine the signs, but also pin down the values of U_{Cr}/U_C and U_{Yr}/U_Y. The separate identification of these two variables in the data is critical and stems from the dynamic structure of our model. In the static framework, one could match the observed income distribution by positing a complementarity between ranks and either the marginal utility of consumption ($U_{Cr} \neq 0$), or the marginal disutility of effort ($U_{Yr} \neq 0$), or a combination of both. The estimation of static optimal income tax formulas would then require choosing functional form restrictions on preferences, e.g. $U_{Cr} = 0$. But such a choice is far from innocuous, as optimal income taxes depend *independently* on both redistribution motives

[18] By contrast, 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 the values of the calibrated parameters imply that the right-hand side of (7) 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 assumptions that preferences are GHH, $U = u\left(g\left(C\right) - v\left(Y/\theta\left(r\right)\right)\right)$ for some convave 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 restriction (7), or impose that $\rho_C = \rho_Y$, which as we discuss below is not consistent with empirical evidence.

[19] Saez (2002) and Ferey, Lockwood, and Taubinsky (2021) emphasize different sufficient statistics, namely the cross-sectional variation of savings with income net of the causal effect of income on savings. These two representations provide equivalent but alternative ways of capturing the correlation between ranks and tastes for savings from the data.

U_{Yr}, U_{Cr}.[20] Introducing a second behavioral margin in the model—consumption-savings in addition to consumption-leisure—is what allows us to identify both of these complementarity variables and, therefore, correctly evaluate optimal income taxes.

3.2 Optimal Top Tax Rates

We now proceed to expressing 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 earnings, consumption and savings have unbounded support and upper Pareto tails with coefficients* ρ_Y, ρ_C, ρ_S*, respectively. In addition, the elasticities* $\zeta_C, \zeta_S, \zeta_Y, \zeta_{CY}$ *and the parameter* s_C *converge to finite limits* as $r \to 1$.

Lemma 1, along with Assumption 2, allows us to derive empirical counterparts for the marginal benefits terms B_C, B_Y, B_S that appear in the optimal tax formulas of Theorem 1 (equation (5)). We find[21]

$$\lim_{r\to 1} B_C(r) = \left[1 - \frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y}\right]^{-1} \tag{9}$$

and

$$\lim_{r\to 1} B_Y(r) = \left[1 + \frac{\zeta_Y}{\rho_Y} - \frac{s_C\zeta_{CY}}{\rho_C}\right]^{-1} \tag{10}$$

and

$$\lim_{r\to 1} B_S(r) = \left[1 - \frac{\zeta_S}{\rho_S}\right]^{-1}. \tag{11}$$

Abstracting for now from complementarities, these expressions show that there is a natural mapping between consumption (resp., earnings, 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 (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 earnings inequality

[20] It is straightforward to see that U_{Yr} and U_{Cr} matter independently for optimal taxes in the static setting. Suppose that the planner is not Rawlsian (see Section 5). Then the marginal utility of consumption determines the marginal social welfare weights at each income level below the top. For instance, with GHH preferences, the concavity of the outer utility function affects the level of optimal taxes but is irrelevant for the difference $U_{Cr}/U_C - U_{Yr}/U_Y$ that the static model allows us to identify in the data.

[21] As 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.

$1/\rho_Y$; intuitively, high earnings inequality indicates that top earners are hard-working and have relatively little leisure. Together with Theorem 1, this implies that, *ceteris paribus*, high earnings and consumption inequality both lead to high optimal tax rates on top incomes τ_Y, while high wealth inequality but low consumption inequality lead to high optimal tax rates of top savings τ_S. Finally, the complementarity between consumption and earnings ζ_{CY} lowers (raises) the marginal benefits of redistributing consumption (leisure). As a result, a higher degree of complementarity unambiguously lowers the optimal top income tax rate τ_Y, and raises the optimal top savings tax rate τ_S. These are familiar results: When preferences are non-separable, it is optimal to tax less heavily the goods that are complementary to labor (Corlett and Hague (1953)).

Expressions (9), (10) and (11) immediately lead to the following theorem, which is the second main result of this paper:

Theorem 2. *Suppose that the optimal allocation satisfies Assumption 2. Then the optimal labor wedge on top income earners* $\overline{\tau}_Y \equiv \lim_{r\to 1} \tau_Y(r)$ *satisfies*

$$1-\overline{\tau}_Y = \frac{1-\zeta_C/\rho_C+\zeta_{CY}/\rho_Y}{1+\zeta_Y/\rho_Y-s_C\zeta_{CY}/\rho_C} \tag{12}$$

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

$$1+\overline{\tau}_S = \frac{1-\zeta_C/\rho_C+\zeta_{CY}/\rho_Y}{1-\zeta_S/\rho_S}, \tag{13}$$

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

Equation (12) provides a very simple generalization of the standard top income tax rate formula of Saez (2001) to a dynamic environment, and equation (13) provides an analogous sufficient statistics formula for savings taxes. The static optimum derived by Saez (2001) (equation (8)) is expressed in terms of the Hicksian and Marshallian elasticities of labor supply ζ_Y^H, ζ_Y^M as:

$$\overline{\tau}_Y^{Saez} = \frac{1}{1-(\zeta_Y^H-\zeta_Y^M)+\rho_Y\zeta_Y^H}, \tag{14}$$

where the map between ζ_Y^H, ζ_Y^M and $\zeta_C, \zeta_Y, \zeta_S, \zeta_{CY}$ is given in the Appendix. In constrast to the static formula (14), the optimal income tax rate in the dynamic environment (12) depends explicitly on the Pareto tail coefficient of *consumption* in addition to that of labor income. Recall that 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. For the same reason, this variable is also implictly present in equation (14). However, in a static economy, 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 (14) 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 this formula in terms of ρ_C and estimate it using consumption data. As soon as we break the equivalence between consumption and after-tax income by adding a consumption-savings margin to the model, it becomes clear that both coefficients ρ_Y and ρ_C matter independently for the level of optimal labor income taxes. 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 redistributional arbitrage representation gives a transparent interpretation of this result.

3.3 A Tale of Three Tails

The budget constraint in our model imposes that earnings are 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. In particular, this restriction implies that one cannot choose all three Pareto coefficients freely from the data. This is the analogue of the condition $\rho_Y = \rho_C$ in the static setting, which follows from the fact that consumption is equal to after-tax income. Our model is thus consistent with the following three scenarii:[22]

1. $\rho_Y = \rho_C < \rho_S$, so that savings are strictly more evenly distributed than earnings and consumption. Equivalently, the budget share of consumption s_C converges to 1 for top earners.

2. $\rho_Y = \rho_S < \rho_C$, so that consumption is strictly more evenly distributed than earnings 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 earnings, 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). Conversely, variations of our model that would

[22] If $\rho_C < \rho_Y$ (resp., $\rho_S < \rho_Y$), then the consumption (resp., savings) shares of after-tax earnings 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 consumption and savings shares must both converge to 0, which violates the inter-temporal budget constraint. Thus, $\min\{\rho_C, \rho_S\} = \rho_Y$. Differentiating 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.

be consistent with higher wealth inequality than income inequality, such as Scheuer and Slemrod (2021), implicitly require that consumption be as unequally distributed as savings, $\rho_C = \rho_S < \rho_Y$. We analyze such a model in Section 5.[23]

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 $s_C = 1$. Using the map between ζ_Y, ζ_C and the Hicksian and Marshallian elasticities ζ_Y^H, ζ_Y^M,[24] it is straightforward to show that formula (12) reduces to the static optimum (14). 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, 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 earnings on current consumption. Unfortunately, as we argue below, this case is not the empirically relevant one.

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 $s_C = 0$. Using the map between ζ_Y, ζ_C and ζ_Y^H, ζ_Y^M, we obtain that, in this case, the optimal top labor income tax rate $\overline{\tau}_Y$ no longer coincides with the static optimum $\overline{\tau}_Y^{Saez}$ given by equation (14), unless the relationship $\zeta_S/\rho_Y = \zeta_C/\rho_C - \zeta_{CY}/\rho_Y$ holds—that is, unless the Atkinson-Stiglitz theorem applies, so that the optimum savings tax rate $\overline{\tau}_S$ is equal to zero. More generally, we have

$$1 - \overline{\tau}_Y^{Saez} = \frac{1 - \overline{\tau}_Y}{1 + \overline{\tau}_S} = \frac{1 - \zeta_S/\rho_S}{1 + \zeta_Y/\rho_Y}. \tag{15}$$

This result implies that the static optimum $\overline{\tau}_Y^{Saez}$ overstates the correct optimum $\overline{\tau}_Y$ whenever the optimal savings tax rate $\overline{\tau}_S$ is strictly positive. Conversely, the static framework underestimates the optimum top labor income tax rate if savings are subsidized. Intuitively, when the consumption share converges to zero at the top, so that top earners save most of their earnings, the optimal allocation is determined by a static trade-off between the two variables Y and S. Thus, the static optimum $\overline{\tau}_Y^{Saez}$ now characterizes the optimal wedge between earnings and savings, which is the

[23] Picking the three parameters freely from the data and allowing for $\rho_S < \rho_Y < \rho_C$ would require introducing an additional source of heterogeneity, which can be rates of return or endowments. While this may certainly be empirically plausible, incorporating such heterogeneity leads to complex multidimensional screening issues that the literature has not yet been able to fully address; for recent explorations of these questions, see e.g. Rothschild and Scheuer (2014), Spiritus, Lehmann, Renes, and Zoutman (2021), and Boerma, Tsyvinski, and Zimin (2022).

[24] In Case 1, we have $\zeta_Y' = \left(1 - \zeta_Y^I\right)/\zeta_Y^H$ and $\zeta_C' = \zeta_Y^I/\zeta_Y^H$ where $\zeta_Y' \equiv \zeta_Y - \zeta_{CY}$ and $\zeta_C' \equiv \zeta_C - \zeta_{CY}$, or conversely, $\zeta_Y^H = 1/(\zeta_Y' + \zeta_C')$ and $\zeta_Y^I = \zeta_C'/(\zeta_Y' + \zeta_C')$. Hence, $1 - \overline{\tau}_Y^{Saez} = \frac{1-\zeta_C'/\rho_Y}{1+\zeta_Y'/\rho_Y}$.

combination of the labor and savings wedges $\overline{\tau}_Y$ and $\overline{\tau}_S$ (equation (15)). By Corollary 1, whenever the optimal tax on savings is strictly positive, i.e., if preferences are such that $U_{Cr} < 0$, then the optimal tax on labor income is strictly lower than that predicted by the formula of Saez (2001).[25]

Case 3. Income, Consumption, and Savings have Identical Tails. Suppose finally that the distributions of earnings, 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 (12) generally differs from the static optimum (14). The dynamic adjustments can only be neglected when the first-period utility is quasilinear in consumption, so that $U_{CC} = U_{CY} = 0$.[26] 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 inccome tax rate, and the standard formula of Saez (2001) ceases to apply.

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: ρ_Y, ρ_C, 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 equal to 1.6 and that on wealth is equal to 1.4 (Diamond and Saez (2011)). Since our model imposes $\rho_Y = \rho_S$, we take their common value to be equal to 1.5. Turning to measures of consumption inequality, Toda and Walsh (2015) argue that consumption is also Pareto distributed at the top, and they estimate an upper tail coefficient of $\rho_C = 3.65$, i.e., $\rho_Y/\rho_C = 0.41$. Straub (2019) finds that the income elasticity of consumption is equal to 0.7, which pins down the ratio of Pareto coefficients of earnings and consumption, $\rho_Y/\rho_C = \frac{C'/C}{Y'/Y} = 0.7$ or $\rho_C = 2.14$. Both of these estimates suggest that consumption has a substantially thinner tail than income, so that $s_C \to 0$ as $r \to 1$: Top earners save most of their earnings. Below we evaluate our optimal tax formulas for $\rho_Y = \rho_S = 1.5$ and $\rho_Y/\rho_C \in \{0.4, 0.6, 0.75\}$.

[25] The fact that the optimal labor tax rate is smaller (resp., larger) than the static optimum if capital is taxed (resp., subsidized) can also be understood by considering a marginal reduction in the labor income tax rate, starting from the static optimum. In addition to the standard effects on labor supply, this tax reform now also raises savings, which in turn raises (resp., lowers) government revenue if savings are taxed (resp., subsidized). The static model fails to account for this fiscal externality, hence overstates (resp., understates) the optimal tax rate on labor income.

[26] Indeed, 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.

Elasticities: ζ_Y, ζ_S. Recall that 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 and average tax rates. 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 earnings 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 $\zeta_Y = \left(1 - \zeta_Y^I\right)/\zeta_Y^H = 2.25$ and $\zeta_S = \zeta_Y^I/\zeta_Y^H = 0.75$, a reasonable value for the risk-aversion parameter 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 $\zeta_Y^I = 1/3$ (so that $\zeta_S = 1$).

Elasticities: ζ_C, ζ_{CY}. Because the combined wedge on income and savings is equal to the static wedge (equation (15)), 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 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 $(\rho_Y, \zeta_Y^H, \zeta_Y^I, \zeta_C, \zeta_{CY}) = (\frac{3}{2}, \frac{1}{3}, \frac{1}{4}, \frac{3}{4}, 0)$ and three values of the Pareto coefficient on consumption ρ_C. We also report the

Table 1: Optimal Taxes in Case 2

	$\rho_Y/\rho_C = 0.4$		$\rho_Y/\rho_C = 0.6$		$\rho_Y/\rho_C = 0.75$		
	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y$	$\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$	$\overline{\tau}_Y^{Saez}$
Baseline	68%	38%	72%	29%	75%	20%	80%
$\zeta_Y^H = 0.5$	60%	17%	65%	5%	69%	−7%	67%
$\zeta_Y^I = 1/3$	66%	58%	70%	52%	73%	47%	86%
$\zeta_C = 1.25$	73%	25%	80%	0%	85%	−33%	80%
$\zeta_{CY} = 0.15$	64%	44%	68%	38%	71%	31%	80%

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 earnings, $\bar{\tau}_S$ 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 $\frac{\overline{\tau}_S}{1+\overline{\tau}_S}$, 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 $\frac{\overline{\tau}_S}{1+\overline{\tau}_S} = 40\%$, 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 *a priori*: Our calibration of the elasticities and Pareto tails implicitly determines the underlying structure of preferences. 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 $\overline{\tau}_Y^{Saez}$ between savings and income taxes $\overline{\tau}_Y, \overline{\tau}_S$ is pinned down by the ratios of elasticities (or risk-aversion parameters) and Pareto coefficients ζ_C/ρ_C, ζ_Y/ρ_Y, ζ_S/ρ_S that respectively drive the marginal benefits of redistributing consumption, leisure, and savings. For low values of the first-period risk aversion or a very thin consumption tail ($\rho_Y/\rho_C = 0.4$), 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 grows 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 $\overline{\tau}_Y$ exceeds $\overline{\tau}_Y^{Saez}$ and savings are subsidized, $\overline{\tau}_S < 0$. 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 $\overline{\tau}_Y \geq 1 - B_Y = \frac{1}{1+\rho_Y/\zeta_Y} = 60\%$ and $\overline{\tau}_S \leq B_S - 1 = \frac{1}{\rho_S/\zeta_S - 1}$ so $\frac{\overline{\tau}_S}{1+\overline{\tau}_S} \leq 52\%$ 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 earnings 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 (12) and (13) 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 increases the ratio of tail coefficients from $\rho_Y/\rho_C = 0.40$ to $\rho_Y/\tilde{\rho}_C = 0.25$. 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 $\rho_S/\rho_C = 0.4$ (resp., 0.75) imply that savings should be taxed unless the first-period risk-aversion coefficient for top earners ζ_C is larger than $\frac{\rho_C}{\rho_S}\zeta_S = 1.875$ (resp., 1). With the complementarity correction, we have $\rho_S/\tilde{\rho}_C = 0.25$ (resp., 0.6), so risk aversion ζ_C would need to exceed 3 (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 correction 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

In this last section, we extend our baseline model, our redistributional arbitrage formulas, and their sufficient-statistic representations, to a general environment with a one-dimensional preference type, arbitrary preferences and commodities.

5.1 General Preferences and Multiple Commodities

We have assumed that preferences were additively time-separable, so that the benefits of "savings" were independent of rank r, "consumption" and "earnings". As we discuss formally below, it is straightforward to generalize Theorem 1 to general preferences of the form $U(C, S, Y; r)$: A further incentive-adjustment $\mu_S(r, r')$ to the inverse marginal utilities will appear in the computation of B_S in equation (5). In addition, 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 earnings. With unrestricted preferences, the analysis will require estimates for two additional preference elasticities to account for the complementarity of consumption and earnings with savings. Finally, 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.

Formally, suppose that agents' preferences are defined as $U(\boldsymbol{X}; r)$, where $\boldsymbol{X}$ 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 $\boldsymbol{X}$ is $C(\boldsymbol{X})$, and we let $p_n = \frac{\partial C}{\partial x_n}$ denote the "price" of good n. The planner's problem reads

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

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.

To characterize the optimal allocation, 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 $r' \geq r$. Any such perturbation around the optimal allocation must keep the planner's objective

function unchanged, for otherwise the perturbation or its negative would lead to a strict welfare improvement. The unique perturbation $\{\delta x_n(r'), \delta x_m(r')\}$ that satisfies these four requirements is given by

$$\delta x_k(r') = \frac{1}{U_k(r')} \exp\left(\int_r^{r'} \frac{U_{kr}}{U_k} dr''\right) \Delta \equiv \frac{1}{U_k(r')} \mu_k(r, r') \Delta, \quad \text{for } k \in \{n, m\}$$

for small positive Δ. This perturbation impacts the planner's objective by

$$\left\{ \int_r^1 \hat{\omega}(r') \left[\mu_n(r, r') - \mu_m(r, r')\right] dr' - p_n \int_r^1 \frac{1}{U_n(r')} \mu_n(r, r') dr' + p_m \int_r^1 \frac{1}{U_m(r')} \mu_m(r, r') dr' \right\} \Delta$$

where $\hat{\omega}(r) = \omega(r) G'(U(r))$ represents the marginal welfare weight on rank r. Setting this expression equal to 0 and rearranging terms yields

$$\frac{U_m(r)}{U_n(r)} \frac{p_n}{p_m} \equiv 1 - \tau_{m,n}(r) = \frac{B_m(r)}{B_n(r)}, \tag{16}$$

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

$$B_k(r) = \mathbb{E}\left[\frac{U_k(r)}{U_k(r')} \mu_k(r, r') \,|r' \geq r\right] \left(1 - \frac{\mathbb{E}\left[\hat{\omega}(r') \mu_k(r, r') \,|r' \geq r\right]}{p_k \mathbb{E}\left[(U_k(r'))^{-1} \mu_k(r, r') \,|r' \geq r\right]}\right) \tag{17}$$

represents the marginal benefits of reducing the consumption of commodity k for ranks above r while preserving incentive-compatibility. In other words, equation (16) implies that the planner's marginal benefits of redistribution must be equalized between any pair of goods. Note that the Inada conditions ensure that the ratio in brackets in (17) vanishes at the top of the type distribution: We thus recover the Rawlsian representation of optimal wedges of Theorem 1.

Next, we rewrite $M_n(r') \equiv \frac{1}{U_n(r')} \mu_n(r, r')$ in terms of observables. Taking derivatives w.r.t. r yields

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

or $M_n(r') \sim \prod_{k=1}^N x_k(r')^{\zeta_{nk}}$ as $r \to 1$, where $\zeta_{nk} = \frac{U_{nk}}{U_n} x_k(r')$ is assumed to converge to a constant. It then follows that $\mathbb{E}\left[\frac{U_n(r)}{U_n(r')} \mu_n(r, r') \,|r' \geq r\right] = \left[1 - \sum_{k=1}^N \frac{\zeta_{nk}}{\rho_k}\right]^{-1}$ where ρ_k denotes the Pareto coefficient on good k. Hence, we obtain, for top earners,

$$1 - \tau_{m,n}(r) = \frac{B_m(r)}{B_n(r)} = \frac{\left[1 - \sum_{k=1}^N \frac{\zeta_{mk}}{\rho_k}\right]^{-1}}{\left[1 - \sum_{k=1}^N \frac{\zeta_{nk}}{\rho_k}\right]^{-1}}. \tag{18}$$

Equation (18) 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); and (ii) the full matrix of income and substitution effects of all commodities which is summarized by $\{\zeta_{nk}\}_{1\leq k\leq N}$.

To conclude this section, as we discussed in the context of Corollary 1, our model reveals a potential rationale for non-uniform commodity taxation for redistributive objectives, 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.

5.2 Heterogeneous Initial Capital

Scheuer and Slemrod (2021) derive a formula for the optimal top capital tax rate in a model that is similar to, but slightly different than our baseline environment. In their framework, as in ours, agents are indexed by a one-dimensional productivity type r. Preferences satisfy the restrictions of Atkinson and Stiglitz (1976); namely, they are separable between consumption and earnings and homogeneous across consumers. However, agents receive an exogenous endowment $Z(r)$ that is perfectly correlated with productivity. This alternative framework makes it possible to break the equality between the Pareto coefficients on income and wealth that the budget constraint in our baseline model imposes. In this section, we study in more detail a special case of our general environment of Section 5.1 that nests both our baseline model and that of Scheuer and Slemrod

(2021). The proofs are in the Appendix.

We show that the characterization of optimal labor and savings wedges of Theorem 1 is the same as in our baseline model, except that we must adjust the definition of the incentive-adjustments $\mu_C(r, r')$ and $\mu_Y(r, r')$. To express the top tax rates in terms of sufficient statistics, we now need to keep track of one additional key parameter, which depends on the consumption to endowment ratio at the top, as well as the respective Pareto coefficients: $s_Z = \frac{\rho_C}{\rho_Z} \lim_{r\to 1} \frac{Z(r)}{C(r)}$. Under Assumption 2, the marginal benefits of redistributing earnings, consumption, and savings at the top are then respectively given, as $r \to 1$, by

$$B_Y(r) \to \frac{1}{1 + \frac{\zeta_Y}{\rho_Y} - (1 - s_Z)\frac{s_C \zeta_{CY}}{\rho_C}}, \quad B_C(r) \to \frac{1}{1 - (1 - s_Z)\frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y}}, \quad B_S(r) \to \frac{1}{1 - \frac{\zeta_S}{\rho_S}}.$$

Thus, the adjustment of marginal benefits for rank-dependent endowments reduces both B_Y and B_C, resulting in strictly higher savings wedges, and a higher combined wedge. We can then distinguish four scenarios, based on the relative values of the Pareto coefficients of the income, consumption, and wealth distributions.

First, suppose that initial wealth endowments have a thinner tail than consumption, so that $\rho_C < \rho_Z$ and $s_Z \to 0$ as $r \to 1$. Then the results of our baseline model continue to hold. Intuitively, if endowments have a strictly thinner tail than consumption, then it must be that they simply do not matter at the top of the distribution—i.e., top earners get all of their income from labor—and we recover the analysis from our baseline framework. Second, take the opposite case where consumption has a thinner tail than wealth endowments, so that $\rho_C > \rho_Z$ and $s_Z \to \infty$ as $r \to 1$. It follows that B_C and B_Y both converge to 0, with $\lim \frac{B_Y(r)}{B_C(r)} > 1$ and $\lim \frac{B_S(r)}{B_C(r)} = \infty$. Hence $\overline{\tau}_Y < 0$ (earnings should be subsidized) and $\overline{\tau}_S \to \infty$ (savings should be taxed at 100%). In the limit where $s_C \zeta_{CY} \to 0$ (in particular, if preferences are separable between consumption and earnings) the earnings subsidy becomes infinitely large.

Third, suppose that endowments, consumption and earnings all have the same tail coefficient, so that $\rho_C = \rho_Z = \rho_Y$ and $0 < s_C, s_Z < \infty$. Optimal taxes are then given by $\overline{\tau}_Y = 1 - \frac{B_Y}{B_C}$ and $\overline{\tau}_S = \frac{B_S}{B_C} - 1$, where B_Y, B_C, B_S are given by the above expressions. The combined wedge in income and savings $\frac{1-\overline{\tau}_Y}{1+\overline{\tau}_S}$ equals the optimal static wedge when $s_Z = 1$, i.e., when top income earners consume just their endowment and save all their labor income; it grows larger when $s_Z > 1$ (top income earners consume less than their endowment), and smaller when $s_Z < 1$ (top income earners consume less than their endowment). Fourth, suppose that endowments and consumption have the same tail coefficient, while income has a strictly thinner tail, so that $\rho_C = \rho_Z < \rho_Y$ and

$s_C \to \infty$ with $s_Z < \infty$. If $\zeta_{CY} > 0$ and $s_Z > 1$ (top earners save a positive fraction of their income), it follows that both $B_Y \to 0$ and $B_C \to 0$. Hence, $\overline{\tau}_Y = 0$ while the optimal savings tax converges to 100%.

To summarize: With $\zeta_{CY} > 0$, the scenario with $\rho_C > \rho_Z$ prescribes 100% savings coupled with an earnings subsidy at the top, whereas the opposite scenario reduces to the one in our baseline model. The equal tails scenario $\rho_C = \rho_Z$ allows for two subcases: Either all tails are equal with interior (finite) solutions for optimal taxes, or the labor income tail is strictly thinner than the consumption/endowment tail, and then we again recover no distortion at the top for labor and full expropriation of marginal savings. Scheuer and Slemrod (2021) consider the case where preferences are separable ($\zeta_{CY} = 0$) and endowments and consumption have an equal tail ($\rho_Z = \rho_C$). This leads to an interior solution for both labor and savings taxes, which depends on the top earners' budget shares of consumption and labor earnings. However, when earnings have a thinner tail than endowments or wealth ($\rho_Y > \rho_C = \rho_Z$) this result is "knife-edge": Even small degrees of preference complementarity lead to discontinuous changes in optimal income and savings taxes.

5.3 Further Extensions

We conclude by briefly discussing other extensions that are outside the scope of the present paper.

Richer Dynamics. A natural extension consists of following the lead of dynamic Mirrlees models (Golosov, Kocherlakota, and Tsyvinski (2003), Farhi and Werning (2013), and Golosov, Troshkin, and Tsyvinski (2016)) and allowing for stochastic evolution of types over multiple periods. Hellwig (2021) develops the implications of such a model for labor and savings wedges with non-separable preferences, and derives conditions, linked to the persistence of types and information rents, under which the optimal labor wedge can still be represented as $1 - \tau_Y\left(r^t\right) = \frac{B_Y\left(r^t\right)}{B_C\left(r^t\right)}$, where r^t represents a t-period sequence of type realizations, leading to an analogous sufficient-statistic representation to that of Theorem 2. The interpretation of this formula as a redistributional arbitrage is the same as here, but in contrast to the present static setting, the current marginal benefits—and hence, in the sufficient-statistic representation, the Pareto coefficients on each good—are now based on distributions of earnings and consumption growth conditional on the prior sequence of types, or equivalently, the prior earnings history.

Multi-Dimensional Types. 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, earnings, 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, preliminary results in Hellwig (2022) for a multi-good monopolist problem suggest that core ideas from the present analysis can be generalized, in particular the representation of local incentive compatibility through incentive-adjusted probability measures, the characterization of optimal relative price distortions through an arbitrage of information rents, and a general representation of optimal distortions that generalizes the B_j/B_k-formula presented here. These preliminary results suggest that there is scope to generalize the analysis, and that the core idea of redistributional arbitrage across different dimensions of the commodity space also applies in multi-dimensional type spaces.

6 Conclusion

We developed 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 redistributional arbitrage. As our main result, we derived a simple new formula for optimal tax distortions based on redistributional arbitrage. We show how to infer the respective marginal benefits of redistibution 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 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. Redistributional arbitrage formalizes how these trade-

offs are resolved by optimal tax policies. In practice, many policy makers probably develop an intuitive understanding for redistributional arbitrage, when determining what policies are popular with their voters and matter for their welfare. In fact, the Roman emperors are perhaps the first rulers on record to perform redistributional 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![27]

[27] To 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.

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.

Boerma, Job, Aleh Tsyvinski, and Alexander P Zimin (2022). *Bunching and Taxing Multidimensional Skills*. Tech. rep. National Bureau of Economic Research.

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.

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.

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.

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

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

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

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

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

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

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

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 nonlinear income taxation". In: *Journal of Public Economics* 117, pp. 201–210.

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

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

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

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.

Rothschild, Casey and Florian Scheuer (2014). *A theory of income taxation under multidimensional skill heterogeneity*. Tech. rep. National Bureau of Economic Research.

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

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

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 (2019a). "How would a progressive wealth tax work? Evidence from the economics literature". In: *Brookings Institution*.

— (2019b). "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*.

Spiritus, Kevin, Etienne Lehmann, Sander Renes, and Floris Zoutman (2021). *Optimal Taxation with Multiple Incomes and Types*. Working Paper.

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.

A Appendix: 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:

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

subject to the ex-ante promise-keeping constraint

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

the promise-keeping constraint

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

and the local incentive compatibility constraint

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

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:

$$\begin{aligned}\mathcal{H} &= \{C(r) - Y(r) + R^{-1}S(r) + \lambda(v_0 - W(r))\omega(r)\} \\ &\quad + \psi(r)\{W(r) - U(C(r), Y(r); r) - \beta V(S(r))\} + \phi(r) U_r(C(r), Y(r); r).\end{aligned}$$

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

$$\psi(r) = \frac{1}{U_C(r)} + \phi(r)\frac{U_{Cr}(r)}{U_C(r)} = \frac{1}{-U_Y(r)} + \phi(r)\frac{U_{Yr}(r)}{U_Y(r)} = \frac{1}{\beta R V'(S(r))}.$$

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 $1/(\beta R V'(S(r)))$ of increasing rank r utility through higher consumption, lower earnings or higher savings, and a second term that measures how such a consumption or earnings increase affects $U_r(r)$ and thereby tightens or relaxes the local incentive compatibility constraint at r by $\frac{U_{Cr}(r)}{U_C(r)}$ or $\frac{U_{Yr}(r)}{U_Y(r)}$. The latter is weighted by the multiplier $\phi(r)$ and added to the former; it is missing from the first-order condition for savings since preferences are separable in savings.

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 any of the three first-order conditions:

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

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

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

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

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

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 earnings 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 yields $\lambda = \frac{\mathbb{E}[m_Y(-U_Y^{-1})]}{\mathbb{E}[m_Y\,\omega]}$ to get

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

We obtain the second intra-temporal optimality condition ("ABC") $\tau_Y(r) = A(r)\cdot B_Y(r)$.

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

$$\phi'(r) = \lambda\omega(r) - \frac{1}{\beta RV'(S(r))},$$

which can be integrated and solved along the same lines as above to find

$$\frac{\phi(r)}{1-r} = \mathbb{E}\left[\frac{1}{\beta RV'(S(r'))}|r' \geq r\right] - \mathbb{E}\left[\frac{1}{\beta RV'(S(r))}\right]\mathbb{E}\left[\omega(r')\,|r' \geq r\right] = \frac{1}{\beta RV'(S(r))}B_S(r)$$

with $\lambda = \mathbb{E}\left[1/\left(\beta RV'(S(r))\right)\right]$. 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 $\lim_{r\to 1}\tau_Y(r) < 1$, then optimal allocations satisfy the Inada condition $\lim_{r\to 1}U_C(r) = \lim_{r\to 1}(-U_Y(r)) = \lim_{r\to 1}V'(S(r)) = 0$. The last equality follows from the Inada condition on V. Moreover, $\lim_{r\to 1}(-U_Y(r)) = \lim_{r\to 1}\frac{B_Y(r)}{B_S(r)}\beta RV'(S(r))$. It is easy to check that $\lim_{r\to 1}B_S(r) \geq 1$ and $\lim_{r\to 1}B_Y(r) \leq 1$, and hence $\lim_{r\to 1}(-U_Y(r)) \leq \lim_{r\to 1}\beta RV'(S(r)) = 0$. Finally, $\lim_{r\to 1}U_C(r) = \lim_{r\to 1}\frac{(-U_Y(r))}{1-\tau_Y(r)} = 0$. □

Proof of Corollary 1. We saw in the proof of Theorem 1 that

$$\frac{1}{\beta RV'(S(r))} = \frac{1}{U_C(r)} + \phi(r)\frac{U_{Cr}(r)}{U_C(r)},$$

with $\phi(r) > 0$ for all r. Since $U_{Cr}(r)$ has a constant sign, we get $U_C(r) \lesseqgtr \beta R V'(S(r))$, or $\tau_S(r) \gtreqless 0$ for all r, if and only if $U_{Cr}(r) \lesseqgtr 0$ for all r. □

Proof of Lemma 1. Totally differentiating the two first-order conditions $-\frac{U_Y}{U_C} = 1 - \tau_Y$ and $U_C = \frac{\beta R}{1+\tau_S} V'$ gives respectively

$$\frac{CU_{CY}}{U_Y}\frac{C'(r)}{C(r)} + \frac{YU_{YY}}{U_Y}\frac{Y'(r)}{Y(r)} + \frac{U_{Yr}}{U_Y} + \frac{\tau_Y'}{1-\tau_Y} = \frac{CU_{CC}}{U_C}\frac{C'(r)}{C(r)} + \frac{YU_{CY}}{U_C}\frac{Y'(r)}{Y(r)} + \frac{U_{Cr}}{U_C}$$

and

$$\frac{CU_{CC}}{U_C}\frac{C'(r)}{C(r)} + \frac{YU_{CY}}{U_C}\frac{Y'(r)}{Y(r)} + \frac{U_{Cr}}{U_C} + \frac{\tau_S'}{1+\tau_S} = \frac{SV''(S)}{V'(S)}\frac{S'(r)}{S(r)}.$$

Using the elasticities and Pareto coefficients introduced in the text, 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 two equations can be rewritten as

$$\begin{aligned} -\frac{s_C\zeta_{CY}}{\rho_C} + \frac{\zeta_Y}{\rho_Y} + \frac{U_{Yr}}{U_Y} + \frac{\tau_Y'}{1-\tau_Y} &= -\frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y} + \frac{U_{Cr}}{U_C} \\ -\frac{\zeta_C}{\rho_C} + \frac{\zeta_{CY}}{\rho_Y} + \frac{U_{Cr}}{U_C} + \frac{\tau_S'}{1+\tau_S} &= -\frac{\zeta_S}{\rho_S}. \end{aligned}$$

Equations (7) and (8) follow immediately. □

Proof of Theorem 2. Let $M_C(r) = \frac{1}{U_C(r)} e^{-\int_r^1 \frac{U_{Cr}(r')}{U_C(r')}dr'}$, $M_Y(r) = \frac{1}{-U_Y(r)} e^{-\int_r^1 \frac{U_{Yr}(r')}{U_Y(r')}dr'}$, and $M_S(r) = \frac{1}{\beta R V'(S(r))}$. Differentiating $V'(S(r))$ with respect to r implies

$$\frac{\frac{d}{dr}V'(S(r))}{V'(S(r))} = \frac{V''(S(r))}{V'(S(r))}S'(r) = -\zeta_S(r)\frac{S'(r)}{S(r)},$$

so that $M_S(r) = (\beta R)^{-1} e^{-\int_r^1 \zeta_S(r')\frac{S'(r')}{S(r')}dr'}$. Next, equation (7) leads to

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

where the second equality uses the first-order condition $\frac{1}{U_C(r)}\frac{\beta R V'(S(r))}{1+\tau_S(r)} = 1$. Analogously, equation

(8) leads to

$$\begin{aligned} M_Y(r) &= \frac{1}{-U_Y(r)} e^{\int_r^1 \left\{ -\frac{\frac{d}{dr}V'(S(r'))}{V'(S(r'))} - \frac{\frac{d}{dr}(1-\tau_Y(r'))}{1-\tau_Y(r')} + \frac{\frac{d}{dr}(1+\tau_S(r'))}{1+\tau_S(r')} \right\} dr'} e^{\int_r^1 \left\{ \zeta_Y(r') \frac{Y'(r')}{Y(r')} - s_C(r')\zeta_{CY}(r') \frac{C'(r')}{C(r')} \right\} dr'} \\ &= e^{\int_r^1 \left\{ \zeta_Y(r') \frac{Y'(r')}{Y(r')} - s_C(r')\zeta_{CY}(r') \frac{C'(r')}{C(r')} \right\} dr'}. \end{aligned}$$

Now, it follows from the Inada condition and the condition that elasticities converge to finite limits that

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

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

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

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

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

Along the same lines,

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

and therefore

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

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 fol-

lows automatically that $\lim_{r\to 1}\tau_Y(r) < 1$. To prove the second part of Theorem 2, combine $\lim_{r\to 1} B_S(r) = \frac{1}{1-\zeta_S/\rho_S}$ for $\zeta_S/\rho_S < 1$ with $\lim_{r\to 1} B_C(r) = \frac{1}{1-\frac{\zeta_C}{\rho_C}+\frac{\zeta_{CY}}{\rho_Y}}$ to get

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

This concludes the proof. □

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 that $\beta = R = 1$ and that the tax schedules are locally linear in the top bracket, $T_Y''(Y) = T_S''(S) = 0$. A perturbation $\delta\hat{T}_Y$ with $\delta \in \mathbb{R}$ leads to behavioral responses $(\delta\hat{Y}, \delta\hat{C}, \delta\hat{S})$ that satisfy the perturbed first-order conditions

$$-\frac{U_Y\left[C+\delta\hat{C}, Y+\delta\hat{Y}; r\right]}{U_C\left[C+\delta\hat{C}, Y+\delta\hat{Y}; r\right]} = 1 - T_Y'(Y) - \delta\hat{T}_Y'(Y)$$

and

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

with

$$\hat{C} + (1+T_S'(S))\hat{S} = (1-T_Y'(Y))\hat{Y} - \hat{T}_Y(Y).$$

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

$$\tilde{\zeta}_Y\frac{\hat{Y}}{Y} + \tilde{\zeta}_C\frac{\hat{C}}{C} = -\frac{\hat{T}_Y'}{1-T_Y'}$$

and

$$\tilde{\zeta}_S\frac{\hat{Y}}{Y} - \left[s_S\tilde{\zeta}_C + s_C\tilde{\zeta}_S\right]\frac{\hat{C}}{C} = \zeta_S\frac{\hat{T}_Y}{(1-T_Y')Y}$$

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

$$\frac{\hat{Y}}{Y} = -\zeta_Y^H\frac{\hat{T}_Y'}{1-T_Y'} + \zeta_Y^I\frac{\hat{T}_Y}{(1-T_Y')Y},$$

with

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

In particular, when $s_C \to 1$ and $s_S \to 0$ (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 $s_C \to 0$ and $s_S \to 1$ (Case 2), we have $\zeta_Y^H = \frac{1}{\zeta_Y + \zeta_S}$ and $\zeta_Y^I = \frac{\zeta_S}{\zeta_Y + \zeta_S}$. □

Calibration for Case 3. In case 3, the Pareto coefficients of consumption, earnings, 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

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

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 $\overline{\tau}_S = 0$ and an optimal labor wedge $\overline{\tau}_Y = \overline{\tau}_Y^{Saez} = 80\%$. 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 $\overline{\tau}_Y = 78\%$ and $\overline{\tau}_S = 17\%$. □

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 earnings 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 $W(r) = U(C(r), Y(r); r) + \beta V(S(r))$. The planner's problem is stated as follows:

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

such that

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

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

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

The last equation is the local incentive compatibility constraint. We solve the planner's problem 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, the promise-keeping and local incentive constraints given r, the Hamiltonian for this problem is stated as follows:

$$\begin{aligned} \mathcal{H} &= C(r) - Y(r) + R^{-1} S(r) + \lambda (v_0 - W(r)) \omega(r) \\ &\quad + \psi(r) \{ W(r) - U(C(r), Y(r); r) - \beta V(S(r)) \} \\ &\quad + \phi(r) \{ U_C(C(r), Y(r); r) Z'(r) + U_r(C(r), Y(r); r) \}. \end{aligned}$$

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

$$\begin{aligned} \psi(r) &= \frac{1}{U_C(r)} + \phi(r) \left(\frac{U_{CC}(r)}{U_C(r)} Z'(r) + \frac{U_{Cr}(r)}{U_C(r)} \right) \\ &= \frac{1}{-U_Y(r)} + \phi(r) \left(\frac{U_{CY}(r)}{U_Y(r)} Z'(r) + \frac{U_{Yr}(r)}{U_Y(r)} \right) = \frac{1}{\beta R V'(S(r))}. \end{aligned}$$

Combining the first two FOCs and rearranging terms yields the following static optimality condition:

$$\frac{1}{U_C(r)} \frac{\tau_Y(r)}{1 - \tau_Y(r)} = \frac{1}{-U_Y(r)} - \frac{1}{U_C(r)} = A(r) \phi(r),$$

where

$$A(r) = \frac{U_{Cr}(r)}{U_C(r)} - \frac{U_{Yr}(r)}{U_Y(r)} + \left(\frac{U_{CC}(r)}{U_C(r)} - \frac{U_{CY}(r)}{U_Y(r)} \right) Z'(r).$$

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 any of the three first-order conditions:

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

along with the boundary conditions $\phi(0) = \phi(1) = 0$. Define $\frac{U_{Cr}(r)}{U_C(r)} + \frac{U_{CC}(r)}{U_C(r)}Z'(r) = \frac{m_C'(r)}{m_C(r)}$, or $m_C(r) = e^{-\int_r^1 \left(\frac{U_{Cr}(r')}{U_C(r')} + \frac{U_{CC}(r')}{U_C(r')}Z'(r')\right)dr'}$. 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

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

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

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

Substituting this expression into the static optimality condition then yields the first intra-temporal optimality condition $\frac{\tau_Y(r)}{1-\tau_Y(r)} = A(r)\cdot B_C(r)$. The FOC for earnings yields an analogous ODE,

$$\mu'(r) = \lambda\omega(r) - \frac{1}{-U_Y(r)} - \phi(r)\left(\frac{U_{Yr}(r)}{U_Y(r)} + \frac{U_{CY}(r)}{U_Y(r)}Z'(r)\right).$$

Let $m_Y(r) = e^{-\int_r^1 \left(\frac{U_{Yr}(r')}{U_Y(r')} + \frac{U_{CY}(r')}{U_Y(r')}Z'(r')\right)dr'}$ and apply the same steps as above yields $\lambda = \frac{\mathbb{E}[m_Y(-U_Y^{-1})]}{\mathbb{E}[m_Y\,\omega]}$ to get

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

We obtain the second intra-temporal optimality condition ("ABC") $\tau_Y(r) = A(r)\cdot B_Y(r)$. 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 FOC for savings yields

$$\phi'(r) = \lambda\omega(r) - \frac{1}{\beta RV'(S(r))},$$

which can be integrated and solved along the same lines as above to find

$$\frac{\phi(r)}{1-r} = \mathbb{E}\left[\frac{1}{\beta RV'(S(r'))}|r' \geq r\right] - \mathbb{E}\left[\frac{1}{\beta RV'(S(r))}\right]\mathbb{E}\left[\omega(r')\,|r' \geq r\right] = \frac{1}{\beta RV'(S(r))}B_S(r)$$

with $\lambda = \mathbb{E}\left[1/\left(\beta RV'(S(r))\right)\right]$. Equating this last expression to $\frac{1}{U_C(r)}B_C(r)$ then yields the expression for the savings wedge:

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

It then follows that the characterizations of optimal labor and savings wedges are the same as in our baseline model, except that we must adjust the definition of incentive-adjustments.

Next, we map the incentive-adjustmentss to our observable sufficient statistics. Totally differentiating $U_C(r)$ yields

$$\begin{aligned}\frac{\frac{d}{dr}U_C(r)}{U_C(r)} &= \frac{U_{Cr}(r)}{U_C(r)} + \frac{U_{CC}(r)}{U_C(r)}C'(r) + \frac{U_{CY}(r)}{U_C(r)}Y'(r)\\ &= \frac{U_{Cr}(r)}{U_C(r)} - \zeta_C(r)\frac{C'(r)}{C(r)} + \zeta_{CY}(r)\frac{Y'(r)}{Y(r)}.\end{aligned}$$

It follows that

$$\begin{aligned}e^{-\int_r^1\left(\frac{U_{Cr}(r')}{U_C(r')}+\frac{U_{CC}(r')}{U_C(r')}Z'(r')\right)dr'} &= U_C(r)\,e^{-\int_r^1\left(\zeta_C(r')\left[\frac{C'(r')}{C(r')}-\frac{Z'(r')}{C(r')}\right]-\zeta_{CY}(r')\frac{Y'(r')}{Y(r')}\right)dr'}\\ &= U_C(r)\,e^{-\int_r^1\left(\zeta_C(r')(1-s_Z(r'))\frac{C'(r')}{C(r')}-\zeta_{CY}(r')\frac{Y'(r')}{Y(r')}\right)dr'}\end{aligned}$$

where $s_Z(r) \equiv \frac{Z'(r)}{C'(r)}$. Applying the same steps to $-U_Y(r)$ yields

$$\begin{aligned}\frac{\frac{d}{dr}U_Y(r)}{U_Y(r)} &= \frac{U_{Yr}(r)}{U_Y(r)} + \frac{U_{CY}(r)}{U_Y(r)}C'(r) + \frac{U_{YY}(r)}{U_Y(r)}Y'(r)\\ &= \frac{U_{Yr}(r)}{U_Y(r)} - s_C(r)\,\zeta_{CY}(r)\frac{C'(r)}{C(r)} + \zeta_Y(r)\frac{Y'(r)}{Y(r)},\end{aligned}$$

and hence

$$\begin{aligned} e^{-\int_r^1\left(\frac{U_{Yr}(r')}{U_Y(r')}+\frac{U_{CY}(r')}{U_Y(r')}Z'(r')\right)dr'} &= -U_Y(r)\,e^{\int_r^1\left(-s_C(r')\zeta_{CY}(r')\left[\frac{C'(r')}{C(r')}-\frac{Z'(r')}{C(r')}\right]+\zeta_Y(r')\frac{Y'(r')}{Y(r')}\right)dr'} \\ &= -U_Y(r)\,e^{\int_r^1\left(-s_C(r')\zeta_{CY}(r')(1-s_Z(r'))\frac{C'(r')}{C(r')}+\zeta_Y(r')\frac{Y'(r')}{Y(r')}\right)dr'}. \end{aligned}$$

Let us now calculate the marginal benefits of redistribution. Assume that Assumption 2 is satisfied. Then, using the Inada conditions which imply that $\lambda(-U_Y)$, λU_C and $\lambda\beta RV'$ converge to zero, we obtain that the marginal benefit of redistributing income (or leisure) satisfies

$$\begin{aligned} B_Y(r) &\sim \mathbb{E}\left[\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_Y}\left(\frac{C(r')}{C(r)}\right)^{s_C(1-s_Z)\zeta_{CY}}|r'\geq r\right] \\ &= \frac{1}{1+\zeta_Y/\rho_Y-s_C(1-s_Z)\zeta_{CY}/\rho_C}. \end{aligned}$$

The marginal benefit of redistributing savings satisfies

$$B_S(r)\sim\mathbb{E}\left[\left(\frac{S(r')}{S(r)}\right)^{\zeta_S}|r'\geq r\right]=\frac{1}{1-\zeta_S/\rho_S}.$$

The marginal benefit of redistributing consumption is given by

$$\begin{aligned} B_C(r) &\sim \mathbb{E}\left[\left(\frac{C(r')}{C(r)}\right)^{(1-s_Z)\zeta_C}\left(\frac{Y(r')}{Y(r)}\right)^{-\zeta_{CY}}|r'\geq r\right] \\ &= \frac{1}{1-(1-s_Z)\zeta_C/\rho_C+\zeta_{CY}/\rho_Y}. \end{aligned}$$

It follows that

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

and

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

where $1>(1-s_Z)\zeta_C/\rho_C-\zeta_{CY}/\rho_Y$, and $1>-\zeta_Y/\rho_Y+s_C(1-s_Z)\zeta_{CY}/\rho_C$, and $1>\zeta_S/\rho_S$. □