Benefit-Based Taxation can Include Substantial Redistribution

Jo Ellery, Matthew Weinzierl*

March 2026

Abstract

Contrary to a long-standing concern, we demonstrate that classical benefit-based labor income taxation in the tradition of Smith (1776) is consistent with substantial income redistribution through both transfers and negative average income tax rates. As redistribution is widespread in practice, our results remove an obstacle to the benefit principle's plausibility as the prevailing normative basis of tax policy.

Introduction

Economists since before Smith (1776) have been drawn to the principle of benefit-based taxation as a foundation—even *the* foundation—for the design of tax policy. However, the widely held belief that the benefit principle is incompatible with redistribution has prevented many modern scholars from embracing it. The idea behind benefit-based taxation is that taxpayers ought to pay taxes according to the benefit they obtain from the activities of the state. It seems fair to assume that all persons obtain some such benefit from the state's activities, from which it would appear to follow that all should pay positive taxes; that is, that there should be neither transfers nor negative average income tax rates in a benefit-based tax system.

Given the ubiquity of redistribution in modern income tax systems, a principle incompatible with it, no matter how appealing for other reasons, seems an inappropriate guide to how taxes are or should be designed. For example, Liam Murphy and Thomas Nagel write: "For though the very poor benefit less from government than the rich, they still benefit greatly as against the baseline of the war of all against all—especially in a country with at least a minimal welfare system. According to the benefit principle, then, the poor must pay for this benefit in proportion to its size... The benefit principle is, in fact, incompatible, as a matter of political morality, with every account of social justice that requires government to provide any kind of income support or welfare provision whatsoever." (Murphy and Nagel, 2002, p. 18-19)

We dispel this concern about benefit-based taxation. To do so, we demonstrate that optimal benefit-based labor income taxation can include transfers and negative tax rates both theoretically and—in plausibly-specified simulations—quantitatively at levels that explain a substantial share of redistribution prevailing in U.S. policy. Notably, we show that a benefit-based approach, in contrast to the utilitarian objectives commonly assumed in optimal tax research, can readily generate transfers throughout the income distribution, as seen in practice.

*Harvard Business School; jellery@hbs.edu and mweinzierl@hbs.edu. Neither author received funding other than from Harvard University in support of this research, and neither has relevant financial relationships or conflicts of interest related to this research. We thank Jim Hines, Joel Slemrod, Dmitry Taubinsky, and participants at the University of Michigan and National Tax Association Annual Conference for helpful comments.

The key idea behind our results is that redistribution can contribute to overall prosperity, for example by preserving public support for a market economy, which benefits all taxpayers. In a benefit-based tax system, taxpayers may therefore choose to fund redistribution as a part of overall productive government spending. This idea has a long history, from the founding of the modern welfare state under German Chancellor Otto von Bismarck in the late nineteenth century,[1] to the New Deal legislation of U.S. President Franklin Delano Roosevelt in the mid-twentieth century,[2] to the broad reach of many social safety net programs today.[3] It is no surprise, then, that there exists a large body of research in political science providing evidence of this connection.[4] But this idea has not, as far as we are aware, been previously studied formally in the optimal tax literature.

In Section 1, we build our formal model. We adapt the canonical Lindahl (1919) model of benefit-based taxation and allow each taxpayer's income-earning ability to depend on the full set of the state's activities, thereby formalizing the "classical" benefit-based view of income taxation that Musgrave (1956) attributed to Smith. In keeping with the idea of productive redistribution, we include transfers as a component of those activities.

In Section 2, we specify functional forms for the model and characterize optimal policy. We derive an expression for optimal tax shares and uncover the novel result that a benefit-based income tax policy will include negative tax shares (and thus negative average tax rates) for low-income earners whenever an expansion of government spending sufficiently increases the spread of income relative to its average level.[5] In the benefit-based approach, obtaining low-income workers' agreement on policy may require higher earners to pay for more than all of government spending.

In Section 3 we generalize our model to include multiple components of government spending (specifically non-transfer spending and transfers) and simulate it numerically. We show that, using plausible parameter values, our model can generate a policy equilibrium that is consistent with U.S. data on incomes, tax shares, the wage distribution, and per household government spending, including a substantial share of transfers.

[1] Explaining the need for his *Sozialstaat*—with its national health insurance, pensions, and unemployment insurance—Bismarck said (see Dawson (1891)) "I maintain that it is [the state's] duty... there are purposes which only the State as a whole can fulfill. To these belong national defense, the general system of communications, and... the help of the necessitous and the removal of those just complaints which provide Social Democracy with really effective material for agitation." It is important to note that Bismarck's reasoning is not simply political; he views the poor as having "just complaints" that can lead to "agitation," and he views it as a duty of the state to address those complaints if it is to retain its legitimacy to all its citizens. In this way, we see how a benefit-based logic for redistribution grants equal respect to all citizens, including those who fund and receive it.

[2] Against a background of rising skepticism toward the market economy, FDR's policies dramatically expanded the social safety net. Historian H.W. Brands writes of FDR's policies: "In the generations that followed, as the American economy continued to thrive and as the benefits of America's material fortune rained down on the wealthy even more than on persons of moderate means, the objective and honest of those who had once denounced Roosevelt for class betrayal recognized that in a decade rife with fascists, militarists, and communists abroad and irresponsible demagogues at home, he was the best thing that could have happened to them." (Brands, 2008).

[3] For example, approximately 60 percent of U.S. households receive more in transfers than they pay in federal taxes (see https://taxfoundation.org/blog/60-percent-households-now-receive-more-transfer-income-they-pay-taxes/)

[4] See Baccini, Hicks, and Rettl (2025), who write that "a large literature establishes a convincing link between social spending and political stability... For example, these policies increase support for openness... and decrease the risk of political backlash..." A related point was made more than a half-century ago by Lester Thurow in Thurow (1971): "Preventing crime and creating social or political stability may depend on preserving a narrow distribution of income or a distribution of income that does not have a lower tail." See also Pauly (1973).

[5] Though this scenario may seem to be unlikely in practice, recent research on trade (D. H. Autor, Dorn, and Hanson (2013)) and task automation (Acemoglu and Restrepo (2019)) demonstrates that policy can increase overall economic efficiency while creating costly short-run dislocations for low-income workers; see also advances in artificial intelligence (Acemoglu, D. Autor, et al. (2022)). To address concerns that this result relies on marginal measures of benefit, in appendix A.3 we discuss why Lindahl's model is immune to them and how the alternative to Lindahl proposed by Moulin (1987) can extend our ideas inframarginal measures.

1 A model of benefit-based labor income taxation

In this section, we develop a model of benefit-based taxation that builds on canonical past contributions. We begin with the first maxim of taxation in Smith (1776):

> Subjects of every state ought to contribute towards the support of the government, as nearly as possible, in proportion to their respective abilities; that is, in proportion to the revenue which they respectively enjoy under the protection of the state.

Musgrave (1956), the leading twentieth century scholar of the history of public finance theory, advanced the interpretation that Smith's maxim combines what tax theorists typically call the "benefit" principle and the "ability to pay" principle. It does so by having the primary benefit individuals obtain from the state's activities flow through their income-earning abilities, thereby enshrining the logic that one's benefit from the state is one's ability to pay. Musgrave called this idea classical benefit-based taxation.[6]

Our formalization of Smith's classical benefit-based taxation builds this link between benefit and ability into the approach of Lindahl (1919), whose "positive solution" attempted to capture the political process by which taxpayers collectively decide on a level of government spending and taxes. Conceptually, Lindahl's solution works as follows. We present to each taxpayer all feasible schedules of individual tax shares: that is, all sets of shares of the total tax bill that sum, across individuals, to one. We then ask taxpayers to choose, for each such schedule, their individually preferred level of government spending. When one such tax schedule induces all taxpayers to choose the same level of government spending, a Lindahl equilibrium is found.[7]

Formally, we define a Lindahl equilibrium as follows.

Definition 1. *A Lindahl equilibrium is a set of tax shares τ_i for all i and a level of government spending G (which includes required transfers $\{R^i\}_{i\in I}$) that solve, for all $i \in I$:*

$$\max_{c_i,l_i,G_i} U(c_i, l_i) = u\left(c_i, \frac{y_i}{f(a_i, G_i)}\right) \tag{1}$$

subject to:

$$y_i + R^i - c_i - \tau_i G_i \geq 0, \tag{2}$$

$$\sum_i \tau_i = 1, \tag{3}$$

and

$$G_i = G \ \ \forall i. \tag{4}$$

In this definition, utility in expression (1) is a function of consumption c_i and labor effort $l_i = y_i/w_i$, where y_i is income and $w_i = f(a_i, G)$ is income-earning ability. Smith's idea that benefit flows through ability is captured by income-earning ability being a function of the individual's

[6] We focus on how government spending benefits the individual through income-earning ability, but many other channels of benefit exist and can be understood as corresponding to other tax instruments. For example, government spending might impact the utility from consumption through complementary public goods and services, or it might directly generate consumption-like utility. These and other channels of benefit may correspond to benefit-based sales taxes (e.g., VATs), wealth or property taxes, and so on.

[7] Lindahl additionally shows that such an equilibrium always exists, regardless of how many agents are in the bargaining process. Because we focus on whether benefit-based taxation can support redistribution as a normative matter, we start with a focus on the "first-best" case, where equilibrium tax shares can be implemented without causing behavioral distortions. We show that our results are robust to incorporating second-best effects, following Weinzierl (2018), in the Appendix.

innate ability a_i and government spending G.[8] Note, for future reference, that the form of $f(\cdot)$ determines the specific ways in which individuals benefit from the activities of the state. The budget constraint for i in expression (2) states that income y_i and transfers received from the government R^i (which the individual takes as given in this section's model) fund consumption c_i and i's share of government spending $\tau_i G_i$. Expressions (3) and (4) together capture Lindahl's equilibrium requirement that all individuals collectively agree to fully fund the same preferred level of government spending.[9]

2 Functional Forms and Optimal Tax Shares

In this section, we specify functional forms for the benefit-based tax model introduced in Section 1 to characterize optimal policy and to enable us to study the model numerically in the next section.

We assume that utility is separable and isoelastic, such that individual i's utility is:

$$U(c_i, l_i) = \frac{c_i^{1-\rho} - 1}{1-\rho} - \frac{1}{\theta}\left(\frac{y_i}{f(a_i, G)}\right)^{\theta}, \tag{5}$$

with scalars ρ controlling the concavity of consumption utility and θ the elasticity of labor effort.[10]

2.1 The ability production function

A crucial question for our analysis is the form of $f(a_i, G)$, the "production function" determining how innate ability and the activities of the state combine to produce individual i's income-earning ability w_i. An obstacle to specifying this form is that we observe neither the underlying distribution of a_i nor (as is familiar from the Mirrlees tradition) the distribution of w_i. Simply backing out a form for the ability production function $f(\cdot)$ is thus infeasible, and to our knowledge no literature has convincingly identified the form of such an ability production function.

Our approach is as follows: rather than directly assuming a functional form for how income-earning abilities w_i depend on innate ability a_i and government spending G through $f(a_i, G)$, we assume an empirically plausible functional form for the distribution of w_i, allow the parameters of that distribution to be functions of G, and then assign individuals indexed by i the value of w_i in that distribution corresponding to i's percentile in the distribution of innate abilities a_i. This procedure assumes that the function $f(a_i, G)$ is order-preserving, that is that $w_j > w_i$ if and only if $a_j > a_i$, so that people of higher innate ability earn higher wages. Otherwise, however, this method provides substantial flexibility in allowing for the underlying functional relationships to be determined by the data.

[8] In this and the next section, we treat G as a monolith, abstracting from its composition, how that composition is determined, and how different components of G affect the economy, with one exception. That exception is in expression (2), where we include transfers R^i as an exogenous input to the individual's budget constraint. Separating out transfers as a component of G in this way is a placeholder; in section (3) we will endogenize the choice of transfers as part of generalized version of this model that includes multiple components of government spending. For now, it simplifies both our analysis and description of the model to work with a unified measure of government spending, G, but section (3) demonstrates that the extension to multiple components is technically straightforward (if notationally messy).

[9] In this paper, we consider a scenario of political equality. All must consent to chosen policy bundles in the tradition of Wicksell (1896) (who was, not coincidentally, Lindahl's teacher). Extending the model to a world with unequal bargaining power or asymmetric information remains an avenue for future work.

[10] Note that we could add an additional parameter k, multiplying labor disutility, to adjust the levels of consumption and labor. We do not add this parameter here, which is equivalent to setting $k = 1$ throughout our simulations.

Specifically, we assume that the distribution of w_i is log logistic, closely related to the Pareto and Champernowne forms commonly used in the optimal tax literature and a form used to describe incomes since at least Fisk (1961).[11] Let $L(\mu, \sigma)$ denote a logistic distribution with location parameter μ and scale parameter σ. Then, we have:

$$\ln(f(a_i, G)) \sim L(\mu(G), \sigma(G)),$$

where G serves as an argument to the location and scale parameters. Let Λ denote the CDF of the standard logistic distribution (i.e., with location parameter 0 and scale parameter 1),[12] and note that its value at $f(a_i, G)$, appropriately standardized by the mean and variance $\mu(G)$ and $\sigma(G)$, equals the percentile $P(a_i)$ of individual i's innate ability in the population:

$$P(a_i) = \Lambda\left(\frac{\ln(f(a_i, G)) - \mu(G)}{\sigma(G)}\right).$$

Rearranging terms, i's income-earning ability can be expressed as:

$$ln(f(a_i, G)) = \mu(G) + \sigma(G)\Lambda^{-1}(P(a_i)). \tag{6}$$

Expression (6) for w_i shows how we can use the dependence on G of the two parameters of the logistic distribution to create an indirect formula for income-earning ability as a function of innate ability a_i and G. Note that the first term, $\mu(G)$, is independent of a_i and depends only on the level of government spending G, while the second term, $\sigma(G)\Lambda^{-1}(P(a_i))$, is the product of the spread parameter of the w_i distribution and the value of the inverse CDF of the standard logistic at the individual's innate ability percentile.

To simplify our notation going forward, define the inverse CDF of the logistic distribution as $q_i = \Lambda^{-1}(P(a_i))$, a zero-centered marker of ability. Because the logistic distribution described above is mean zero, the inverse CDF will return a negative number for persons whose innate ability is below the mean (that is, persons for whom $P(a_i) < 0.5$). Thus q_i will be negative for low-ability individuals and positive for high-ability individuals.

Next, we assume functional forms for $\mu(G)$ and $\sigma(G)$:

$$\mu(G) = \bar{\mu} + \ln(G^{\beta_1}) - \beta_2 G^{\beta_3}, \tag{7}$$

$$\sigma(G) = \bar{\sigma} + \alpha_1 G^{\alpha_2}, \tag{8}$$

where all parameters other than the "baseline" values $\bar{\mu}$ and $\bar{\sigma}$ are assumed non-negative, that is:

$$\{\beta_1, \beta_2, \beta_3, \alpha_1, \alpha_2\} \geq 0.$$

We can now write i's income-earning ability as:

$$f(a_i, G) = G^{\beta_1} e^{\bar{\mu} - \beta_2 G^{\beta_3}} e^{(\bar{\sigma} + \alpha_1 G^{\alpha_2}) q_i}. \tag{9}$$

This specification for income-earning ability is empirically tractable (in logs) and has several appealing conceptual features.

First, the core idea of classical benefit-based taxation is that the activities of the state have the potential to magnify individuals' income-earning abilities. That idea is at its most compelling when

[11] For a thorough explanation of how these distributions are related, see appendix C.

[12] The spread of the logistic distribution is characterized by a scale parameter which is not equal to the variance. In this case, a scale parameter of one would imply a variance of $\frac{\pi^2}{3}$.

no state exists, a point made forcefully by Murphy and Nagel (2002). That is, even minimal-state libertarians recognize that some activities of the state are essential for modern economic production (see, for example, Nozick (1974)). To capture that reality formally, when $G = 0$ our functional forms should yield both negligible income-earning abilities and a positive impact of government spending on those abilities for all individuals. The above forms of $\mu(G)$ and $\sigma(G)$ serve that purpose. The first factor of $f(a_i, G)$, G^{β_1}, implies that for any i:

$$\lim_{G \to 0^+} f(a_i, G) = 0.$$

In words, as we approach anarchy, all earning abilities decrease to nil. We can also show that

$$\lim_{G \to 0^+} f_G(a_i, G) = \infty$$

for all i as long as $\beta_1 \in (0, 1)$ (i.e., as long as ability is concave in G).[13] That is, when starting in anarchy, adding government spending has a large positive impact on earning abilities.

Second, though increasing G from low levels raises income-earning abilities, at some scale we would expect further increases in government spending to be unproductive. The third term of expression (7), $-\beta_2 G^{\beta_3}$, allows us to capture that effect. Taking the partial derivative of expression (9) with respect to G, we have:

$$f_G(a_i, G) = \left(\beta_1 - \beta_2\beta_3 G^{\beta_3} + \alpha_1\alpha_2 G^{\alpha_2} q_i\right) G^{\beta_1 - 1} e^{\bar{\mu} - \beta_2 G^{\beta_3} + (\bar{\sigma} + \alpha_1 G^{\alpha_2}) q_i} \tag{10}$$

We see that $f_G(a_i, G) < 0$ for the individual with $q_i = 0$ when $\beta_1 < \beta_2\beta_3 G_3^{\beta}$. That is, at a large enough level of G, marginal government spending lowers, rather than raises, median income-earning ability. Notably, the parameter values from our preferred calibrations in Section 3 imply that this condition would hold for values of G only slightly greater than the 2019 value.[14]

Third, in the simplifying case of $\beta_2 = \alpha_1 = 0$, expression 9 reduces to

$$f(a_i, G) = G^{\beta_1} e^{\bar{\mu}} e^{\bar{\sigma} q_i}.$$

This Cobb-Douglas form suggests that $G's$ "factor share" in output is tied to β_1 and, as we will see in the next section, gives a rough sense for the optimal benefit-based levels of government spending and tax rates.

Fourth, the form for $\sigma(G)$ in expression (8) means that government spending may increase the spread of income-earning abilities. That is, the assumption that $\{\alpha_1, \alpha_2\} \geq 0$ implies that $\sigma'(G) \geq 0$.[15] This feature has implications for optimal tax rates, as we will see below.

These forms do have limitations, of course. For example, there is no reliable way for us to increase the impact of G on high-earner incomes without simultaneously impacting low-earner incomes. In other words, we cannot reliably increase spread asymmetrically. Future work may generalize these forms in a principled way; here, we aim to show that the model using these relatively simple forms can match modern U.S. taxes and incomes.

[13] The derivation of these results is in appendix D.

[14] Note that if $\beta_1, \beta_2, \beta_3 > 0$, then $\lim_{G \to \infty} f_G(a_i, G) = 0$. Though for most values of β_2, β_3 the derivative will become negative at large G levels, it will also decline in magnitude as G approaches infinity. This derivation is in appendix D.

[15] A more nuanced form could allow $\sigma'(G)$ to be positive for some levels of G and negative for others.

2.2 Optimal tax shares

In the general model of Section 1, individual i's first-order conditions imply that i's optimal tax share τ_i is:[16]

$$\tau_i = \left(\frac{y_i}{f(a_i, G)} \right) f_G(a_i, G). \tag{11}$$

In words, expression (11) implies that i pays a tax share corresponding to the marginal effect of government spending on i's income-earning ability.

One implication of this result helps build intuition for our functional form assumptions. Consider the simplifying case from Section 2.1 in which $\beta_2 = \alpha_1 = 0$, where the expression for $i's$ income-earning ability reduces to $f(a_i, G) = G^{\beta_1} e^{\bar{\mu}} e^{\bar{\sigma} q_i}$. In that case, we can show the following result.

Proposition 1. *In the Lindahl equilibrium of expressions (1) through (4), with functional forms from expressions (5) and (9), if $\beta_2 = \alpha_1 = 0$, the optimal effective tax rate $t_i = \tau_i G / y_i$ equals β_1 for all i, and government spending as a share of total income also equals β_1.*

In words, this simplifying case allows us to interpret β_1 as the factor share of G in output. Because we are solving for optimal *benefit-based* taxes and that factor share measures the marginal effect of government spending on output (i.e., its benefit), the factor ought to receive a corresponding income share, just as in a Cobb-Douglas production function. With the assumption that $\alpha_1 = 0$, there is no room for government spending to have heterogeneous effects on individuals, so all pay the same rate.

A second implication, new to this paper, relates more directly to our focus on redistribution.

Proposition 2. *In the Lindahl equilibrium of expressions (1) through (4), with functional forms from expressions (5) and (9), the tax share of individual i is negative if government spending lowers i's income-earning ability on the margin. That is: if $f_G(a_i, G) < 0$ then $\tau_i < 0$.*

In words, the Lindahl process of agreement among taxpayers results in a collective willingness to compensate, through assigning them a negative tax share, individuals whose income-earning ability is reduced (on the margin) by the activities of the state. This result reflects the analogy of the Lindahl approach to conventional markets for goods and services, where prices are set on the margin.[17] Proofs for propositions 1 and 2 are included in appendix B.

Because we have specified functional forms in this section, we can say more about when and for whom there will be negative tax rates. Using expression (6) for $f(a_i, G)$ in expression (11), i's tax share becomes:

$$\tau_i = \frac{\beta_1 - \beta_2 \beta_3 G^{\beta_3} + \alpha_1 \alpha_2 G^{\alpha_2} q_i}{G} y_i. \tag{12}$$

We can convert this tax share into i's tax *rate* on income:

$$t_i = \frac{\tau_i G}{y_i} = \beta_1 - \beta_2 \beta_3 G^{\beta_3} + \alpha_1 \alpha_2 G^{\alpha_2} q_i. \tag{13}$$

Recall that q_i is increasing in innate ability, so that the final term of expression (13) implies that the tax rate t_i is weakly increasing in innate ability given our functional form and parameter sign assumptions. Note also that if $\beta_1 < \beta_2 \beta_3 G_3^{\beta}$, then t_i will be negative for $q_i = 0$, consistent with the earlier observation about the condition under which government spending becomes counterproductive for the average citizen.

[16] See appendix A.1 for the full derivation of this condition.

[17] In appendix A.3, we address concerns that our results may rely too heavily on Lindahl's focus on the margin.

Expression (12) and expression (13) show that the optimal benefit-based tax share and tax rate for i may be pushed below zero by the effects of government spending for low-income individuals, as their values for q_i will be negative. In fact, we can use these expressions to determine a threshold of ability below which individual i will face a negative tax rate, denoted q_i^*:

$$q_i^* = -\frac{\beta_1 - \beta_2\beta_3 G^{\beta_3}}{\alpha_1\alpha_2 G^{\alpha_2}}. \tag{14}$$

In expression (14), the numerator reflects the effect of G on the average level of income-earning ability while the denominator reflects its effect on the spread of income-earning ability.[18] We can use this expression to see two distinct effects of government spending on the share of taxpayers with a negative tax share in equilibrium.

First, as the familiar saying goes, a rising tide lifts all boats. If G has a larger positive impact on the average level of income-earning ability, the numerator of expression (14) will be larger and the absolute magnitude of q_i^* will increase. Because q_i^* is typically negative, when the absolute magnitude of q_i^* rises, q_i^* will be more negative and a smaller share of individuals will fall below the threshold and face a negative tax rate. For example, as β_1 increases, the benefit of G increases for all individuals, and thus q_i^* and the share of individuals with negative tax rates both fall.

In contrast, rising inequality can result in smaller pieces of the economic pie for some, even if the overall pie is larger. If G increases the spread of income-earning abilities because α_1, α_2, or their product is larger, the denominator of equation (14) will be larger, decreasing the absolute magnitude of q_i^*. In other words, when the spread of income-earning ability reacts more strongly to government spending, more agents will be "left behind" and will obtain negative tax rates in the policy equilibrium.

3 Calibration and Simulation Results

In this section, we demonstrate that a generalized version of our model, using the functional forms described above and plausible parameter values, can generate numerical results consistent with U.S. policy and outcomes data, including a substantial share of prevailing levels of redistribution through transfers and negative tax rates.[19]

3.1 Generalizing to multiple components of government spending

To study the extent of redistribution, we generalize our model to allow taxpayers to choose multiple components of government spending. In the original model of section (1), for expositional simplicity we assumed transfers were an exogenously determined part of the monolithic G over which taxpayers exerted choice. Now, we allow taxpayers to choose the levels of transfers $\{R^j\}_{j\in I}$, one for each ability type, as well as a catch-all category of non-transfer spending which we now label X for "expenditure."[20] One interesting path for future research is to further divide this non-transfer spending into distinct components with their own implications for taxpayers.

[18] The quantity q_i^* will typically be negative in our context; $q_i^* = 0$ would imply that half of agents receive negative tax rates, far more than is typically true. See appendix E for more comparative statics on q_i and the related $P(a_i^*)$.

[19] In this version of the paper, we focus on recent data. In future revisions, we will also use data from 2000 and 1980 to extend the explanatory exercise.

[20] We could mimic the section (1) approach in numerical simulations by imposing an exogenous constraint that taxes must fund transfers at empirical levels. In a previous version of this paper, we showed that this approach can closely match prevailing policies in the United States over the past several decades. Results are available from the authors upon request.

With this more general setup, we follow essentially the same process as in our original model to define a Lindahl equilibrium. For each of the components of government spending, we present to each taxpayer all feasible schedules of individual tax shares for that component: that is, all sets of shares of the tax bill for each component of spending that sum, across individuals, to one. We then ask taxpayers to choose simultaneously, for all such combinations of schedules, their individually preferred levels of all components of government spending. When a set of such tax schedules induces all taxpayers to agree on the levels of each component of government spending, a Lindahl equilibrium is found.

Formally, we now define a Lindahl equilibrium as follows.

Definition 2. *A Lindahl equilibrium is a level of government non-transfer spending X, a set of transfers $\{R^j\}_{j\in I}$, and a set of tax shares τ_i^X and $\{\tau_i^{R^j}\}_{j\in I}$ for all j, i that solve, for all $i \in I$:*

$$\max_{c_i, l_i, X_i, \{R^j\}_{j\in I}} U(c_i, l_i) = u\left(c_i, \frac{y_i}{f(a_i, X_i, \{R_i^j\}_{j\in I})}\right) \tag{15}$$

subject to:

$$y_i + R^i - c_i - \tau_i^X X_i - \sum_{j\in I} \tau_i^{R^j} R_i^j \geq 0, \tag{16}$$

$$\sum_i \tau_i^X = 1 \quad \textit{and} \quad \sum_i \tau_i^{R^j} = 1 \quad \textit{for all } j \in I, \tag{17}$$

where we set

$$\tau_i^{R^i} = 0, \tag{18}$$

and

$$X_i = X \quad \textit{and} \quad R_i^j = R^j \ \ \forall i. \tag{19}$$

This definition closely resembles that in expressions (1) through (4) of section (1), with a few adjustments. As shown in $f(a_i, X_i, \{R_i^j\}_{j\in I})$ from expression (15), income-earning ability now depends separately on transfers and non-transfer government spending, and the budget constraint for i in expression (16) includes not only the transfer that i receives from the government R^i but also i's tax payments that fund transfers to others $\sum_{j\in I} \tau_i^{R^j} R^j$. Note that expression (18) sets each individual $i's$ tax share of the transfer *it receives* to zero. This constraint ensures that transfers to any type are voluntarily funded by all other types.[21]

Also in keeping with our approach from section (1), we assume that income-earning ability follows a log-logistic distribution as described in expressions (6) through (9), but now we allow transfers and non-transfer spending to each exert their own influence on the location and scale parameters of the logistic function. Formally, we convert expression (6), $ln(f(a_i, G)) = \mu(G) + \sigma(G)\Lambda^{-1}(P(a_i))$, which used a single measure of G, into an expression that shows the effects of the separate components of government spending:

$$ln(f(a_i, X, \{R^j\}_{j\in I}) = \left[\mu(X) + \sum_{j\in I} \mu(R^j)\right] + \left[\sigma(X) + \sum_{j\in I} \sigma(R^j)\right] \Lambda^{-1}(P(a_i)). \tag{20}$$

[21] This constraint is without loss of generality, in that any positive tax share of a type's own transfer would simply reduce the net transfer to it. Setting the share to zero is a natural normalization. Note that transfers to individual i will be funded by other taxpayers only to the extent that they benefit those taxpayers. Transfers thus act as part of productive government spending, not negative taxes, consistent with the statutory distinction between the tax and transfer system; after all, if transfers were equivalent to negative taxes, there would be no need to have both.

We can also extend the parametric specifications of these functional form assumptions to this multi-component setting, with notation for the β and α parameters now including superscripts to denote whether they apply to non-transfer spending X or transfers R^j:

$$ln(f(a_i, X, \{R^j\}_{j\in I}) = \left(X^{\beta_1^X}\prod(R^j)^{\beta_1^j}\right)\left(e^{\bar{\mu}-\beta_2^X X^{\beta_3^X}-\sum_j \beta_2^j (R^j)^{\beta_3^j}}\right)\left(e^{\left[\bar{\sigma}+\alpha_1^X X^{\alpha_2^X}+\sum \alpha_1^j (R^j)^{\alpha_2^j}\right]q_i}\right). \tag{21}$$

3.2 Empirical targets

To see how well our model can match reality, we first need to establish our empirical targets. As this is a model of income taxation with a focus on the role of the state and redistribution, we will target household incomes and income tax shares, along with levels of transfers and non-transfer government spending and the shape of the income-earning ability distribution.

To measure incomes and tax shares, we gather annual data on household market income (income before taxes and transfers, not including any forms of social insurance) and average income tax rates, which together give us tax shares, by household population quintile from the Congressional Budget Office for 1979 through 2019.[22] The difference between final income and market income, plus what was paid in taxes, gives us redistributive transfers.

For reference, Figure 1 shows coincident trends in the data of increasing income inequality and increasing redistribution over the past several decades. We plot each household quintile at its range's midpoint.

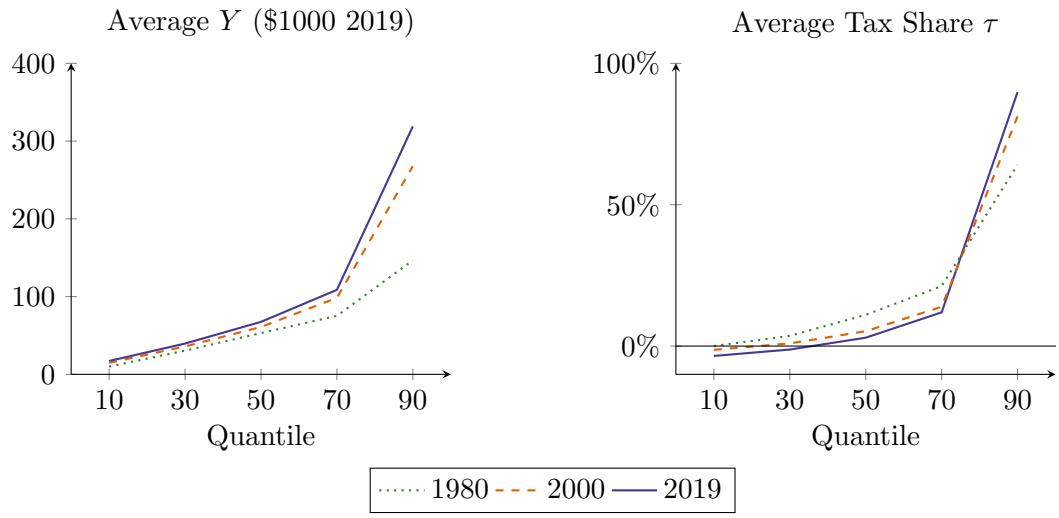


Figure 1: Average market income and average tax share by quintile for 1980, 2000 and 2019. The data is taken from the CBO. The plot on the left highlights both the across-the-board increase in incomes over time and the increase in inequality over time. The plot on the right shows the increasingly redistributive U.S. tax system.

As this figure shows, since 1987 average income tax rates for the bottom quintile have been negative, reaching rates below -10% in the last decade. The second-lowest quintile has also had negative average tax rates for the past two decades. In Figure 1, we convert these tax rates into tax

[22]Both datasets can be found in the CBO's Distribution of Household Income reports. Negative tax rates reflect the EITC and other tax credits.

shares.[23] A combination of rising inequality and increasingly negative tax shares at the bottom of the income distribution have meant that top quintiles pay an increasingly large share of the total income tax bill, even as tax *rates* on top earners have largely stayed flat or fallen.

To measure government non-transfer spending, we take the difference between federal net outlays (which includes transfer payments)[24] and the transfers estimated from CBO data as described above. For the shape of the income-earning ability distribution, we feed the data on incomes, tax shares,[25] transfers, and government spending into the household first-order conditions. Those first-order conditions, using the functional forms from Section 2, depend on utility function parameters, which we set at values within conventional ranges: the curvature of consumption utility $\rho = 1.5$ and the curvature of labor disutility $\theta = 3$ (so that the Frisch elasticity of labor supply is 0.5). We then back out implied values for w_i[26] and fit a log logistic distribution to them, providing the target estimates of the mean and spread parameters, which we denote $\hat{\mu}$ and $\hat{\sigma}$.

3.3 Parameter values

Once we have established these empirical targets, we turn to the task of calibrating the parameters of the ability production function so that the resulting Lindahl equilibrium values for incomes, tax shares, transfers and non-transfer government spending, and wages match those targets as closely as possible. We calibrate our parameters using three quantiles (the 10th, 50th, and 90th percentiles) in order to improve convergence speed.

We start by hypothesizing plausible starting values for some of the parameters. First, consider parameters relevant to non-transfer government spending, X. As in our analysis of the model in section (1) with monolithic G, when $\beta_2^X = 0$, the functional form for $\mu(X)$ takes a familiar Cobb-Douglas structure, where β_1^X corresponds to non-transfer spending's "factor share" of total output. Our measure of X is approximately 48 percent of our measure of total government spending in 2019, which in turn is approximately 21 percent of U.S. gross domestic product and 39 percent of labor income (the only form of income in our model economy), implying X is approximately 19 percent of total labor income.[27] We therefore consider values around 0.2 for β_1^X. Once that value is chosen, we can use it, along with with the empirical values of non-transfer spending and our estimate for $\hat{\mu}$, to choose a starting value for $\bar{\mu}$. We can then adjust β_1^X, β_2^X and β_3^X so that the values predicted by expression (7) align roughly with the data on non-transfer spending. Similarly, for $\hat{\sigma}$, we assume that $\bar{\sigma} > 0$, and we adjust α_1^X and α_2^X so that the values predicted by expression (8) roughly match our target $\hat{\sigma}$ (given empirical values of non-transfer spending).

Next, we do the same for transfer spending. We apply the same values for $\beta_2^{R^j}$ and $\beta_3^{R^j}$ as we used for non-transfer spending, but we assume that transfers are less effective at raising the total mean earning ability than general spending, implying smaller $\beta_1^{R^j}$. We further assume that transfers do not increase earning ability variance independent of their direct impact on earnings,

[23] As a reminder, if t_i is the average tax rate for quintile i and y_i its average income, its tax *share* is $\tau_i = t_i y_i / \left(\sum_j t_j y_j\right)$.

[24] See series FYONGDA188S at https://fred.stlouisfed.org/

[25] As in expression (2), we use the tax share, not the tax rate, in the budget constraint to maintain consistency across the models with which we infer our empirical targets and try to match them.

[26] Using our model, we back out the distribution of wages w_i assuming that agents maximize by using the labor first order condition $w_i = \left(c^\rho y^{\theta-1}\right)^{\frac{1}{\theta}}$. The tax rate τ enters in consumption. Because we do not ask individuals to make choices subject to a linear tax rate, we do not include a linear tax rate in our moment conditions.

[27] In 2019, federal net outlays were 20.6 percent of gross domestic product, which was $21.5 trillion, while compensation of employees (labor income) was $11.4 trillion. See Table 1.1.5 at https://apps.bea.gov and series FYONGDA188S and A033RC1A027NBEA at https://fred.stlouisfed.org

implying that $\alpha_1^{R^j}$ and $\alpha_2^{R^j}$ equal zero for all $j \in I$. Of course, these assumptions may be modified in future work to explore alternative parameterizations, including based on estimation from data.

We report our preferred ability production function parameter values in Table 1, along with the resulting values of $\mu(G)$ and $\sigma(G)$ (where G is total spending), and the target values of $\hat{\mu}$ and $\hat{\sigma}$.

Table 1: Calibrated values of model parameters (columns) by types of government spending (rows) for 2019. The first four columns represent the calibrated parameters for $\mu(G)$, while columns (5)-(7) give the calibrated parameters of $\sigma(G)$. Finally, columns (8)-(11) show the total resulting values of $\mu(G)$ and $\sigma(G)$, and the target values of $\hat{\mu}$ and $\hat{\sigma}$.

	$\bar{\mu}$	β_1	β_2	β_3	$\bar{\sigma}$	α_1	α_2	$\mu(G)$	$\hat{\mu}$	$\sigma(G)$	$\hat{\sigma}$
Total	3.75				0.25			5.0	5.01	0.56	0.53
Non-transfer spending X		0.26	0.001	1.1		0.03	0.52				
Transfer to p10		0.05	0.001	1.1		0	0				
Transfer to p50		0.04	0.001	1.1		0	0				
Transfer to p90		0.02	0.001	1.1		0	0				

Though this calibration is only rudimentary, it yields some reassuring results. The value of β_1, 0.26 for non-transfer spending, implies a reasonable "factor share" of government in determining ability. The calibration is able to match $\hat{\mu}$ and $\hat{\sigma}$ well.

3.4 Results

In Figure 2, we show the fit of the model's equilibrium results for incomes and tax shares to the data for 2019. Top incomes are slightly lower than reality, while tax shares are slightly too progressive. Overall, however, the model results closely resemble empirical tax and income distributions.

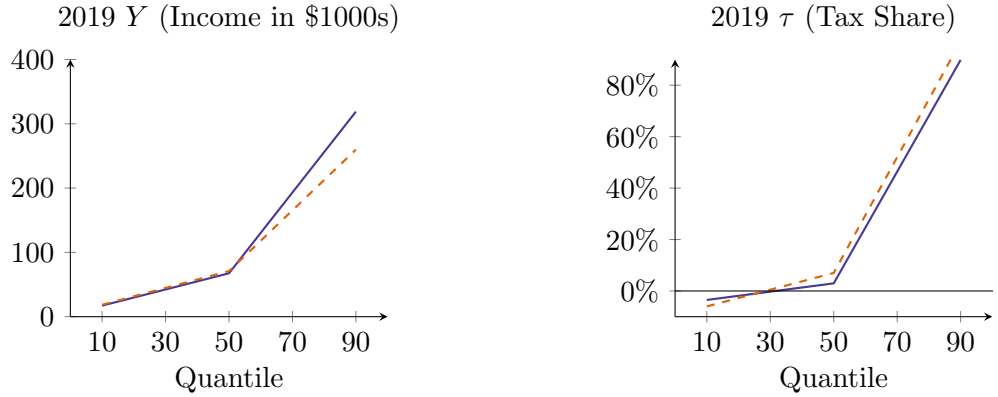


Figure 2: Model Calibration Results for 2019. This figure compares the empirical levels of market income and tax shares (solid lines), taken from CBO income data, to the levels in the calibrated model solution (dashed lines).

In Table 2, we show the fit of the model's equilibrium results for government spending. The overall level of government spending is a close match, while the model somewhat overshoots the level of non-transfer spending and somewhat undershoots—but still delivers a substantial share of—transfer spending.[28] Perhaps most important for the purposes of this paper, the model generates large redistributive transfers to low-income households, contrary to the conventional belief

[28]Per-household data is calculated from per-capita data, multiplied by household size in that year. Household size is calculated as the total U.S. population divided by the number of households; both series come from the Federal Reserve. Household size ranges from 2.9 in 1980 to 2.55 in 2019.

that benefit-based taxation is inconsistent with redistribution. Notably, our results also produce substantial transfers for middle- and high-income households, as our parameters imply that they also contribute to overall productivity (e.g., perhaps through increasing political support for - and thus the legitimacy of - the activities of the state). In contrast, the utilitarian objective typically assumed in optimal tax theory would reject such transfers unless, tellingly, political economy considerations (of, for example, the benefit-based type) were added to the analysis.

Table 2: Annual per household government spending for 2019, overall and by component, in model outcomes and data (shown in $2019 dollars).

	Variable	Model	Data
Per household total spending	G	$38,170	$34,371
Non-transfer spending	X	$29,820	$16,590
P10 mean transfers per household	R_1	$11,650	$21,800
P50 mean transfers per household	R_2	$9,460	$17,900
P90 mean transfers per household	R_3	$3,960	$14,300

Overall, these results suggest that this paper's model of classical benefit-based taxation can successfully reproduce high-level features of the U.S. income tax system over the past several decades, including substantial income redistribution.

4 Conclusion

Contrary to a long-standing concern, we have demonstrated that benefit-based income taxation is consistent with redistribution at the scale and scope of the modern U.S. tax system both theoretically and quantitatively. Arguably, the mechanism behind this result resonates with at least one way to understand modern discussions of economic inequality and the response of the tax system to it: if low-income taxpayers feel left behind by a system that seems to generate large increases in aggregate (and even average) incomes but small increases (or even decreases) in their own incomes, they may condition their support for that system on obtaining agreement to redistribute some of those gains from those who gain the most. In agreeing to redistribution that thereby preserves the legitimacy of the market system and the prosperity it enables, taxpayers can be said to be "protecting capitalism from itself."

This benefit-based justification for redistribution, whether through transfers or negative tax rates, is fundamentally distinct from that of the welfarist view now standard in optimal tax theory built on Mirrlees (1971). In the benefit-based approach, redistribution is not driven by a centralized pursuit of social welfare maximization in which individuals are assumed to have stepped behind a Rawlsian veil of ignorance, but by taxpayers collectively choosing policy in their enlightened self-interest. Given this contrast, it is notable that our numerical results' ability to match prevailing policy compares favorably to that of conventional optimal tax models, where one must often use extreme parameter values to match similar features of reality (see Lockwood and Weinzierl (2016))

We focus, in this paper, on testing whether the benefit principle is consistent with redistribution, not justifying benefit-based taxation more generally or Lindahl's formal approach to it. As discussed in Stewart (2015), Weinzierl (2018), and Voyiakis (2025), however, both have many appealing features. We hope that this paper's results therefore help to rekindle interest in benefit-based taxation, which we believe has substantial promise as a principle for tax policy design.

References

Acemoglu, Daron, David Autor, et al. (Apr. 2022). "Artificial Intelligence and Jobs: Evidence from Online Vacancies". *Journal of Labor Economics* 40.S1, S293–S340.

Acemoglu, Daron and Pascual Restrepo (May 2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor". *Journal of Economic Perspectives* 33.2, pp. 3–30.

Autor, David H., David Dorn, and Gordon H. Hanson (Oct. 2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States". *American Economic Review* 103.6, pp. 2121–2168.

Baccini, Leo, John Hicks, and Paula Rettl (Feb. 2025). "Populism and Political Trust: Evidence from Latin America".

Brands, H. W. (2008). *Traitor to his class: the privileged life and radical Presidency of Franklin Delano Roosevelt*. First edition. New York: Doubleday Publishing Group.

Dawson, William Harbutt (1891). *Bismarck and State Socialism: An Exposition of the Social and Economic Legislation of Germany Since 1870*. English. S. Sonnenschein & Co.

Fisk, Peter R. (1961). "The Graduation of Income Distributions". *Econometrica* 29.2, pp. 171–185.

Lindahl, Erik (1919). "Just Taxation—A Positive Solution". *Classics in the Theory of Public Finance*. Ed. by Richard A. Musgrave and Alan T. Peacock. International Economic Association Series. London: Palgrave Macmillan UK, pp. 168–176.

Lockwood, Benjamin B. and Matthew Weinzierl (Feb. 2016). "Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions". *Journal of Monetary Economics* 77, pp. 30–47.

Mirrlees, J. A. (1971). "An Exploration in the Theory of Optimum Income Taxation". *The Review of Economic Studies* 38.2, pp. 175–208.

Moulin, Herve (1987). "Egalitarian-Equivalent Cost Sharing of a Public Good". *Econometrica* 55.4, pp. 963–976.

Murphy, Liam and Thomas Nagel (Apr. 2002). *The Myth of Ownership: Taxes and Justice*. Oxford University Press.

Musgrave, Richard A. (1956). "A Multiple Theory of Budget Determination". *FinanzArchiv / Public Finance Analysis* 17.3, pp. 333–343.

Nozick, Robert (1974). *Anarchy, State, and Utopia*. Basic Books.

Pauly, Mark V. (1973). "Income Redistribution as a Local Public Good". *The Journal of Public Economics* 2, pp. 35–58.

Scherf, Robert and Matthew Weinzierl (2020). "Understanding Different Approaches to Benefit-Based Taxation". *Fiscal Studies* 41.2, pp. 385–410.

Smith, Adam (1776). *Wealth of Nations*. Prometheus Books, 1991.

Stewart, Miranda (2015). "The tax state, benefit and legitimacy". *Studies in the History of Tax Law, Volume 7*. Ed. by Peter Harris and Dominic de Cogan. London: Hart Publishing.

Thurow, Lester C. (1971). "The Income Distribution as a Pure Public Good". *The Quarterly Journal of Economics* 85.2, pp. 327–336.

Voyiakis, Emmanuel (2025). "The moral force of the benefit principle". *Economics Philosophy*.

Weinzierl, Matthew (July 2018). "Revisiting the Classical View of Benefit-based Taxation". *The Economic Journal* 128.612, F37–F64.

Wicksell, Knut (1896). "A New Principle of Just Taxation". *Classics in the Theory of Public Finance, pub 1958*. Ed. by Richard A. Musgrave and Alan T. Peacock. International Economic Association Series. London: Palgrave Macmillan UK, pp. 72–118.

Appendix A: Derivation of Equilibrium Tax Rates

A.1: General Form of Equilibrium Tax Share

Recall that c_i denotes an individual i's consumption and y_i their income, such that their work effort is given by $l_i = y_i/w_i$. The individual chooses consumption, income and government spending to maximize their utility, as follows:

$$\max_{c_i,l_i,G_i} U(c_i, l_i) = u\left(c_i, \frac{y_i}{f(a_i, G_i)}\right)$$

subject to their budget constraint:

$$y_i - c_i - \tau_i G_i + R_i \geq 0$$

where τ_i is the share of government spending paid for by an individual of type i. Agents are also subject to two equilibrium constraints:

$$G_i = G^* \ \ \forall i$$
$$\sum_i \tau_i = 1.$$

The presence of R_i as a separate quantity in the individual's budget constraint should not be confused as implying that it is separate from government spending G. Instead, R_i is included in G. If all of G were spent, for example, on public goods such as national defense, R_i would be zero for all i. If all of G were spent on transfers to i, R_i would equal G. To give a numerical example that may help clarify this specification, suppose there are only two individuals, i and j, in the economy, with values for $y_i, c_i, R_i, \tau_i G$ of 100,0,50,50 and 20,10,20,10, respectively. The transfer of 10 to individual j is included in the overall G of 60, so that total income of 120 funds total G of 60, 50 of which is non-transfer spending and 10 of which is sent to j as a transfer, which together with the remaining 60 of income funds 70 in total consumption.

Agents solve a Lagrangian of the form:

$$\mathcal{L} = u\left(c_i, \frac{y_i}{f(a_i, G_i)}\right) + \lambda(y_i - c_i - \tau_i G_i + R_i).$$

This yields three first order conditions:

$$\frac{\partial \mathcal{L}}{\partial c_i} : \ \ u_c\left(c_i, \frac{y_i}{f(a_i, G_i)}\right) = \lambda \tag{A1}$$

$$\frac{\partial \mathcal{L}}{\partial y_i} : \ \ u_l\left(c_i, \frac{y_i}{f(a_i, G_i)}\right)\frac{1}{f(a_i, G_i)} = -\lambda \tag{A2}$$

$$\frac{\partial \mathcal{L}}{\partial G_i} : \ \ u_l\left(c_i, \frac{y_i}{f(a_i, G_i)}\right)\left(\frac{y_i}{f(a_i, G_i)^2}\right) f_G(a_i, G_i) = -\lambda\tau_i. \tag{A3}$$

Combining (A2) and (A3), we see that:

$$\left(\frac{y_i}{f(a_i, G_i)}\right) f_G(a_i, G_i) = \tau_i.$$

We can use this optimality condition, along with our equilibrium constraints, to determine the equilibrium values of G and τ_i. When G_i is at its equilibrium value G, tax shares are given by:

$$\left(\frac{y_i}{f(a_i, G)}\right) f_G(a_i, G) = \tau_i. \tag{A4}$$

A.2: Specific Form of Equilibrium Tax Share

From the previous section, we have:

$$y_i = \frac{\tau_i f(a_i, G_i)}{f_G(a_i, G_i)}. \tag{A5}$$

Recall our chosen form for $f(a_i, G_i)$:

$$f(a_i, G) = G^{\beta_1} \exp\{\bar{\mu} - \beta_2 G^{\beta_3}\} \exp\{(\bar{\sigma} + \alpha_1 G^{\alpha_2}) q_i\}.$$

Plugging this into equation (A5), we can see that income is given by:

$$y_i = \frac{\tau_i G_i}{\beta_1 - \beta_2 \beta_3 G_i^{\beta_3} + q_i \alpha_1 \alpha_2 G_i^{\alpha_2}}.$$

Tax rates are the share of government spending paid by the agent τ_i multiplied by the total government spending G_i and taken as a percentage of total income y_i. Thus we see in this case that tax rates are:

$$t_i = \frac{\tau_i G_i}{y_i} = \beta_1 - \beta_2 \beta_3 G_i^{\beta_3} + q_i \alpha_1 \alpha_2 G_i^{\alpha_2}.$$

Setting $t_i = 0$, our threshold for redistribution can then be written as q_i such that:

$$q_i \leq -\frac{\beta_1 - \beta_2 \beta_3 G_i^{\beta_3}}{\alpha_1 \alpha_2 G_i^{\alpha_2}}.$$

A.3 Discussion of marginal effects

Here, we address a natural concern that the result in Proposition 1 may rely too heavily on marginal benefit.

Two related observations may provide some reassurance on this point.

First, Lindahl's equilibrium is a highly abstract version of a fundamentally political process that is far more complex than the elegant, marginalist structure of the model tempts us to think. The core of Lindahl's idea is that individuals are participating in a process that reconciles their different willingness to pay for what the government does and could do. In this model, what matters is not what a theorist, analyst, or even politician determines the marginal dollar is spent doing. What matters is what voters, the individuals who will ultimately bear the burden and benefit of each policy, see as the effect of the marginal expansion in government on them. This view then informs what bundle of government spending and tax payments they will support.[29] Modern debates over policies like immigration, automation and trade reform provide abundant evidence that low-income voters believe they should be compensated for the harm caused them by policies that have benefited those with higher incomes.

Second, the ideas we explore here are not dependent on measuring benefit at the margin, as they can apply just as well in the model of Moulin (1987) that assigns taxes according to measuring benefit from inframarginal government spending. As Scherf and Weinzierl (2020) explain, an individual's tax share in Moulin's model depends on the benefit the individual obtains from government spending at the equilibrium level G minus the benefit that individual would have obtained from government spending at a level $g^* < G$ (the details of g^* are not important for this paper). Redistributive transfers can play the same productive role in this setting as when taxes are assigned in the Lindahl way, and negative marginal tax rates can arise in Moulin's model if the expansion of the state from g^* to G leaves a taxpayer worse off.

[29] We may question the intertemporal consistency of these beliefs, given policy outcomes, but holding voters accountable for teasing out the effects of policy that econometricians have trouble identifying is questionable in itself.

Appendix B: Proofs of Propositions

Proposition 1

The optimal tax share is given by:

$$\tau_i = \left(\frac{y_i}{f(a_i, G)} \right) f_G(a_i, G). \tag{22}$$

When $\beta_2 = \alpha_1 = 0$, income earning ability is given by the equation:

$$f(a_i, G) = G^{\beta_1} e^{\bar{\mu}} e^{\bar{\sigma} q_i}. \tag{23}$$

The derivative of income earning ability with respect to G is then:

$$f_G(a_i, G) = \beta_1 G^{\beta_1 - 1} e^{\bar{\mu}} e^{\bar{\sigma} q_i}. \tag{24}$$

Combining equations 22 and 24, we see that the optimal tax share is:

$$\tau_i = y_i \frac{\beta_1}{G} \tag{25}$$

Then the optimal tax rate is, for all i:

$$t_i = \beta_1. \tag{26}$$

This implies that government spending as a share of total income is given by:

$$\frac{G_i}{\sum_i y_i} = \frac{\sum_i t_i y_i}{\sum_i y_i} = \frac{\beta_1 \sum_i y_i}{\sum_i y_i} = \beta_1. \tag{27}$$

Proposition 2

The optimal tax share is given by equation 22. We assume that income is always weakly positive. The functional form for earning ability implies that it is positive for all $G > 0$. This means that the sign of the tax share is determined by the sign of:

$$f_G(a_i, G) = \left(\beta_1 - \beta_2 \beta_3 G^{\beta_3} + \alpha_1 \alpha_2 G^{\alpha_2} q_i \right) G^{\beta_1 - 1} e^{\bar{\mu} - \beta_2 G^{\beta_3} + (\bar{\sigma} + \alpha_1 G^{\alpha_2}) q_i} \tag{28}$$

Thus, if $f_G(a_i, G) < 0$, then the tax share $\tau_i < 0$.

Appendix C: Relationships Between Distributions

A series of equivalences connects the distributions discussed in this paper. The log logistic distribution parameterized by a shape and scale parameter describes a random variable, the logarithm of which has a logistic distribution. Suppose X is a random variable with a log logistic distribution with scale e^{μ} and shape $1/\sigma$. Then the variable $Z \sim \ln(X)$ has a logistic distribution with location μ and scale σ.

The log logistic distribution with scale parameter $\alpha = e^{\mu}$ and shape $\beta = 1/\sigma = 1$ (call the density of this $LL(\alpha, \beta)$) has density:

$$LL(x; \alpha = e^{\mu}, \beta = 1) = \frac{(\beta/\alpha)(x/\alpha)^{\beta - 1}}{(1 + (x/\alpha)^{\beta})^2} = \frac{e^{-\mu}}{(1 + x e^{-\mu})^2}.$$

This is equivalent to a generalized Pareto distribution with location parameter $\alpha = 0$, scale $\beta = e^{\mu}$ and shape $\gamma = 1$ (call this $GPD(\alpha, \beta, \gamma)$). We can see this as follows. The generalized Pareto distribution has density:

$$GPD(x; \alpha, \beta, \gamma) = \frac{1}{\beta}\left(1 + \gamma\frac{x-\alpha}{\beta}\right)^{-(1/\gamma+1)}$$

which then implies that:

$$GPD(x; 0, e^{\mu}, 1) = e^{-\mu}\left(1 + xe^{-\mu}\right)^{-2}$$

and thus that

$$LL(x; e^{\mu}, 1) = GPD(x; 0, e^{\mu}, 1).$$

The familiar Pareto distribution with scale α and shape β is itself related to the generalized version as follows. A random variable $X \sim Pareto(\alpha, \beta)$ has density:

$$Pareto(x; \alpha, \beta) = \frac{\beta\alpha^{\beta}}{x^{\beta+1}}$$

which then implies that:

$$Pareto(x; \alpha, \beta) = GPD\left(x; \alpha, \frac{\alpha}{\beta}, \frac{1}{\alpha}\right)$$

as:

$$GPD\left(x; \alpha, \frac{\alpha}{\beta}, \frac{1}{\beta}\right) = \frac{\beta}{\alpha}\left(1 + \frac{x-\alpha}{\alpha}\right)^{-(\beta+1)} = \frac{\beta\alpha^{\beta}}{x^{\beta+1}}$$

The Champernowne distribution, when describing the log of a variable, has three parameters: α, β, γ, and a normalizing constant (defined by the other parameters) n. These parameters describe a density function of the form:

$$Champ(x; \alpha, \beta, \gamma) = \frac{n}{x\left[\frac{1}{2}\left(\frac{x}{\gamma}\right)^{-\alpha} + \frac{1}{2}\left(\frac{x}{\gamma}\right)^{\alpha} + \beta\right]}.$$

When we have $\alpha = 1$, $\beta = 1$ and $\gamma = 1$, then the Champernowne distribution has the same density as a standard log logistic (where $n = \frac{1}{2}$):

$$Champ(x; 1, 1, 1) = \frac{n}{x\left[\frac{1}{2}x^{-1} + \frac{1}{2}x + 1\right]} = \frac{1}{(1+x)^2}.$$

This implies that for $n = \frac{1}{2}$:

$$\begin{aligned} Champ(x; 1, 1, e^{\mu}) &= \frac{n}{\left[\frac{1}{2}e^{\mu} + \frac{1}{2}x^2e^{-\mu} + 1\right]} \\ &= \frac{n}{\frac{1}{2}e^{\mu}\left[1 + x^2e^{-2\mu} + 2xe^{-\mu}\right]} \\ &= \frac{e^{-\mu}}{\left(1 + xe^{-\mu}\right)^2} \end{aligned}$$

Thus, we can translate between the Champernowne and log logistic as:

$$Champ(x; 1, 1, e^{\mu}) = LL(x; e^{\mu}, 1)$$

Appendix D: Results on the behavior of Earning ability for extreme values of G

In this section, we consider a world without transfers. For all results below, we assume that $|q_i| < \infty$, i.e. that $0 < P(a_i) < 1$. This condition ensures that we are considering only interior points of the ability distribution, which seems reasonable if ability is unbounded. Note also that we assume throughout the paper that all parameters $\{\beta_1, \beta_2, \beta_3, \alpha_1, \alpha_2\} \geq 0$; anywhere these parameters are referred to as non-zero, they are assumed positive.

The partial derivative of $f(a_i, G)$ with respect to G is:

$$f_G(a_i, G) = \left(\beta_1 - \beta_2\beta_3 G^{\beta_3} + \alpha_1\alpha_2 G^{\alpha_2} q_i\right) G^{\beta_1 - 1} e^{\bar{\mu} - \beta_2 G^{\beta_3} + (\bar{\sigma} + \alpha_1 G^{\alpha_2})q_i}$$

D.1: Limit as $G \to 0^+$

The value for the marginal value of government spending in the absence of a state (when $G = 0$) will depend on the choice of values of our parameters.

First consider the main case of interest: the limit as $G \to 0^+$ of $f_G(a_i, G)$ when $\beta_1 \in (0, 1)$. This condition implies that we are considering earnings ability that is concave in government spending ($\beta_1 < 1$) and that government spending will strictly raise mean earnings ability when the channel of $\beta_2 G^{\beta_3}$ is suppressed.

As $G \to 0^+$, we see that:

$$\lim_{G \to 0^+} \left(\beta_1 - \beta_2\beta_3 G^{\beta_3} + \alpha_1\alpha_2 G^{\alpha_2} q_i\right) = \beta_1$$

$$\lim_{G \to 0^+} e^{\bar{\mu} - \beta_2 G^{\beta_3} + (\bar{\sigma} + \alpha_1 G^{\alpha_2})q_i} = e^{\bar{\mu} + \bar{\sigma} q_i}$$

$$\lim_{G \to 0^+} G^{\beta_1 - 1} = \infty$$

The terms other than $G^{\beta_1 - 1}$ limit to non-zero constants by assumption, regardless of the value of parameters other than β_1. Therefore by the product rule for limits, we have:

$$\lim_{G \to 0^+} f_G(a_i, G) = \infty$$

If instead $\beta_1 > 1$, each term above will limit to zero, and the entire limit will be zero. In this case, earnings ability is convex in government spending; thus, the first dollar will be the least effective.

If instead $\beta_1 = 0$, the limit will depend on the value of the other parameters. We have:

$$f_G(a_i, G) = \left(\alpha_1\alpha_2 G^{\alpha_2 - 1} q_i - \beta_2\beta_3 G^{\beta_3 - 1}\right) e^{\bar{\mu} - \beta_2 G^{\beta_3} + (\bar{\sigma} + \alpha_1 G^{\alpha_2})q_i}$$

As above, the exponential term will limit to a constant in this case. If $\alpha_2, \beta_3 > 1$, then we have:

$$\lim_{G \to 0^+} \alpha_1\alpha_2 G^{\alpha_2 - 1} q_i - \beta_2\beta_3 G^{\beta_3 - 1} = 0$$

Which implies that:

$$\lim_{G \to 0^+} f_G(a_i, G) = 0.$$

If instead $\alpha_2 < 1$ and $\beta_3 > 1$, the limit will be infinitely positive (as in the case of $\beta_1 \in (0, 1)$). If $\alpha_2 > 1$ and $\beta_3 < 1$, the limit will be infinitely negative. Finally, if both $\alpha_2, \beta_3 < 1$, the limit will

be infinite, with a sign which depends on the relative size of α_2, β_3. If α_2 is smaller, the limit will be positive, and vice versa.

Let $g(\alpha_2, \beta_3)$ map the parameters to the limit of $f_G(a_i, G)$ when $\beta_1 = 0$ as discussed above. Then we have:

$$\lim_{G \to 0^+} f_G(a_i, G) = \begin{cases} g(\alpha_2, \beta_3) & \beta_1 = 0 \\ \infty & \beta_1 \in (0, 1) \\ 0 & \beta_1 > 1 \end{cases}$$

For the purposes of this paper, we are interested in the case $\beta_1 \in (0, 1)$.

D.2: Limit as $G \to \infty$

Now consider the case where G goes to infinity; that is, the case where the state becomes infinitely large. We look for the limit:

$$\lim_{G \to \infty} f_G(a_i, G)$$

As in the case of the limit at 0, this will depend on the values of parameters. To simplify this analysis, we consider only the case of the average agent; that is, the agent for whom $q_i = 0$. The partial of earning ability with respect to G for this agent is then:

$$f_G(a_i, G) = \left(\beta_1 - \beta_2 \beta_3 G^{\beta_3}\right) G^{\beta_1 - 1} e^{\bar{\mu} - \beta_2 G^{\beta_3}}$$

Suppose first that all parameters are zero. Clearly the limit will be zero. This will also be the case if only β_3, $\bar{\mu}$, or both are non-zero.

Next, suppose that only $\beta_1 > 0$, and all other parameters are zero. Then:

$$\lim_{G \to \infty} f_G(a_i, G) = \lim_{G \to \infty} \beta_1 G^{\beta_1 - 1}$$

If $\beta_1 < 1$, then the power on G will be negative and this limit will be zero; if $\beta_1 > 1$, the limit will be infinity. This is also true if any combination of β_3 and $\bar{\mu}$ are non-zero while $\beta_2 = 0$, or if any combination of β_2 and $\bar{\mu}$ are non-zero while $\beta_3 = 0$, as these cases only add extra constants.

Now suppose that only $\beta_2 > 0$, and all other parameters are zero. The limit will be zero, as $\beta_1 - \beta_2 \beta_3 G^{\beta_3}$ will be identically zero in this case. This is also true if $\bar{\mu} \neq 0$.

Now suppose that $\beta_2, \beta_3 > 0$ and $\bar{\mu}$ takes any value, while $\beta_1 = 0$. Then:

$$\lim_{G \to \infty} f_G(a_i, G) = \lim_{G \to \infty} -\left(\beta_2 \beta_3 G^{\beta_3 - 1}\right) e^{\bar{\mu} - \beta_2 G^{\beta_3}}$$

The limit of the second term in the product is zero:

$$\lim_{G \to \infty} e^{\bar{\mu} - \beta_2 G^{\beta_3}} = 0$$

If $\beta_3 < 1$, the limit of the first term is also zero, and the overall limit is zero. If instead $\beta_3 > 1$, we can evaluate the limit using L'Hôpital's rule:

$$\lim_{G \to \infty} f_G(a_i, G) = \lim_{G \to \infty} -\frac{\left(\beta_2 \beta_3 G^{\beta_3 - 1}\right) e^{\bar{\mu}}}{e^{\beta_2 G^{\beta_3}}} = \lim_{G \to \infty} -\frac{\left(\beta_2 \beta_3 (\beta_3 - 1) G^{-1}\right) e^{\bar{\mu}}}{\beta_2 e^{\beta_2 G^{\beta_3}}} = 0$$

Thus the limit is zero regardless of the values of β_2, β_3, as long as they are greater than zero.

Finally, suppose that $\beta_1, \beta_2, \beta_3 > 0$ and $\bar{\mu}$ takes any value. The limit can be decomposed into two terms:

$$\lim_{G\to\infty} f_G(a_i, G) = \lim_{G\to\infty} \left(\beta_1 G^{\beta_1 - 1} e^{\bar{\mu} - \beta_2 G^{\beta_3}} - \beta_2\beta_3 G^{\beta_3+\beta_1-1} e^{\bar{\mu} - \beta_2 G^{\beta_3}}\right)$$

The limit of the exponential function in both terms is zero. If $\beta_3 + \beta_1 < 1$, both power terms have negative exponents and limit to zero, implying that the overall limit is zero. If instead $\beta_3 + \beta_1 > 1$, we can evaluate each term using L'Hôpital's rule. Consider first the first term, assuming that $\beta_1 > 1$ (if $\beta_1 < 1$, this limit is clearly zero). We have for some constant κ:

$$\lim_{G\to\infty} e^{\bar{\mu}} \frac{\beta_1 G^{\beta_1 - 1}}{e^{\beta_2 G^{\beta_3}}} = \lim_{G\to\infty} e^{\bar{\mu}} \frac{\beta_1(\beta_1 - 1) G^{\beta_1 - \beta_3 - 1}}{\beta_2\beta_3 e^{\beta_2 G^{\beta_3}}} = \lim_{G\to\infty} \kappa \frac{G^{\beta_1 - 1 - n\beta_3}}{e^{\beta_2 G^{\beta_3}}} = 0$$

Where the second equality follows from taking the nth derivative of the numerator and denominator (applying L'Hôpital's rule n times). For n large enough, we see that both numerator and denominator go to zero for any positive values of β_3 and β_1.

Now consider the second term, applying L'Hôpital's rule once more:

$$\lim_{G\to\infty} -e^{\bar{\mu}} \frac{\beta_2\beta_3 G^{\beta_3+\beta_1-1}}{e^{\beta_2 G^{\beta_3}}} = \lim_{G\to\infty} -e^{\bar{\mu}} \frac{\beta_2\beta_3(\beta_3 + \beta_1 - 1) G^{\beta_1 - 1}}{\beta_2\beta_3 e^{\beta_2 G^{\beta_3}}} = \lim_{G\to\infty} -\kappa \frac{G^{\beta_1 - 1 - (n-1)\beta_3}}{e^{\beta_2 G^{\beta_3}}} = 0$$

Where the second equality again follows from applying L'Hôpital's rule n times. For n large enough, we again see that both numerator and denominator go to zero for any positive values of β_3 and β_1.

Thus when $\beta_1, \beta_2, \beta_3 > 0$ (regardless of relative magnitudes) and $\bar{\mu}$ takes any value, the limit is:

$$\lim_{G\to\infty} f_G(a_i, G) = 0$$

The results of this section are summarized in the table below.

Table 3: This table collects values of $\lim_{G\to\infty} f_G(a_i, G)$ for various values of parameters, given $q_i = 0$. Note that all of the below limits hold regardless of the value of $\bar{\mu}$.

Param. Value							
	β_1	0	$\in (0,1)$	> 1	$\in (0,1)$	> 1	> 0
	β_2	≥ 0	0	0	≥ 0	≥ 0	> 0
	β_3	≥ 0	≥ 0	≥ 0	0	0	> 0
$\lim_{G\to\infty} f_G(a_i, G)$		0	0	∞	0	∞	0

Given that we focus on the case where $\beta_1 \in (0, 1)$, we see that regardless of the other parameter values:

$$\lim_{G\to\infty} f_G(a_i, G) = 0$$

Appendix E: Parameter Value Impact on q_i^* Threshold

Below we include the graphs describing the impact of each parameter value on the threshold q_i^*, below which agents receive negative tax rates, and the corresponding percentile of the income distribution, $P(a_i^*)$. Again, for intuition, we consider a world without transfers. Recall that:

$$q_i^* = -\frac{\beta_1 - \beta_2\beta_3 G^{\beta_3}}{\alpha_1\alpha_2 G^{\alpha_2}}.$$

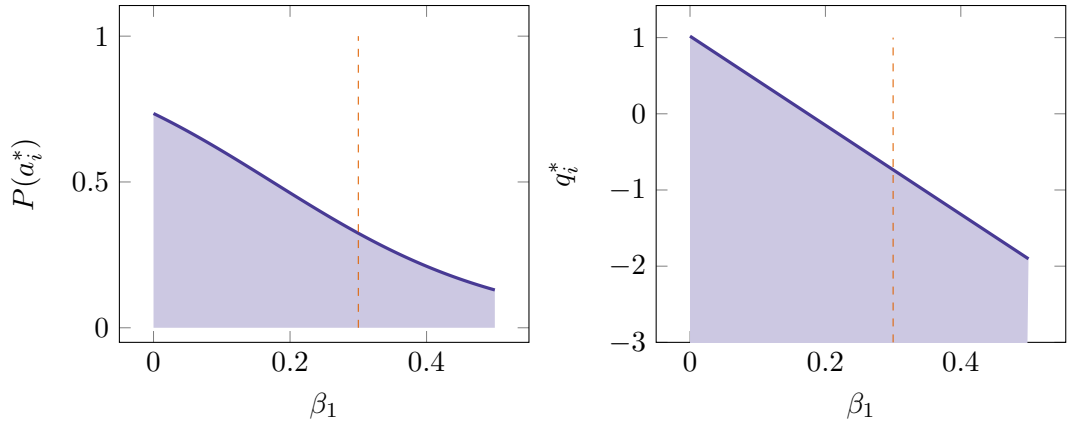


Figure 3: The impact of β_1 on the threshold of ability below which taxes are negative, for values of other parameters fixed at 2019 levels and fixed $G = 100$ (approximately two-thirds of the true 2019 value). The left graph shows the impact of shifts in β_1 on the lowest percentile rank ($P(a_i^*)$) which receives negative taxes, while the right shows the mechanical impact on q_i^*. The higher is β_1, the lower an agent's ability must be ranked to receive negative taxes. The orange dashed lines indicate our chosen parameters for 2019.

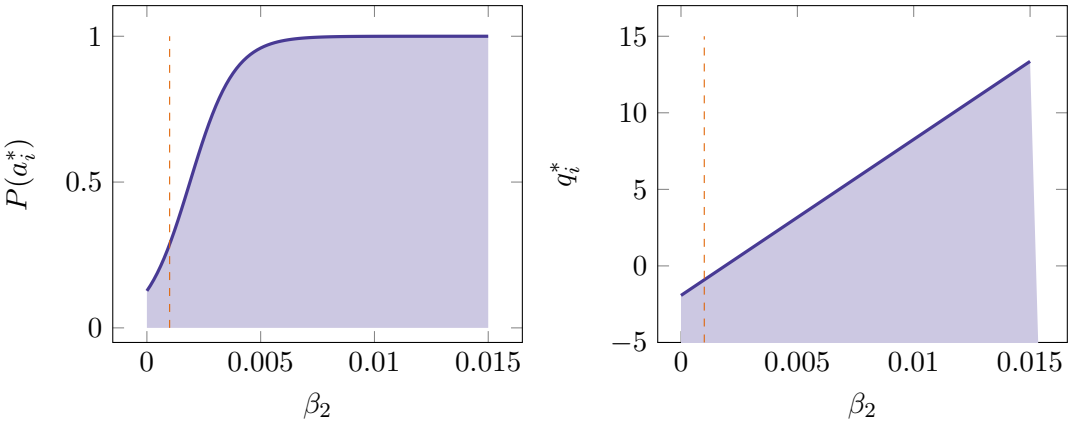


Figure 4: The impact of β_2 on the threshold of ability below which taxes are negative, for values of other parameters fixed at 2019 levels and fixed $G = 100$. The left graph shows the impact of shifts in β_2 on the lowest percentile rank ($P(a_i^*)$) which receives negative taxes, while the right shows the mechanical impact on q_i^*. The higher is β_2, the more agents will receive negative tax rates. The orange dashed lines indicate our chosen parameters for 2019.

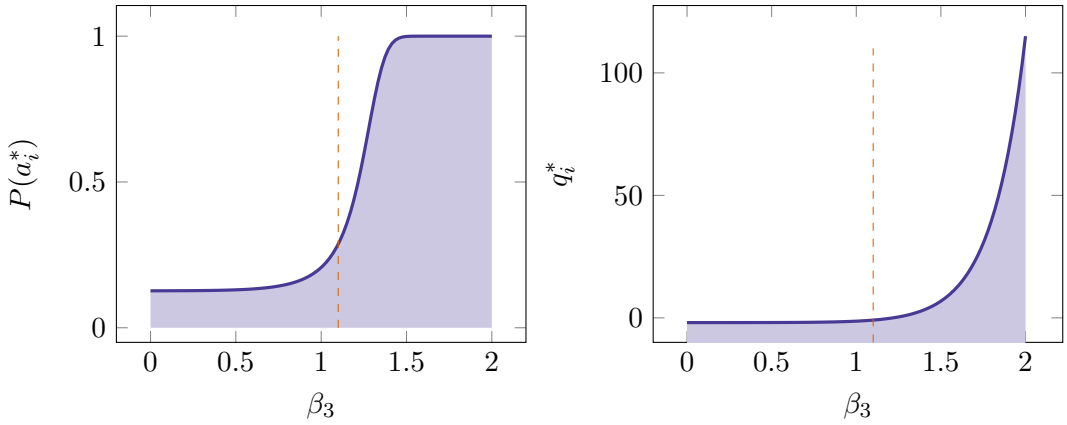


Figure 5: The impact of β_3 on the threshold of ability below which taxes are negative, for values of other parameters fixed at 2019 levels and fixed $G = 100$. The left graph shows the impact of shifts in β_3 on the lowest percentile rank ($P(a_i^*)$) which receives negative taxes, while the right shows the mechanical impact on q_i^*. The higher is β_3, as with β_2, the more agents will receive negative tax rates. The orange dashed lines indicate our chosen parameters for 2019.

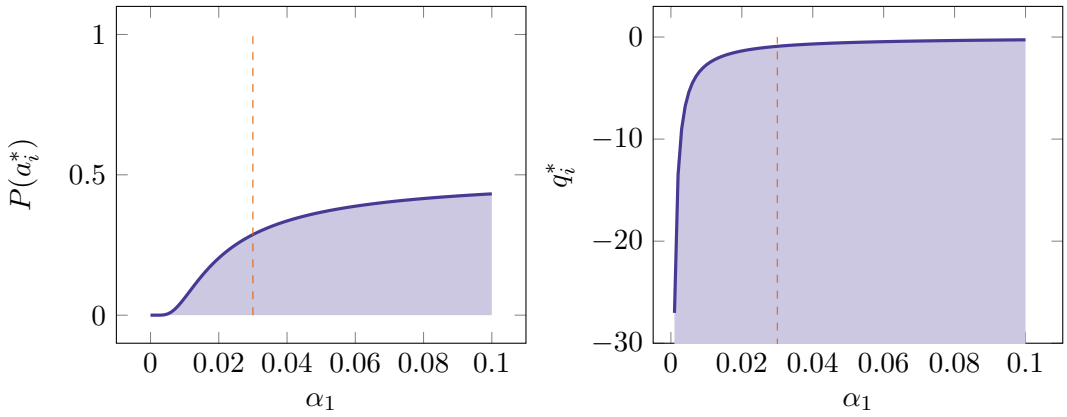


Figure 6: The impact of α_1 on the threshold of ability below which taxes are negative, for values of other parameters fixed at 2019 levels and fixed $G = 100$. The left graph shows the impact of shifts in α_1 on the lowest percentile rank ($P(a_i^*)$) which receives negative taxes, while the right shows the mechanical impact on q_i^*. The higher is α_1, the higher is q_i^*, implying that agents of higher ability will receive negative tax rates—although q_i^* is always below zero in this case, as (given positive $\alpha_1\alpha_2$) only the numerator determines the sign of the threshold. The orange dashed lines indicate our chosen parameters for 2019.

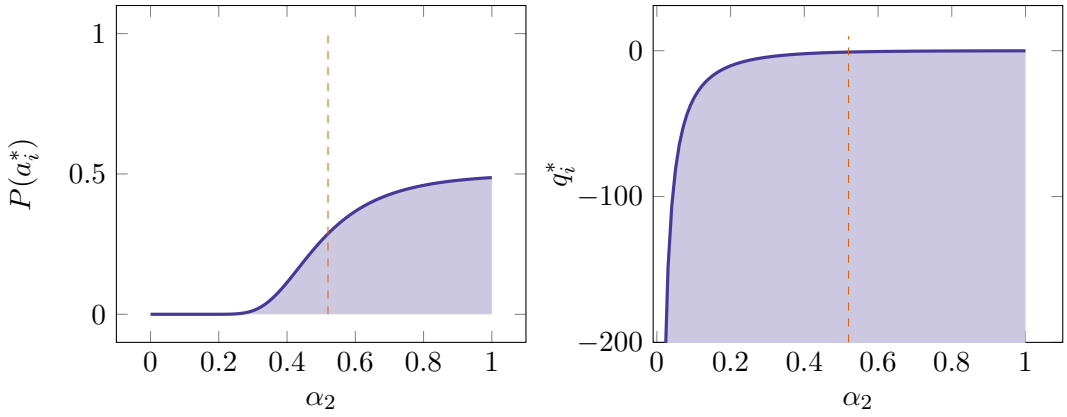


Figure 7: The impact of α_2 on the threshold of ability below which taxes are negative, for values of other parameters fixed at 2019 levels and fixed $G = 100$. The left graph shows the impact of shifts in α_2 on the lowest percentile rank ($P(a_i^*)$) which receives negative taxes, while the right shows the mechanical impact on q_i^*. The higher is α_2, the higher is q_i^*, implying that agents of higher ability will receive negative tax rates, though q_i^* is always below zero in this case, as in the case of varying α_1 above. The orange dashed lines indicate our chosen parameters for 2019.