Equality of Opportunity vs Equality of Outcomes: A Theory of Optimal Predistribution*

Raj Chetty†
Henrik Kleven‡
Damián Vergara§
Danny Yagan¶

August 2026

Abstract

We generalize the Mirrleesian framework to allow for two features: (i) social preferences for equality of opportunity, and (ii) predistribution of labor market skill. The predistribution instrument may be interpreted as early education provision or another pre-market instrument. In our model, income inequality derives from heterogeneity in two dimensions: preferences and family background. The government respects inequality due to preferences (responsibility principle), but not inequality due to family background (compensation principle). Taxes and transfers redistribute income regardless of the source of income differences, whereas education provision predistributes skill based on family background. We provide intuitive policy rules that depend on estimable parameters. Governments should focus on predistribution over redistribution when the earnings elasticity with respect to taxes is large, when the earnings elasticity with respect to education is large, and when a large share of income inequality reflects preference heterogeneity. A calibration analysis suggests an important role for predistribution. Government policy is more progressive than in standard models, but less of the progressivity is achieved through traditional taxes and transfers.

*We thank John Sturm Becko and Jacob Goldin, as well as seminar and conference participants at the Australian National University, the Bank of Italy, GRAPE Warsaw, Harvard University, NBER TAPES, the NTA Annual Meetings, Princeton University, the University of Chicago, the University of Copenhagen, and the University of Michigan for helpful comments and suggestions.

†Harvard University, NBER, and Opportunity Insights. Email: chetty@opportunityinsights.org.
‡Princeton University and NBER. Email: kleven@princeton.edu.
§University of Michigan and NBER. Email: damiv@umich.edu.
¶UC Berkeley and NBER. Email: yagan@berkeley.edu.

1 Introduction

The canonical normative framework in public economics asks how a social planner should use taxes and transfers to redistribute income in a world with heterogeneous skill (Mirrlees 1971). This framework has generated important insights into the design of tax policy and provided tractable formulas for optimal tax rates based on estimable parameters (Diamond 1998; Saez 2001). Income taxes are highly progressive when preferences for income equality are strong and when labor supply elasticities are small. In the extreme, if labor supply elasticities are zero, optimal tax rates equal 100%.

While the Mirrleesian framework is a valuable benchmark, it abstracts from two key features of debates about government policies. First, the planner's objective in the Mirrleesian framework is equality of outcomes: welfare weights depend on marginal utilities of income, not on *why* individuals earn what they earn. In practice, the public appears to prioritize equality of opportunity over equality of outcomes, with greater concern about inequality that arises from circumstances (such as luck or family background) than from choices (such as effort). For example, when asked in the General Social Survey whether they favored promoting equal opportunity or equal outcomes, about 88% of respondents favored policies promoting equal opportunity.[1] Similarly, Alesina, Stantcheva, and Teso (2018) show that concerns about unequal opportunities are broadly held and largely bipartisan in the United States and Western Europe: in each of the countries they study, more than 80% of respondents view unequal opportunities as a problem or a serious problem.

Second, the Mirrleesian framework focuses on a planner whose only instrument is redistribution of market incomes via taxes and transfers. In practice, governments devote substantial resources to reducing pre-market inequality—so-called predistribution—through early education provision, place-based policies, labor market institutions, etc. Contrary to the traditional view, Blanchet, Chancel, and Gethin (2022) show that the lower inequality in Europe relative to the US is driven entirely by predistribution, rather than by differences in tax and transfer systems.

These departures from the Mirrleesian approach have been engaged by two separate literatures. On the normative side, a long tradition in social choice and political philosophy has formalized equality of opportunity around two principles (Roemer 1998; Fleurbaey 2008; Fleurbaey and Maniquet 2011, 2018; Roemer and Trannoy 2015). According to the *compensation principle*, people should be compensated for unequal circumstances such as luck, genes, and family background. According to the *responsibility principle*, people should be held responsible for unequal choices such as effort. While this literature has made important conceptual progress on formalizing non-welfarist values, it has not been closely connected to estimable parameters in the style of the sufficient statistics approach of modern public

[1] This evidence is presented in Figure 1, which we discuss in Section 2.

economics (Chetty 2009; Kleven 2021). Furthermore, it focuses largely on tax and transfer instruments, rather than pre-market instruments that target unequal opportunities directly.

On the empirical side, a large literature in labor economics studies intergenerational mobility and its determinants. While outcomes are generally persistent across generations, the degree of persistence varies substantially across space and is correlated with measures of both family structure and pre-market policies (e.g., Chetty *et al.* 2014, 2026). Moreover, experimental and quasi-experimental studies find sizable effects of early childhood interventions on earnings in adulthood (e.g., Heckman *et al.* 2010; Heckman, Pinto, and Savelyev 2013; Chetty *et al.* 2011*b*; Chetty, Friedman, and Rockoff 2014; Kline and Walters 2016; Barr and Gibbs 2022; Gray-Lobe, Pathak, and Walters 2023). But this literature is not directly tied to a normative theory of optimal policy, leaving it unclear which empirical parameters are of greatest importance to policy design and what the estimates imply for the appropriate mix of redistributive and predistributive instruments.

We develop a normative framework that bridges these different literatures. Our framework departs from the canonical Mirrleesian approach in two ways. First, we consider a government with preferences for equality of opportunity. To separate opportunities from outcomes, we assume that realized earnings depend on preferences for working (denoted by A) and family background (denoted by ω_0), both of which are heterogeneous in the population. The government respects inequality arising from differences in preferences (consistent with the responsibility principle), but not inequality arising from differences in family background (consistent with the compensation principle).[2] These principles are captured by specifying social welfare as a function of individual utilities evaluated at the *average* preferences in the population. This flattens the profile of social marginal welfare weights with respect to earnings, because higher earnings reduce welfare weights only through their correlation with family background.[3] We demonstrate that our social welfare function respects the Pareto principle, which addresses a classic concern about non-welfarist social objectives (Kaplow and Shavell 2001*a*,*b*).

Second, we endogenize labor market skill to depend on pre-market government investments such as early education provision. We will refer to this instrument as "education," but it may capture any policy that affects the mapping between family background ω_0 and realized earnings. Our baseline model assumes that the government can target education provision directly to family background (perfect tagging) and that parental education investments are exogenous to government policy, but both of these assumptions are relaxed in model extensions. These extensions introduce additional fiscal externalities

[2]As discussed below, rather than labeling A and ω_0 as "preferences" and "family background," respectively, we could alternatively label them as meritocratic and unmeritocratic skill. The key assumption is that earnings are driven by two-dimensional heterogeneity, where the government respects one dimension but not the other.

[3]A reduced-form alternative to our approach is to incorporate the equal-opportunity principles directly in the social marginal welfare weights (see e.g., Lockwood and Weinzierl 2015; Saez and Stantcheva 2016). In our approach, the welfare weights are explicitly tied to (re-scaled) individual utilities.

without changing the core insights of our analysis.

The government sets redistribution using income taxes and predistribution using education provision to maximize its equal-opportunity objective. The tax schedule is characterized by a generalized Diamond-Saez rule (Diamond 1998; Saez 2001; Diamond and Saez 2011), while the education schedule is characterized by a generalized Samuelson rule (Samuelson 1954).[4] The two rules are interlinked and must be solved simultaneously.

The optimal tax rule depends on social marginal welfare weights measured at average preferences and on the elasticity of earnings with respect to the marginal net-of-tax rate. The source of earnings inequality is critical for optimal tax progressivity. For example, if earnings inequality is driven entirely by preference heterogeneity, optimal marginal tax rates equal zero at *any* earnings elasticity. When allowing for endogenous parental education investments, the tax rule depends on a macro elasticity that includes both the direct effect of children's supply responses to their marginal tax rate and the indirect effect of parents' investment responses to their children's marginal tax rate.

The optimal education rule equates the social marginal benefit and social marginal cost of education provision at each level of family background. The benefit side is governed by the elasticity of earnings with respect to education spending, scaled by the direct welfare effect of higher skill and a fiscal externality of higher skill. Allowing for endogenous parental investments does not change the optimality condition in a reduced-form sense, but it implies that the relevant education elasticity includes any crowding-out effects from parental responses to government policy. We note that reduced-form estimates of the impact of childhood interventions on adult earnings include such crowding-out and can therefore be used to calibrate our model.

Taking a step back, pre-market government investments play three roles in our framework. First, they provide a targeting advantage by being conditioned on family-background skill ω_0 (perfect tagging) or parental earnings (imperfect tagging). For governments interested in equality of opportunity, this is preferred to income taxes based on current earnings regardless of their source.[5] Second, they generate public-goods benefits as long as government and parental education investments are imperfect substitutes. In other words, the specific education good provided by the government is underprovided by the private market. Third, because education investments are complementary to labor supply, they violate the

[4]The Samuelson rule in our setting accounts for the presence of distortionary taxation and distributional concerns (as in, e.g., Stiglitz and Dasgupta 1971; Atkinson and Stern 1974; Hylland and Zeckhauser 1979; Kaplow 1996, 2004; Kreiner and Verdelin 2012) and for the fact that education is a publicly-provided *private* good.

[5]If income taxes were conditioned on background skill or parental earnings, the targeting advantage of pre-market investments would disappear. In fact, if taxes were allowed to depend on exogenous background skill ω_0, we would be in the first-best world of the second welfare theorem where governments can costlessly redistribute endowments (opportunities). We focus on governments setting taxes based on the current earnings of tax-liable individuals, because this is the universal form of redistribution across all countries. Conversely, predistribution policies such as early education generally depend on parental characteristics rather than on the children's characteristics.

conditions of the Atkinson-Stiglitz theorem (Atkinson and Stiglitz 1976) under which governments should use only income taxes. In general, it is second-best optimal to either subsidize labor supply complements (Christiansen 1984; Saez 2002) or provide them in-kind (Nichols and Zeckhauser 1982).

We provide numerical simulations based on a calibrated version of our model. Specifically, we compare Gini coefficients of the income distribution at the optimal policy and the laissez-faire equilibrium, decomposing the difference into redistribution and predistribution. We conduct comparative statics analyses of how this decomposition changes as the key parameters of the model change. The role of predistribution relative to redistribution is increasing in the earnings elasticity with respect to taxes, the earnings elasticity with respect to education spending, and the share of income inequality driven by preference heterogeneity. In our baseline calibration, roughly a third of the income-equalizing effect of government policy is due to predistribution. Government policy is more progressive than in the Mirrleesian framework (under the same calibration), but much less of the progressivity is achieved through traditional taxes and transfers.

Our framework identifies the share of earnings variation explained by preference heterogeneity as a key statistic for policy design. Related, an empirical literature measures (in)equality of opportunity by decomposing observed inequality into the parts explained by circumstances and effort choices, respectively (e.g., Hufe *et al.* 2017; Hufe, Kanbur, and Peichl 2022). A conceptual challenge of such decompositions is that choices are functions of circumstances, making it difficult to cleanly separate the two. Our framework is not based on a distinction between circumstances and choices per se, but rather between background skill ω_0 and preferences A, both of which are primitives of the model. To separate the two sources of earnings variation, we propose a simple method based on the R^2 of two sequential regressions: a regression of wage income on family-background variables (opportunity regression), and a regression of wage income on both family background and labor supply variables (opportunity+choice regression). The role of preferences is based on the difference between the two R^2-estimates. The sequential procedure implies that the covariance between family background and labor supply is assigned to the effects of family background.[6] To put it differently, the role of preferences is based on the correlation between earnings and labor supply choices above and beyond what can be explained by family-background variables. Using NLSY79 data, we find that at least 18.5% of earnings inequality is due to preferences, with a preferred estimate of 44.7%.

Our paper contributes to a large literature on optimal redistribution, including papers in the Mirrleesian tradition (e.g., Mirrlees 1971; Diamond 1998; Saez 2001; Piketty and Saez 2013) and papers that allow for non-welfarist fairness principles (e.g., Saez and Stantcheva 2016; Fleurbaey and Maniquet

[6]For example, if family background and labor supply variables were *perfectly* correlated, our sequential procedure would assign none of the earnings variation to preference heterogeneity.

2018). To the best of our knowledge, our paper is the first to model the optimal mix of redistribution and predistribution, and to characterize how this policy mix interacts with equal-opportunity concerns. Related to our modeling of social welfare, Lockwood and Weinzierl (2015) develop a model with heterogeneity in both abilities and preferences, characterizing the optimal tax implications of preference-neutral welfare weights. They show that preference heterogeneity reduces redistribution under the condition that preferences for working rise with income. Gelber and Weinzierl (2016) study optimal taxation in an intergenerational model where children's labor market abilities depend on parents' disposable incomes. In their model, income taxes achieve both redistribution and predistribution, thereby raising optimal tax progressivity. Berg, Håvarstein, and Stubhaug (2025) introduce social welfare functions that value reducing income inequality due to circumstances, but not income inequality due to merit. Estimating the role of merit using Norwegian data, they argue that optimal redistribution is substantially lower than in standard models.

A number of papers have investigated the conditions under which public provision of private goods can be justified on redistributive grounds. Arrow (1971) considered public expenditures targeted to individual types in a setting without income taxes or transfers. Subsequent research merged the analysis of public goods provision with the analysis of optimal income taxation (e.g., Blomquist and Christiansen 1995; Boadway and Marchand 1995; Pirttilä and Tuomala 2002), using the two-type model of Stiglitz (1982). The main question in this literature is whether public provision can relax the incentive-compatibility constraint that restricts optimal redistribution. Their results are linked to the literature on taxing/subsidizing specific goods (Atkinson and Stiglitz 1976; Christiansen 1984; Saez 2002) and providing in-kind transfers (Nichols and Zeckhauser 1982).

Finally, a literature has extended the Mirrleesian framework to allow for human capital investments (e.g., Bovenberg and Jacobs 2005; Jacobs 2013; Stantcheva 2017). This body of work considers models where workers acquire human capital to increase their wages, and where the government jointly optimizes income taxes and education subsidies. In such models, income taxes distort the return to education investments and lead to underinvestment. In general, it is socially optimal to subsidize human capital investments to offset these disincentives, leaving redistribution to the income tax instrument. For example, the classic Bovenberg-Jacobs result prescribes making human capital expenses fully tax deductible, which is regressive by itself. In contrast to this literature, our model considers education investments made by parents and the government *pre-market*—there are no education investments made by adult workers. Government education provision creates equity gains by reducing inequality in pre-market skill, depending on the degree to which such provision crowds out parental investments. To summarize, while prior work and our paper both consider human capital investments, they analyze different policy

instruments (subsidies vs public provision), have different timings (post-market vs pre-market), and are motivated by different rationales (efficiency vs equality of opportunity). Hence, their policy implications run in opposite directions: regressive subsidies paired with more redistributive taxes, as opposed to progressive public provision paired with less redistributive taxes.

The rest of the paper is organized as follows. Section 2 presents descriptive evidence on beliefs about equality of opportunity in the United States. Section 3 sets up our baseline model with fixed parental investments and perfect tagging of government investments. Section 4 extends the model to allow for endogenous parental investments and imperfect tagging. Section 5 presents numerical simulations based on empirically-grounded calibrations. Section 6 concludes.

2 Beliefs About Equality of Opportunity

To motivate our normative framework, we present evidence on beliefs about equality of opportunity in the United States. Figure 1 considers two questions included in the General Social Survey (GSS).[7] The first question elicits redistributive preferences: do you favor promoting equal opportunity or equal outcomes? Respondents overwhelmingly support equalizing opportunities (88% in favor) over equalizing outcomes (12% in favor). These views suggest that Americans prefer "opportunity policies" to traditional redistribution via taxes and transfers. Consistent with this, Stantcheva (2021) finds that support for raising taxes—among both Democrats and Republicans—is much lower when the revenue is spent on out-of-work transfers than on better schools for children from low-income families.

The next question elicits beliefs about success: why does a person's life turn out well or poorly? The respondents are asked to rate the importance of luck, genes, early-life advantages, and hard work. The first three may be interpreted as dimensions of opportunity—factors outside the individual's control—while the last (hard work) may be interpreted as a choice. The figure shows the share of respondents who consider each factor either "very important" or "important."[8] Virtually everyone believes that hard work is important for success in life. Sizable shares of respondents believe that opportunity factors are important too, particularly early-life advantages, but most people view hard work as the single most important factor. These patterns may explain the strong preference for equalizing opportunities over

[7] The data are drawn from the 1993 wave of the survey—the only wave in which these specific questions were included. While this information may appear somewhat dated, the GSS includes many related questions about beliefs and redistributive preferences that are repeated across waves. In general, responses to such recurring questions are very stable over time.

[8] The exact wording of the GSS question is the following: "I'm going to read some statements that give reasons why a person's life turns out well or poorly. As I read each one, tell me whether you think it is very important, important, somewhat important, or not at all important for how somebody's life turns out? (i) Such things are decided by God, (ii) Some people are born with better genes than others, (iii) Society gives some people a head start and holds others back, (iv) Some people use their will power and work harder than others, and (v) It's just a matter of chance." In the figure, we drop the first category (God decides) and label the third category (society gives some people a head start) as "early-life advantages."

equalizing outcomes. If hard work is critical for how life turns out, equalizing outcomes through taxes and transfers largely redistributes on the basis of choice-based heterogeneity rather than opportunity-based heterogeneity.

Appendix Figure A.1 considers a differently-worded question about the reasons for success.[9] Here respondents are asked to rate the importance of race, gender, connections, family wealth, well-educated parents, and hard work. As before, almost everyone views hard work as important for success, further underpinning the opposition to redistributing outcomes rather than opportunities. Among the different opportunity factors, having connections and well-educated parents are viewed as most important.

To summarize, the evidence presented here highlights two facts about beliefs: that outcomes are driven by both choices (hard work) and opportunities—with hard work being the single most important factor—and that equalizing opportunities is preferred to equalizing outcomes. Our theoretical framework allows for both of these features.

3 Baseline Model

3.1 Workers

A continuum of mass 1 of workers vary in terms of skill ω and preferences for working A. Consumption is given by $z - T(z)$, where z is earnings and $T(z)$ is a tax schedule. Preferences over consumption and effort are characterized by a quasi-linear utility function:

$$u = z - T(z) - \frac{A\omega}{1+1/\varepsilon}\left(\frac{z}{A\omega}\right)^{1+1/\varepsilon}. \tag{1}$$

The assumption of quasi-linearity avoids income effects on effort and is common in the optimal tax literature (e.g., Diamond 1998; Kleven, Kreiner, and Saez 2009). The parameterization of the second term—disutility of effort—has two features worth highlighting: the disutility term is iso-elastic and it is scaled by $A\omega$. As we shall see, the scaling implies that $A\omega$ can be interpreted as a measure of "earnings capacity." This form of scaling has been used in a number of papers (Kleven, Kreiner, and Saez 2009; Saez 2010; Kleven *et al.* 2025; Kleven 2026), albeit in models without preference heterogeneity operating through A. The parameterization in equation (1) is analytically convenient and useful for interpretation, but it is not critical for any of the fundamental insights presented below.

The maximization of utility with respect to earnings gives the following labor supply function

$$z = A\omega\left(1 - \tau(z)\right)^{\varepsilon}, \tag{2}$$

[9] This question is drawn from the 2021 wave of the GSS and is therefore more recent than the questions considered in Figure 1.

where $\tau(z) \equiv T'(z)$ is the marginal tax rate, and ε is the elasticity of earnings with respect to the marginal net-of-tax rate $1-\tau$. Equation (2) implies that, at a marginal tax rate of zero, workers choose $z = A\omega$. In other words, the product of preferences and skill can be interpreted as laissez-faire earnings or earnings capacity. The imposition of positive marginal tax rates depresses actual earnings z below earnings capacity $A\omega$ according to the elasticity parameter ε.

Indirect utility can be written as $u(A\omega)$. In a standard Mirrleesian framework, the planner's social welfare function is defined over $u(A\omega)$ across workers of different types. In this case, the distinction between A and ω—given they enter symmetrically—would be irrelevant. But we consider a planner who evaluates inequality driven by A and ω differently, respecting the former but not the latter. As described below, our social objective provides a tractable way of capturing the principles of *compensation* and *responsibility* from the literature on equality of opportunity (Fleurbaey and Maniquet 2011; Roemer and Trannoy 2015; Fleurbaey and Maniquet 2018). While we label A as preferences and ω as skill, we could alternatively label them as meritocratic and unmeritocratic skill, respectively. Regardless of labels, the substantive assumption is that earnings inequality is driven by two dimensions of heterogeneity, where the government views only one of them (the "opportunity") as worthy of redistribution. The income tax $T(z)$ does not distinguish between the two, blunting its effectiveness for redistributing in a welfare-relevant manner.

3.2 Skill Production

The skill parameter ω is not a primitive of the model, in contrast to Mirrleesian models where skill reflects innate ability. Instead, we assume that skill depends partly on family background (parental skill) and partly on a government instrument (predistribution). For the sake of concreteness, we refer to the predistribution instrument as early-education policy, but it could reflect other policies that influence the pre-market distribution of labor market skill and thus consumption opportunities.

The production function for skill takes the following form

$$\omega = \Omega(\omega_0, E(\omega_0)) \equiv \widehat{\Omega}(\omega_0). \tag{3}$$

Realized skill ω is governed by two factors. First, there is background skill ω_0 drawn from an exogenous distribution—the effect of family background. Realized skill is non-decreasing in background skill, $\Omega'_{\omega_0} \geq 0$. Second, there are government investments in early education $E(\omega_0)$—the effect of predistribution. These investments are made prior to the labor supply choices in equation (2) and they are allowed to vary flexibly by family background ω_0.[10] Realized skill is also non-decreasing in education, $\Omega'_E \geq 0$. Our

[10] In practice, public education provision depends on family-background variables such as location and parental income,

analytical characterizations do not rely on any specific functional form for $\Omega(.)$, but we assume that, at the optimal education schedule, the total derivative of realized skill ω with respect to background skill ω_0 is non-negative, i.e., $\widehat{\Omega}'(\omega_0) \equiv \Omega'_{\omega_0} + \Omega'_E \cdot E'(\omega_0) \geq 0$. In other words, there is intergenerational persistence in the model, even after accounting for predistribution.

Three points are important for interpretation. First, in the baseline model, the transmission of background skill ω_0 from parents to children is exogenous. Parents do not make private education investments in their children, or at least none that are endogenous to government policies. We provide a generalization with parental education investments in Section 4.1. Second, publicly provided education is modeled as a private good. It is both excludable and rival. This implies that education can be targeted to specific individuals without any productivity spillovers to other individuals. While the baseline model assumes perfect tagging to family background, we provide a generalization with imperfect tagging in Section 4.2.

Finally, our model differs in important ways from prior work extending the Mirrleesian framework to allow for human capital investments (e.g., Bovenberg and Jacobs 2005; Stantcheva 2017). This literature considers models where workers acquire human capital to increase their wages, and where the government may tax or subsidize human capital investments in addition to taxing income. An insight from this literature is that, in general, governments should subsidize human capital investments to offset the disincentives created by redistributive income taxation. This is an efficiency argument. Optimal education subsidies are typically regressive, while the full system of income taxes and education subsidies is progressive.[11] In our baseline model, there are no private education investments—and therefore no income-tax distortion of those investments—but rather public education investments made pre-market. A progressive education schedule creates equity gains by reducing inequality in pre-market skill.[12] With an equal-opportunity objective, the equity gains of early-education provision are larger than the equity gains of income taxes—provided the two are equally progressive—as the former redistributes pre-market skill while the latter redistributes income driven by both pre-market skill and preferences.

which map into a relationship between children's education and parental skill ω_0. We consider a government that chooses a flexible and potentially nonlinear education schedule $E(\omega_0)$ as a function of parental skill.

[11] A classic result is that, in a static and deterministic model of human capital acquisition, the optimal policy makes human capital expenses fully tax deductible (Bovenberg and Jacobs 2005). With progressive marginal income tax rates, education subsidies are regressive.

[12] Education provision has efficiency effects as well, both through fiscal externalities and through public goods benefits. The latter effect arises because, in this model, there is no private market for the education good provided by the government. Even with a private education market (Section 4), there will be public goods benefits as long as private and public education are not perfect substitutes.

3.3 Policy Instruments and Levels of Inequality

The model incorporates four levels of inequality, linked by government policies. The first level is heterogeneity in innate types—background skill ω_0 and preferences for working A. The joint distribution of innate types is given by a density function $f_0(\omega_0|A) f_A(A)$. The second level is heterogeneity in realized types (ω, A), which depends on government investments in education. Given the skill production function (3), the joint distribution of realized types is given by the density function $f_\omega(\omega|A) f_A(A) = f_0\left(\widehat{\Omega}^{-1}(\omega)|A\right) f_A(A)$. The difference between the innate and realized type distributions represents *predistribution*. The third level is heterogeneity in before-tax earnings z, which depends on marginal income tax rates. Given the labor supply function (2), the joint distribution of earnings and preferences is given by $f_z(z|A) f_A(A) = f_\omega\left(\frac{z}{A(1-\tau(z))^\varepsilon}|A\right) f_A(A)$. The final level is inequality in after-tax earnings $z - T(z)$, determined by the distribution of before-tax earnings (which depends on marginal tax rates) and the tax schedule $T(z)$. The difference between the distributions of realized types and after-tax earnings represents *redistribution*.

Below we characterize the trade-off between predistribution (equalizing pre-market opportunities) and redistribution (equalizing post-market outcomes). We provide intuitive formulas that highlight the key mechanisms and depend on estimable sufficient statistics.

3.4 Government

We consider a government with preferences for equality of opportunity. A rich social choice literature studies distributional objectives that incorporate fairness principles beyond the standard utilitarian perspective of public economics (e.g., Roemer 1998; Fleurbaey and Maniquet 2011, 2018; Roemer and Trannoy 2015; Saez and Stantcheva 2016). In this literature, the equal-opportunity ethic is typically based on the principles of compensation and responsibility. According to the *compensation principle*, people should be compensated for unequal circumstances such as luck, genes, and family background. According to the *responsibility principle*, people should be held responsible for unequal choices such as effort. Our social objective is grounded in these ideas, but formalized in a novel way.

The equal-opportunity ethic just described, while intuitive, must address an important challenge: how do we separate circumstances and choices? This is not only an empirical challenge but a conceptual one, because choices are functions of circumstances. In the Mirrleesian framework, workers are born with an innate ability (circumstance) which determines labor supply (choice); there is no meaningful sense in which choices can be separated from circumstances. In fact, this is true of all economic models. They characterize choices in terms of primitives and external factors such as preference parameters, market

prices, government policies, and stochastic processes—none of which are chosen by the individual.[13]

To address these conceptual issues, we specify the equal-opportunity objective in terms of primitives that can be clearly separated: background skill ω_0 and preferences for working A. The social planner respects inequality driven by preferences A (responsibility principle), but not inequality driven by family background ω_0 (compensation principle).[14] As noted above, while we label the two primitives as preferences and background skill, we could alternatively label them as meritocratic and unmeritocratic skill.[15] The fact that they enter the worker's problem symmetrically (via earnings capacity $A\omega$) is analytically convenient, but not important. The key assumption is that earnings capacity depends on two factors, only one of which merits redistribution.[16]

To incorporate the idea that redistribution in the preference dimension has no social value, we assume that the government evaluates individual utilities at the *average* preferences for working in the population. Specifically, the social welfare function is given by

$$\begin{aligned} W &= \int_A \int_{\omega_0} \Psi \left[u \left(z \left(A^* \widehat{\Omega} \left(\omega_0 \right) \right), A^* \widehat{\Omega} \left(\omega_0 \right) \right) \right] dF_0 \left(\omega_0 | A \right) dF_A \left(A \right) \\ &= \int_A \int_z \Psi \left[u \left(\frac{A^*}{A} z, A^* \omega \left(z, A \right) \right) \right] dF_z \left(z | A \right) dF_A \left(A \right), \end{aligned} \tag{4}$$

where $\Psi \left[. \right]$ is an increasing and concave transformation of individual utilities, and $A^* \equiv \int_A A dF_A \left(A \right)$ denotes average preferences for working. The first line of equation (4) is specified in terms of background skill ω_0, while the second line is specified in terms of earnings z. The second line clarifies that, by valuing each individual at average preferences, the planner scales their earnings by a factor A^* / A, treating those with strong (weak) preferences for working as if they had lower (higher) earnings than they actually do.[17]

When setting taxes and education provision, the government must satisfy the following budget constraint

$$\begin{aligned} \int_A \int_{\omega_0} T \left(z \left(A \widehat{\Omega} \left(\omega_0 \right) \right) \right) dF_0 \left(\omega_0 | A \right) dF_A \left(A \right) - p \int_A \int_{\omega_0} E \left(\omega_0 \right) dF_0 \left(\omega_0 | A \right) dF_A \left(A \right) &\geq R \\ \int_A \int_z T \left(z \right) dF_z \left(z | A \right) dF_A \left(A \right) - p \int_A \int_{\omega_0} E \left(\omega_0 \right) dF_0 \left(\omega_0 | A \right) dF_A \left(A \right) &\geq R, \end{aligned} \tag{5}$$

[13] A more philosophical perspective on this point is that, in economic models, there is no free will. Economic agents solve decision-making problems, but these are essentially translations of external factors into choices. The difference between any two agents can be explained entirely by differences in factors outside the agent's span of control.

[14] This definition of equal-opportunity preferences is robust to the absence of free will.

[15] With only two goods—consumption and leisure—the parameter A determines the relative weighting of consumption and leisure in utility. Therefore, preferences for working A could also be labeled as preferences for consuming. In this terminology, the equal-opportunity objective puts a zero value on income inequality driven by heterogeneity in consumption preferences.

[16] Our approach is most closely related to that of Lockwood and Weinzierl (2015).

[17] When expressing social welfare in terms of earnings z, $\omega \left(z, A \right) \equiv \frac{z}{A(1-\tau(z))^{\varepsilon}}$ defines the skill level ensuring that a worker with preferences A chooses z given the labor supply function in equation (2).

where p is the price per unit of education, and R is an exogenous revenue requirement (for financing non-education spending). We denote the Lagrange multiplier associated with the government budget constraint by λ. Note that this multiplier corresponds to the Marginal Value of Public Funds (MVPF).

The government sets the tax schedule $T(z)$ and the education schedule $E(\omega_0)$ to maximize social welfare (4) subject to incentive compatibility constraints (2), the skill production function (3), and the government budget constraint (5). The tax and education schedules may be nonlinear, but they are assumed to be continuously differentiable.

The government's redistributive preferences may be summarized by social marginal welfare weights: the social marginal value of income to different individuals ($\Psi'[.]$) in terms of the marginal value of public funds (λ). These welfare weights can be expressed in terms of background skill ω_0, realized skill ω, or earnings z, using the mapping between the three concepts provided by our model. For characterizing optimal income taxes $T(z)$, it is convenient to express welfare weights in terms of earnings. We define $g(z, A, A^*) \equiv \Psi'\left[u\left(\frac{A^*}{A}z, A^*\omega(z, A)\right)\right]/\lambda$ as the social marginal welfare weight on individuals with earnings z and preferences A. This allows us to specify average welfare weights as follows:

$$g(z, A^*) \equiv \frac{\int_A g(z, A, A^*) f_z(z|A) dF_A(A)}{h(z)} \tag{6}$$

$$G(z, A^*) \equiv \frac{\int_z^{\bar{z}} g(z', A^*) dH(z')}{1 - H(z)}, \tag{7}$$

where $h(z) \equiv \int_A f_z(z|A) dF_A(A)$ and $H(z) \equiv \int_A F_z(z|A) dF_A(A)$ denote the *unconditional* density and distribution functions for earnings.[18] These equations define average social marginal welfare weights at earnings level z ($g(z, A^*)$) and above earnings level z ($G(z, A^*)$), each measured at average preferences A^*. With government preferences for equality, the weights g, G will in general be declining in z. However, because we hold preferences fixed at their average, the weights will decline less steeply than in standard models whenever some of the earnings inequality reflects preference heterogeneity. If earnings inequality were driven solely by preference heterogeneity, the welfare weights would be flat in z. In the optimal policy analysis, the equal-opportunity objective is fully captured by the welfare weights defined above. Our approach is consistent with the argument in Saez and Stantcheva (2016) that different fairness principles can be captured by generalized social marginal welfare weights. A difference between the two approaches is that, in our framework, the welfare weights are explicitly tied to individual utilities.

A concern with social objectives that include fairness principles beyond those based on individual utilities is that they are non-Paretian (Kaplow and Shavell 2001*a*,*b*). An advantage of the social welfare

[18] In equation (7), we denote the upper bound of the earnings distribution by $\bar{z}$. With an infinite upper tail, we have $\bar{z} = \infty$.

function in equation (4) is that it respects the Pareto principle:

Remark 1 (Social Welfare Function is Paretian). *Given a distribution of primitives (A, ω_0), any set of individual-specific utility perturbations such that no individual is worse off (better off) and at least one individual is strictly better off (worse off) implies a strictly larger (smaller) value of the social welfare function W defined in equation (4). Because of quasi-linear utility, these utility perturbations correspond to individualized lump-sum taxes or transfers based on type.*

Proof. See Appendix B. □

Hence, even allowing for individual lump-sum instruments—the most flexible policy instruments conceivable—our social welfare function will not generate non-Paretian policy recommendations. This implies that the policy rules derived in this paper respect the Pareto principle.

It is worth noting that there is a trivial sense in which non-Paretian results are achievable, namely if we allow the planner to manipulate the structure of primitives. For example, suppose all heterogeneity in earnings capacity is due to family background ω_0 rather than preferences A. If the planner shifts heterogeneity from ω_0 to A such that all individual utilities are kept constant, the value of social welfare W increases. Then, if the planner destroys a small amount of A for one individual, social welfare remains higher than in the initial equilibrium even though the change in allocation is a Pareto worsening. More generally, we show in Appendix B that *any* change in allocation associated with non-Paretian effects must involve a reallocation of the primitives of the model. Such non-Paretian effects have no policy relevance as they are not grounded in any instrument the government could conceivably have.[19]

3.5 Optimal Policy

We first characterize the optimal tax schedule (redistribution of income), taking the education schedule (predistribution of skill) as given. We solve for the optimal marginal tax rates using a perturbation approach (Saez 2001). This approach is based on the idea that, at the optimum, small changes in the tax schedule have no first-order effects on welfare. Specifically, consider a small reform which increases the marginal tax rate by $d\tau$ on an earnings interval $(z, z + dz)$. Such a reform has three welfare effects: a mechanical revenue effect dM_τ, a behavioral revenue effect dB_τ, and a direct welfare effect dW_τ. We derive each of these effects below, characterizing the optimal tax schedule based on the condition $dM_\tau + dB_\tau + dW_\tau = 0$.

[19] To our knowledge, this distinction between non-Paretian effects from policy-induced vs primitive-induced allocation changes is novel in the social choice literature. We argue that the most appropriate definition of Paretian objectives should be anchored in conceivable policy instruments.

The mechanical revenue effect (evaluated at the marginal value of public funds λ) is given by

$$dM_\tau = \lambda \cdot d\tau \cdot dz \cdot (1 - H(z)). \tag{8}$$

This expression captures that all workers above earnings level z (with mass $1 - H(z)$) experience a tax liability increase of $d\tau \cdot dz$, taking behavior as given.

The behavioral revenue effect (also evaluated at the marginal value of public funds λ) is given by

$$\begin{aligned} dB_\tau &= \lambda \cdot \tau(z) \cdot dz \cdot h(z)\, dz \\ &= -\lambda \cdot \frac{\tau(z)}{1 - \tau(z)} \cdot \varepsilon \cdot z \cdot d\tau \cdot h(z)\, dz, \end{aligned} \tag{9}$$

where we use the earnings supply function (2). This expression captures that workers in the affected range (with mass $h(z)\, dz$) reduce earnings by dz and therefore tax liability by $\tau(z) \cdot dz$. This is the fiscal externality—or efficiency effect—from behavioral responses.

Finally, the direct welfare effect can be written as

$$\begin{aligned} dW_\tau &= -\int_A \int_z^{\overline{z}} \Psi'[.] \cdot d\tau \cdot dz \cdot dF_z(z'|A)\, dF_A(A) \\ &= -\lambda \cdot d\tau \cdot dz \cdot (1 - H(z)) \cdot G(z, A^*) \\ &= -dM_\tau \cdot G(z, A^*), \end{aligned} \tag{10}$$

where we use the definition of $G(z, A^*)$ in equation (7). These derivations allow us to characterize the optimal tax schedule as follows:

Proposition 1 (Optimal Redistribution). *For any given education schedule $E(\omega_0)$, the optimal marginal tax rate at each earnings level z satisfies*

$$\frac{\tau(z)}{1 - \tau(z)} = \frac{1 - G(z, A^*)}{\varepsilon \cdot \alpha(z)}, \tag{11}$$

where ε is the elasticity of earnings with respect to $1 - \tau$, $\alpha(z) \equiv zh(z) / (1 - H(z))$ denotes the local Pareto parameter of the earnings distribution, and $G(z, A^*)$ denotes the average social marginal welfare weight above earnings level z, measured at average preferences A^*, as defined in equation (7).

Proof. The result in equation (11) follows from the optimality condition $dM_\tau + dB_\tau + dW_\tau = 0$, using the characterization of each term in equations (8)-(10). We call $\alpha(z)$ a local Pareto parameter because, for an exact Pareto distribution with parameter a, we have $\alpha(z) = a$. □

Formally, the only difference between our optimal tax rule and the Mirrlees-Diamond-Saez rule (Mirrlees 1971; Diamond 1998; Saez 2001) lies in the social marginal welfare weight $G(z, A^*)$. Our model

implies a different profile of welfare weights because of two features: (*i*) earnings inequality is driven by heterogeneity in both background skill and preferences, and (*ii*) the government has an equal-opportunity objective that respects inequality driven by preferences (responsibility principle). All else equal, this flattens the profile of welfare weights, reducing the progressivity of marginal tax rates.[20]

A second difference between our model and the Mirrleesian model—not directly visible in the reduced-form representation of equation (11)—is the presence of a predistribution instrument. This instrument allows the government to reduce inequality in pre-market skill and therefore in earnings, feeding back into the optimal tax rule via $\alpha(z)$ and $G(z, A^*)$. This effect further reduces the progressivity of marginal tax rates. Importantly, while our model limits the desirability of traditional redistribution, optimal policy is not necessarily less progressive overall due to the role of predistribution.

We now turn to the characterization of the optimal education schedule (predistribution). As with taxes, the derivations are based on a perturbation approach. Consider a small reform which increases education by dE on a skill interval $(\omega_0, \omega_0 + d\omega_0)$. The reform has three welfare effects: a mechanical revenue effect dM_E, a behavioral revenue effect dB_E, and a direct welfare effect dW_E. At the social optimum, we must have $dM_E + dB_E + dW_E = 0$.

The mechanical revenue effect is given by

$$dM_E = -\lambda \cdot p \cdot dE \cdot \int_A f_0(\omega_0|A)\, d\omega_0 dF_A(A). \tag{12}$$

The behavioral revenue effect is given by

$$dB_E = \lambda \cdot \Omega'_E \cdot dE \cdot \int_A \tau(z) \frac{dz}{d\omega} \cdot f_0(\omega_0|A)\, d\omega_0 dF_A(A). \tag{13}$$

This expression characterizes the fiscal externality that arises because affected workers increase earnings and therefore tax liability (by $\tau(z)\frac{dz}{d\omega}$) when their skill increases (by $\Omega'_E(.) \cdot dE$) due to education provision.

Finally, the direct welfare effect is given by

$$\begin{aligned} dW_E &= \Psi'[.] \cdot \frac{du}{d\omega} \cdot \Omega'_E \cdot dE \cdot \int_A f_0(\omega_0|A)\, d\omega_0 dF_A(A) \\ &= \lambda \cdot g(\omega_0, A^*) \cdot \frac{du}{d\omega} \cdot \Omega'_E \cdot dE \cdot \int_A f_0(\omega_0|A)\, d\omega_0 dF_A(A), \end{aligned} \tag{14}$$

[20] These remarks highlight that, in our model, the source of earnings inequality is critical for optimal taxation. In the special case where earnings inequality is driven entirely by preference heterogeneity, all workers get the same welfare weight g^*. Because the average welfare weight in the population equals one (this follows from the optimization of the intercept $T(0)$), we obtain $G(z, A^*) = g^* = 1$ and therefore $\tau(z) = 0$ at any z. This holds even as the earnings elasticity ε approaches zero. In the limit where both $G(z, A^*) = 1$ and $\varepsilon = 0$, social welfare is independent of marginal tax rates: they have neither equity benefits nor efficiency costs. In this case, the optimal tax schedule becomes indeterminate.

where $g(\omega_0, A^*) \equiv \Psi' \left[u\left(z\left(A^* \widehat{\Omega}(\omega_0) \right), A^* \widehat{\Omega}(\omega_0) \right) \right] / \lambda$ denotes the social marginal welfare weight on individuals with background skill ω_0, measured at average preferences A^*.

With these derivations, the optimal education schedule can be characterized as follows:

Proposition 2 (Optimal Predistribution). *For any given tax schedule $T(z)$, the optimal education provision at each level of background skill ω_0 satisfies*

$$\Omega'_E \cdot \left[g(\omega_0, A^*) \frac{du}{d\omega} + \mathbb{E}\left[\tau(z) \frac{dz}{d\omega} \mid \omega_0 \right] \right] = p, \tag{15}$$

where $\frac{du}{d\omega} = A^* (1 - \tau(z))^{1+\varepsilon} / (1+\varepsilon)$ denotes the marginal willingness to pay for higher skill, and $g(\omega_0, A^*)$ denotes the social marginal welfare weight at background skill ω_0, measured at average preferences A^*.

Proof. At the social optimum, we have $dM_E + dB_E + dW_E = 0$. Inserting equations (12)-(14) and rearranging terms gives equation (15). The formula for the marginal willingness to pay $\frac{du}{d\omega}$ follows from the utility specification (1), applying the labor supply function (2) (envelope theorem). □

Education provision is set to ensure that, at each skill level, the social marginal benefit (left-hand side of equation (15)) equals the social marginal cost (right-hand side of equation (15)). The social marginal benefit depends on the skill-effect of education, Ω'_E, scaled by two terms: a direct welfare effect of higher skill (first term) and a fiscal externality of higher skill (second term). The result relates to the Samuelson rule for public goods provision (Samuelson 1954) in settings with distortionary taxation and distributional concerns (e.g., Stiglitz and Dasgupta 1971; Atkinson and Stern 1974; Hylland and Zeckhauser 1979; Kaplow 1996, 2004; Kreiner and Verdelin 2012). Two differences between the optimal education rule (15) and standard results are worth noting. First, our policy rule considers public provision of a private good (excludable and rival). The social marginal benefit of providing education to a specific group of individuals includes effects only for that group, as opposed to effects in the entire population. Second, our policy rule allows for education to be targeted based on family background ω_0 (tagging). The education schedule $E(\omega_0)$ may be progressive and nonlinear.

An alternative interpretation of Proposition 2 notes that education provision is a form of in-kind transfer, which increases skill and is complementary to labor supply. Hence, our result also relates to the literature on in-kind transfers (Nichols and Zeckhauser 1982) or subsidies to specific goods (Atkinson and Stiglitz 1976; Christiansen 1984; Saez 2002). This literature shows that labor supply complementarity creates an argument for in-kind transfers or subsidization. In our model, the degree of complementarity between education and labor supply is governed by Ω'_E. If $\Omega'_E = 0$, equation (15) implies zero education provision.

Because of the government's equal-opportunity objective, the source of earnings inequality is critical for both redistribution and predistribution. It is useful to consider the special case where earnings inequality is driven entirely by preference heterogeneity:

Corollary 1 (Pure Preference Heterogeneity). *Suppose earnings inequality is driven only by preference heterogeneity. In this case, all workers have the same social marginal welfare weight* $(= 1)$ *such that* $G(z, A^*) = 1$. *The optimal tax schedule has* $\tau(z) = 0$ *at any* z, *and the optimal education schedule at background skill* ω_0 *satisfies*

$$\frac{A^* \cdot \Omega'_E}{1+\varepsilon} = p. \tag{16}$$

If skill production is separable in family background ω_0 *and education* E (*i.e.*, $\Omega''_{E,\omega_0} = 0$), *education provision is constant across workers. This corresponds to universal public education, financed by a lump-sum tax,* $T(z) = \overline{T}$.

In this special case, education provision is motivated solely by efficiency considerations as there is no welfare-relevant inequality.[21] Still, providing a constant level of education reduces inequality (as the effect is *proportionally* larger at the bottom), while the lump-sum tax used for financing increases inequality. In other words, predistribution is progressive and redistribution is regressive, but their use is based on efficiency concerns alone.

For calibration purposes, we note that Ω'_E relates to a parameter estimated in the empirical literature. Using the earnings supply function (2), the elasticity of earnings with respect to education spending is given by

$$\zeta \equiv \frac{dz/z}{dE/E} = \Omega'_E \cdot \frac{E}{\omega}. \tag{17}$$

In general, the education elasticity ζ varies by the level of skill and therefore earnings, depending on the shape of $\Omega(.)$. For example, if $\Omega(.)$ is separable in background skill and education (such that Ω'_E is independent of skill), ζ is a decreasing function of ω.[22] Equation (17) shows that, in this case, ζ goes from infinity ($\omega = 0$) to zero ($\omega = \infty$) as skill rises. In the simulation analysis (Section 5), we calibrate the parameters of $\Omega(.)$ to match empirical estimates of the average education elasticity $\mathbb{E}[\zeta]$. Intuitively, the optimal mix of predistribution and redistribution depends on the relative magnitudes of tax and education elasticities, ε and $\mathbb{E}[\zeta]$.

To summarize, the optimal design of tax and education policies—redistribution and predistribution—is

[21] The education instrument has efficiency effects (even absent fiscal externalities), because it compensates for a missing market: our baseline model assumes that households cannot purchase the same education good in the private market. Even with a private education market (Section 4), unless public and private education are perfect substitutes, there will be efficiency effects of public education.

[22] In practice, arguments can be made for either complementarity or substitutability between inputs in skill production. Children with high family-background skill may be better able to acquire knowledge through schooling (implying $\Omega''_{E,\omega_0} > 0$), but they also have less to learn in school (implying $\Omega''_{E,\omega_0} < 0$). If these mechanisms offset each other, the reduced form $\Omega(.)$ would be separable in family-background skill ω_0 and education investments E.

governed by equations (11) and (15). These policy rules provide clear intuition and highlight the central empirical parameters, but they are not closed-form solutions and they interact in complicated ways. To make explicit policy statements, the optimality conditions must be solved simultaneously using calibrated parameters and numerical simulations.

4 Extensions

4.1 Allowing for Parental Education Investments

In our baseline model, parents transmit background skill ω_0 to their children in an exogenous manner. This intergenerational transmission may require parental education investments, but these were implicit and exogenous to government policy and children's behavior. We now generalize the model to allow parents—generation zero—to purchase an education good that increases the earnings capacity of their children. We choose units such that the prices of private education (E_p) and government education (E_g) are both equal to p.

The utility of parents with skill ω_0 is specified as

$$u_0 = z_0 - T_0(z_0) - pE_p - \frac{A_0\omega_0}{1+1/\varepsilon}\left(\frac{z_0}{A_0\omega_0}\right)^{1+1/\varepsilon} + \kappa \cdot u(A\omega), \tag{18}$$

where $u(A\omega)$ is the children's indirect utility and κ is the weight placed by parents on their children. A value of $\kappa = 1$ ($\kappa = 0$) corresponds to full (zero) internalization of the children's welfare. As shown above, the impact of higher skill on children's utility equals $\frac{du(A\omega)}{d\omega} = A(1-\tau(z))^{1+\varepsilon} / (1+\varepsilon)$. Skill production takes the following form

$$\omega = \Omega(\omega_0, E_p, E_g), \tag{19}$$

with derivatives $\Omega'_p, \Omega'_g \geq 0$ and $\Omega''_{pp}, \Omega''_{gg} \leq 0$. We impose no restrictions on the sign of the cross-derivative, Ω''_{pg}, allowing private and public investments to be either technological complements or substitutes.

Parents choose E_p to maximize utility (18) subject to skill production (19), having perfect information about government policies in the children's generation. Parental education investments satisfy the following first-order condition

$$\left[\frac{\kappa A(1-\tau(z))^{1+\varepsilon}}{1+\varepsilon}\right] \cdot \Omega'_p = p, \tag{20}$$

where the left-hand side measures the marginal benefit of parental education investments: the impact on children's skill, Ω'_p, scaled by the parents' marginal willingness to pay for their children's skill. Note that, as $\kappa \to 0$, the left-hand side goes to zero and the first-order condition has no interior solution with $E_p > 0$.

In this case, parents choose $E_p = 0$ and all of our previous results apply.

From equation (20), we may characterize how parents respond to changes in government policy. We obtain

$$\frac{dE_p}{d\tau} = \frac{\Omega'_p}{\Omega''_{pp}} \cdot \frac{1+\varepsilon}{1-\tau} \tag{21}$$

$$\frac{dE_p}{dE_g} = -\frac{\Omega''_{pg}}{\Omega''_{pp}}. \tag{22}$$

The sign of $dE_p/d\tau$ is negative—parents reduce their education investments when a higher tax rate makes those investments less valuable to the children—while the sign of dE_p/dE_g is determined by Ω''_{pg}. When private and public investments are independent in skill production, $\Omega''_{pg} = 0$, parents do not respond to changes in government education provision.

We consider a government maximizing social welfare in the children's generation, taking policies in the parents' generation as given. Hence, the government's optimization problem takes the same general form as before, reinterpreting the revenue requirement R as being net of any leftover revenue from generation zero. The only difference between the optimization problems lies in the presence of the additional incentive compatibility constraint (20). This constraint implies that, whenever the government changes taxes or education provision for the children, parents change their education investments according to equations (21)-(22).

Consider first the optimal tax schedule. As before, we solve for the optimal schedule using a perturbation approach, increasing the marginal tax rate by $d\tau$ on an earnings interval $(z, z+dz)$. The three welfare effects dM_τ, dB_τ, and dW_τ characterized in equations (8)-(10) continue to apply, but there is an additional behavioral effect, dB^p_τ, because of parental investment responses to a higher tax rate on their children. This effect may be written as

$$\begin{aligned} dB^p_\tau &= \lambda \int_A \tau(z) \cdot \frac{dz}{d\omega} \cdot \frac{d\omega}{dE_p} \cdot dE_p \cdot f_z(z|A) \cdot dz \cdot dF_A(A) \\ &= -\lambda \cdot \frac{\tau(z)}{1-\tau(z)} \cdot \varepsilon_p \cdot z \cdot d\tau \cdot h(z) \cdot dz, \end{aligned} \tag{23}$$

where ε_p denotes an indirect earnings elasticity with respect to $1-\tau$, operating through parental education investments. It is defined as follows:

$$\varepsilon_p \equiv \mathbb{E}\left[\frac{dz/z}{dE_p/E_p} \frac{dE_p/E_p}{d(1-\tau)/(1-\tau)} \mid z\right] = -\mathbb{E}\left[\frac{(\Omega'_p)^2}{\Omega''_{pp}} \frac{1+\varepsilon}{\omega} \mid z\right] > 0. \tag{24}$$

The relationship between the reduced-form elasticity ε_p and model primitives—as shown in the last

expression on the right-hand side—uses equations (2), (19), and (21).

The optimal tax schedule may be characterized as follows:

Proposition 3 (Optimal Redistribution with Parental Investments). *For any given education schedule $E(\omega_0)$, the optimal marginal tax rate at each earnings level z satisfies*

$$\frac{\tau(z)}{1-\tau(z)} = \frac{1-G(z,A^*)}{(\varepsilon+\varepsilon_p)\,\alpha(z)}, \tag{25}$$

where ε is the direct earnings elasticity with respect to $1-\tau$ through children's supply responses, and ε_p is the indirect earnings elasticity with respect to $1-\tau$ through parents' education responses, defined in equation (24). The sum $\varepsilon+\varepsilon_p$ represents the macro elasticity of earnings accounting for parental education investments. As above, $\alpha(z) \equiv zh(z)/(1-H(z))$ denotes the local Pareto parameter, and $G(z,A^*)$ denotes the average social marginal welfare weight above earnings level z, measured at average preferences.

Proof. Equation (25) follows from the optimality condition $dM_\tau + dB_\tau + dB_\tau^p + dW_\tau = 0$, using the characterizations provided in equations (8)-(10) and (23). □

The optimal tax rule takes the same general form as in Proposition 1, but specified in terms of the macro elasticity of earnings, $\varepsilon+\varepsilon_p$, rather than the micro elasticity. The macro elasticity accounts for both the direct supply effect of taxes (ε) and the indirect effect operating through parental education investments (ε_p). Parental investment responses increase the earnings elasticity and therefore reduce optimal marginal tax rates, all else equal.

Let us turn to the optimal education schedule. We consider a small reform which increases public education provision by dE_g on the skill interval $(\omega_0, \omega_0+d\omega_0)$. The reform has the same conceptual effects as in the baseline model, but the characterizations need modification to account for parental education investments. Specifically, the impact of public education on skill includes both a direct effect and an indirect effect from parental responses. Using equation (22), we have

$$\frac{d\omega}{dE_g} = \Omega_g' + \Omega_p' \frac{dE_p}{dE_g} = \Omega_g' - \frac{\Omega_{pg}''\Omega_p'}{\Omega_{pp}''}, \tag{26}$$

where the second term captures the indirect effect from changes in parental investments. This term will be negative if $\Omega_{pg}'' < 0$ (crowding-out) and positive if $\Omega_{pg}'' > 0$ (crowding-in).

The three welfare effects of the perturbation to the education schedule can be characterized as follows:

$$dM_E^p = -\lambda \cdot p \cdot dE_g \cdot \int_A f_0(\omega_0|A)\, d\omega_0 dF_A(A) \tag{27}$$

$$dB_E^p = \lambda \cdot \left(\Omega_g' - \frac{\Omega_{pg}''\Omega_p'}{\Omega_{pp}''}\right) \cdot dE_g \cdot \int_A \tau(z) \frac{dz}{d\omega} \cdot f_0(\omega_0|A)\, d\omega_0 dF_A(A) \tag{28}$$

$$dW_E^p = \lambda \cdot g(\omega_0, A^*) \cdot \frac{du}{d\omega} \cdot \left(\Omega_g' - \frac{\Omega_{pg}''\Omega_p'}{\Omega_{pp}''}\right) \cdot dE_g \cdot \int_A f_0(\omega_0|A)\, d\omega_0 dF_A(A). \tag{29}$$

These expressions are similar to those derived in the baseline model, except for the term related to Ω_{pg}''. This term captures the potential crowding-out (or crowding-in) effects of public education provision. With these derivations, we may characterize the optimal education schedule:

Proposition 4 (Optimal Predistribution with Parental Investments). *For any given tax schedule $T(z)$, the optimal education provision at each level of background skill ω_0 satisfies*

$$\left(\Omega_g' - \frac{\Omega_{pg}''\Omega_p'}{\Omega_{pp}''}\right) \cdot \left[g(\omega_0, A^*) \frac{du}{d\omega} + \mathbb{E}\left[\tau(z) \frac{dz}{d\omega} \mid \omega_0\right]\right] = p, \tag{30}$$

where the first term in parentheses is the effect of education provision on skill, net of any crowding-out effects, as characterized in equation (26). As above, $\frac{du}{d\omega} = A^*(1-\tau(z))^{1+\varepsilon}/(1+\varepsilon)$ denotes the marginal willingness to pay for higher skill, and $g(\omega_0, A^*)$ denotes the social marginal welfare weight at background skill ω_0, measured at average preferences A^*.

Proof. At the social optimum, we have $dM_E^p + dB_E^p + dW_E^p = 0$. Inserting equations (27)-(29) and rearranging terms gives equation (30). □

The optimal education rule ensures that the social marginal benefit of education provision equals the social marginal cost. The only difference relative to the baseline model (Proposition 2) lies in the first term in parentheses, which now accounts for the potential crowding-out or crowding-in effects of government provision.[23] Although this term is a complicated function of the derivatives of the Ω-function, it relates to an empirical reduced-form parameter: the elasticity of earnings with respect to government education spending. Using equations (2) and (26), we have

$$\zeta \equiv \frac{dz/z}{dE_g/E_g} = \left(\Omega_g' - \frac{\Omega_{pg}''\Omega_p'}{\Omega_{pp}''}\right) \cdot \frac{E_g}{\omega}. \tag{31}$$

[23] Note that, if parental and government education investments are perfect substitutes, Ω is a function of $E_p + E_g$. In this case, we have $\Omega_p' = \Omega_g'$ and $\Omega_{pg}'' = \Omega_{pp}''$ such that the term $\Omega_g' - \Omega_{pg}''\Omega_p'/\Omega_{pp}''$ on the left-hand side of equation (30) is equal to zero. This corresponds to a scenario with full crowd-out, in which case optimal government education provision (predistribution) is zero. Empirical evidence on the effects of education interventions and other predistribution policies is inconsistent with full crowd-out.

Using this relationship, the education rule in equation (30) is, in a reduced-form sense, identical to the baseline rule in equation (15). The reason is that parental education responses matter only through their implications for the total earnings elasticity with respect to government provision, ζ. Empirically, this elasticity can be inferred from estimates provided in the literature on the impact of childhood education interventions on adult earnings. Any crowd-out effects from parental responses will be subsumed in reduced-form estimates of earnings effects in adulthood. As noted above, ζ is heterogeneous across earnings levels. In our simulation analysis, we calibrate $\Omega(.)$ to match empirical estimates of the average education elasticity $\mathbb{E}[\zeta]$ in the initial equilibrium.

To conclude, Propositions 3-4 show that parental investment responses can be captured by reinterpreting the reduced-form elasticities that govern optimal policy. Optimal redistribution depends on a macro elasticity of earnings with respect to taxes ($\varepsilon + \varepsilon_p$), whereas optimal predistribution depends on an elasticity of earnings with respect to government education spending that includes offsetting effects from crowding-out (ζ). However, while the optimal policy rules are unchanged in a reduced-form sense, the chosen policy parameters are not unchanged. The reduced-form elasticities are endogenous variables, and the policy rules are functions of their values *at the optimum*. For this reason, as emphasized by Kleven (2021), optimal policy analyses are not true sufficient statistics approaches. Even if expressed in terms of elasticities, they require global knowledge of supply and demand functions.

4.2 Allowing for Imperfect Tagging

The baseline model assumes that education provision $E(\omega_0)$ can be targeted based on exogenous family-background type ω_0. This may exaggerate the power of the predistribution instrument relative to taxes and transfers. In reality, family background is not perfectly observable and targeting may have to be based on choice variables. While imperfect tagging can be modeled in different ways, a natural scenario to consider is one where predistribution is based on parental income z_0 rather than parental type ω_0. To separate the role of imperfect tagging from the role of parental education investments, we introduce imperfect tagging in the baseline model without parental investments.

The utility function of both parents and children is given by equation (1), but we allow for different parameters in the two generations. Denoting parents by subscript zero, their earnings choices are characterized by $z_0 = A_0\omega_0 (1 - \tau_0(z_0))^{\varepsilon_0}$. This condition implies a mapping between parental skill and parental earnings, which we denote by $\omega_0 = \omega_0(z_0, A_0)$.

Parents transmit labor market skill but not preferences to their children. The government provides education to children based on parental income, which depends on a combination of parental skill and

preferences. Hence, skill production in the children's generation can be written as

$$\omega = \Omega\left(\omega_0, E\left(z_0\right)\right) = \Omega\left(\omega_0\left(z_0, A_0\right), E\left(z_0\right)\right) \equiv \widehat{\Omega}\left(z_0, A_0\right), \tag{32}$$

where we use the aforementioned link between parental skill ω_0 and parental income z_0. The link depends on parental preferences A_0: if education provision is targeted to children from low-income families, this may reflect either low parental skill or low parental preferences for working. The fact that income-targeting partly captures parental preferences is the reason why tagging is imperfect in this model.

As above, we consider a government maximizing social welfare in the children's generation, taking policies in the parents' generation as given. The social welfare function may be specified in terms of parental income z_0 and parental preferences A_0, i.e.

$$W = \int_{A_0}\int_{z_0} \Psi\left[u\left(A^*\widehat{\Omega}\left(z_0, A_0\right)\right)\right] dF_{z_0}\left(z_0|A_0\right) dF_{A_0}\left(A_0\right). \tag{33}$$

Social welfare depends on *average* preferences for working in the children's generation (A^*) and *individual* preferences for working in the parents' generation (A_0). The former reflects that the government applies the responsibility principle to preference heterogeneity, while the latter reflects that the level of background skill ω_0 affected by income-based education provision $E\left(z_0\right)$ depends on actual parental preferences.

The final piece of the government's problem is the budget constraint, which may be written as

$$\begin{aligned}\int_A\int_{A_0}\int_{z_0} T\left(z\left(A\widehat{\Omega}\left(z_0, A_0\right)\right)\right) dF_{z_0}\left(z_0|A_0\right) dF_{A_0}\left(A_0\right) dF_A\left(A\right) \\ -p\int_{A_0}\int_{z_0} E\left(z_0\right) dF_{z_0}\left(z_0|A_0\right) dF_{A_0}\left(A_0\right) \geq R,\end{aligned} \tag{34}$$

where we assume independence in the preferences for working across generations. The government maximizes social welfare (33) subject to incentive compatibility constraints (2), skill production (32), and the government budget constraint (34).

The imperfect-tagging model does not change the optimal tax rule in equation (11), but it changes the optimal education rule. To characterize the optimal policy, suppose the government increases education provision by dE on the parental-earnings interval $(z_0, z_0 + dz_0)$. As in the baseline model, such a perturbation has a mechanical revenue effect, a behavioral revenue effect, and a direct welfare effect.

These effects may be characterized as follows:

$$dM_E = -\lambda \cdot p \cdot dE \cdot dz_0 \cdot \int_{A_0} f_{z_0}(z_0|A_0)\, dF_{A_0}(A_0) \tag{35}$$

$$dB_E = \lambda \cdot dE \cdot dz_0 \cdot \mathbb{E}\left[\tau(z)\frac{dz}{d\omega}\Omega'_E \mid z_0\right] \cdot \int_{A_0} f_{z_0}(z_0|A_0)\, dF_{A_0}(A_0) \tag{36}$$

$$dW_E = \lambda \cdot dE \cdot dz_0 \cdot \mathbb{E}\left[g(z_0, A_0, A^*)\frac{du}{d\omega}\Omega'_E \mid z_0\right] \cdot \int_{A_0} f_{z_0}(z_0|A_0)\, dF_{A_0}(A_0), \tag{37}$$

where we define $g(z_0, A_0, A^*) \equiv \Psi'\left[u\left(A^*\widehat{\Omega}(z_0, A_0)\right)\right]/\lambda$ as the social marginal welfare weight on individuals with parental background variables z_0, A_0. Using these expressions, we obtain the following result:

Proposition 5 (Optimal Predistribution with Imperfect Tagging). *For any given tax schedule $T(z)$, the optimal education provision at each level of parental income z_0 satisfies*

$$\mathbb{E}\left[\Omega'_E \cdot \left(g(z_0, A_0, A^*)\frac{du}{d\omega} + \tau(z)\frac{dz}{d\omega}\right) \mid z_0\right] = p, \tag{38}$$

where $du/d\omega$ denotes the marginal willingness to pay for higher skill, and $g(z_0, A_0, A^*)$ denotes the social marginal welfare weight on children with parental income z_0 and parental preferences A_0, conditional on average preferences for working in the children's generation A^*.

Proof. The result follows from the optimality condition $dM_E + dB_E + dW_E = 0$, using equations (35)-(37) and rearranging terms. □

The optimal education rule is similar to the baseline rule in Proposition 2. Formally, imperfect tagging changes the result in two ways: (*i*) the social marginal benefit of government provision is an average at parental income z_0, and (*ii*) the social marginal welfare weight is defined over parental income and preferences z_0, A_0 rather than parental skill ω_0. All else equal, the latter effect implies less progressivity in education provision because parental income is only a noisy signal of deservingness (background skill ω_0) due to preference heterogeneity.

4.3 General Dynamic Framework with Imperfect Tagging

The preceding sections simplify dynamics by considering models with only two generations—parents and children—where the government maximizes social welfare in the children's generation, taking policies in the parents' generation as given. It is possible to generalize the analysis to account for multiple generations and a government that sets tax and education schedules in each generation to maximize

dynamic social welfare. This section lays out the general framework, with the characterization of optimal policies presented in Appendix C.

Our focus is on intergenerational dynamics rather than lifecycle dynamics within a generation. We consider a setting in which non-overlapping generations live for one period.[24] Each generation t consists of a continuum of mass 1 of workers with heterogeneous skill ω_t and preferences for working A_t. The joint distribution of skill and preferences is given by the density function $f_t(\omega_t|A_t) f_{A_t}(A_t)$. The distribution of skill ω_t is predetermined within a generation, but it may change from generation to generation in response to education provision.

The utility of a worker of type (A_t, ω_t) is given by

$$u_t = z_t - T_t(z_t) - \frac{A_t\omega_t}{1+1/\varepsilon}\left(\frac{z_t}{A_t\omega_t}\right)^{1+1/\varepsilon}, \tag{39}$$

the maximization of which yields earnings supply:

$$z_t = A_t\omega_t\left(1-\tau_t(z_t)\right)^{\varepsilon}. \tag{40}$$

The earnings capacity of each worker—defined as earnings at $\tau_t = 0$—is given by $A_t\omega_t$. Indirect utility is written as $u_t(A_t\omega_t)$.

As above, the key departure from a Mirrleesian environment is to endogenize skill with respect to a pre-market instrument such as education provision. This instrument affects the intergenerational transmission of skill. Specifically, skill production takes the following form:

$$\omega_t = \Omega\left(\omega_{t-1}, E_{t-1}(z_{t-1})\right), \tag{41}$$

where $\Omega(.)$ is non-decreasing in both arguments. The skill of an individual in generation t depends on the previous generation through two mechanisms. The first is a draw ω_{t-1} from the skill distribution of the previous generation. This captures the effect of family background: ω_{t-1} can be interpreted as the skill level of one's parents. The second mechanism is the effect of government education provision for individuals of that family background, $E_{t-1}(z_{t-1})$. These investments are made during the working life of generation $t-1$ and are therefore pre-market from the perspective of generation t. Note that education provision depends on (endogenous) parental earnings, thereby allowing for imperfect tagging as in the previous section.

Inserting lagged values of skill into equation (41), we may rewrite intergenerational skill production

[24] More precisely, each generation of individuals makes active labor market decisions in only one period, but the individuals are alive during the previous period—i.e., when their parents make labor market decisions and when they may benefit from pre-market education provision.

as follows:

$$\omega_t = \widehat{\Omega}\left(\omega_0, E_0\left(z_0\right), ..., E_{t-1}\left(z_{t-1}\right)\right). \tag{42}$$

Hence, the skill of a worker in generation t depends on the baseline skill of their dynasty (ω_0) and on the entire history of education investments in their ancestors and themselves ($E_0, ..., E_{t-1}$). This implies that, if the government changes education investments in a given generation, the effect cascades through future generations due to intergenerational skill transmission. The strength of such cascading depends on the magnitude of $\Omega'_{\omega_{t-1}}$ in equation (41). If $\Omega'_{\omega_{t-1}}$ is small (large), intergenerational cascading fades quickly (slowly).

The government has an equal-opportunity objective: it respects inequality driven by preferences A_t, but not inequality driven by predetermined skill ω_t. Social welfare is specified as

$$W = \sum_{t=0}^{T} \int_{A_t} \int_{\omega_t} \delta_t \Psi\left[u_t\left(A_t^* \omega_t\right)\right] dF_{\omega_t}\left(\omega_t | A_t\right) dF_{A_t}\left(A_t\right), \tag{43}$$

where δ_t is a Pareto weight on generation t, and $A_t^* \equiv \int_{A_t} A_t dF_{A_t}\left(A_t\right)$ denotes average preferences for working in generation t.[25] When setting tax and education schedules, the government must satisfy a dynamic budget constraint:

$$\sum_{t=0}^{T} \int_{A_t} \int_{\omega_t} \beta^t \left[T_t\left(z_t\left(A_t \omega_t\right)\right) - p E_t\left(z_t\left(A_t \omega_t\right)\right)\right] dF_{\omega_t}\left(\omega_t | A_t\right) dF_{A_t}\left(A_t\right) \geq R, \tag{44}$$

where β is a discount factor. Here we assume that the price of education investments is constant across generations, but this is trivial to generalize.

The government sets generation-specific tax and education schedules $\left\{T_t\left(z_t\right), E_t\left(z_t\right)\right\}_{t=0}^{T}$ to maximize social welfare (43) subject to incentive compatibility (40), skill production (41), and the government budget constraint (44). We solve this optimization problem in Appendix C and provide dynamic policy rules similar in spirit to the static rules presented above. Allowing for general dynamics does not change the fundamental insights of the model, but it increases the empirical and numerical complexity by an order of magnitude. Solving the general model numerically is complicated—especially with many generations—because the tax and education schedules rely on endogenous entities that interact both within and across generations. The optimal policy in any given generation depends on potential effects many generations down the line—effects that are difficult to evaluate and are unlikely to be internalized by policymakers. Hence, the simpler models presented above have greater empirical and policy relevance.

[25] Alternatively, we could specify social welfare as a function of average preferences across all generations: $A^* \equiv \sum_{t=0}^{T} \int_{A_t} A_t dF_{A_t}\left(A_t\right) / T$. This specification choice does not affect the results in any substantive way.

5 Quantifying Optimal Predistribution

In this section we present numerical simulations that illustrate the quantitative importance of predistribution relative to traditional redistribution. Using our baseline model from Section 3, we tailor the calibrations to match a prominent predistribution policy: public expenditure on K-12 education. We first consider a simplified version of the model that abstracts from preference heterogeneity (in Sections 5.1-5.2) and then extend the simulations to allow for heterogeneity in preferences for working (in Sections 5.3-5.4). The full details of the numerical simulations are provided in Appendix D.

5.1 Calibration Details

The simulations require four functional forms. First, normalizing the preference parameter A to one, the utility function (1) may be written as

$$u\left(z;\omega_0\right) = z - T\left(z\right) - \frac{\widehat{\Omega}\left(\omega_0\right)}{1+1/\varepsilon}\left(\frac{z}{\widehat{\Omega}\left(\omega_0\right)}\right)^{1+1/\varepsilon}, \tag{45}$$

where earnings supply equals $z = \widehat{\Omega}\left(\omega_0\right)\left(1-\tau\left(z\right)\right)^{\varepsilon}$ by equations (2)-(3).

Second, the technology of skill production (3) is assumed to be additively separable in family-background skill ω_0 and government education investments $E\left(\omega_0\right)$. Specifically, we have

$$\omega = \widehat{\Omega}\left(\omega_0\right) = \Omega\left(\omega_0, E\left(\omega_0\right)\right) = \omega_0 + B \cdot E\left(\omega_0\right)^{\eta}. \tag{46}$$

The parameter B governs the aggregate productivity of government investments, while η governs their returns to scale. In the baseline model, we abstract from parental education investments, but these will be incorporated in an extension (in progress).

Third, we assume that background skill ω_0 is log-normally distributed with parameters $\left(\mu_{\omega_0}, \sigma_{\omega_0}\right)$.[26] Finally, the social welfare function takes a standard CRRA form, $\Psi\left(u\right) = \frac{u^{1-\rho}}{1-\rho}$, where the coefficient of relative risk aversion ρ governs inequality aversion. Our baseline calibration assumes $\rho = 1$, corresponding to the logarithmic case.

We set parameters based on external evidence or calibration as follows. The earnings elasticity ε with respect to the net-of-tax rate is set equal to 0.5 at baseline, broadly consistent with the empirical literature on intensive margin responses (Chetty *et al.* 2011*a*; Saez, Slemrod, and Giertz 2012; Kleven *et al.* 2025). We investigate comparative statics with respect to ε. The returns-to-scale parameter η is set equal to

[26] While log-normal distributions induce thinner tails than the ones observed empirically (Diamond, 1998; Saez, 2001), we refrain from appending a Pareto tail because, as we detail in Appendix D, the log-normal assumption allows us to exactly decompose the distribution of $A\omega_0$ into log-normal marginal distributions of A and ω_0, facilitating the analysis of the case with preference heterogeneity.

0.75, the midpoint of structural estimates surveyed in Browning, Hansen, and Heckman (1999).[27] We jointly calibrate the skill-distribution parameters $(\mu_{\omega_0}, \sigma_{\omega_0})$ and the productivity of government education spending B to match three empirical moments: the mean and standard deviation of the current U.S. labor income distribution (obtained from the five most recent waves of the Panel Study of Income Dynamics) and an empirical estimate of the elasticity of earnings with respect to government education spending, ζ.

Specifically, interpreting $E(\omega_0)$ as early-education provision, we define the empirical moment ζ as the elasticity of adult earnings with respect to government spending on K-12 education. Table 1 details our construction of ζ from the empirical literature. Jackson and Mackevicius (2024) survey twenty-six quasi-experiments on government K-12 spending in the U.S. to estimate its average per-dollar causal impact on test scores. Chetty, Friedman, and Rockoff (2014) study as-good-as-random assignments of students to teachers to estimate effects of teachers on both childhood test scores and later-life adult earnings, the ratio of which yields an adult earnings impact per childhood test score impact. Combining these two values yields our preferred estimate for ζ of 0.216, local to the below-average parent-income sample of Chetty, Friedman, and Rockoff (2014). Given the decreasing returns to scale $\eta = 0.75$, our calibration implies a population average $\mathbb{E}[\zeta]$ equal to 0.196. We consider the implications of alternative values of $\mathbb{E}[\zeta]$ in comparative statics analyses.

The calibration is implemented based on an approximation of the current policy environment in the U.S. We assume that every type ω_0 receives the same amount of public education spending, $E(\omega_0) = \overline{E}$, equal to the 2024 annual average K-12 spending per pupil of $20,668 (U.S. Census Bureau 2026). We empirically estimate the current tax system $T(z)$ using the microfiles of Piketty, Saez, and Zucman (2018), as in Gaubert *et al.* (2025).

5.2 Baseline Results

Having calibrated the model, we solve numerically for the policy vector $\{T(z), E(\omega_0)\}_{\forall z,\omega_0}$ that satisfies the optimal policy rules in equations (11) and (15). Details on the solution algorithm are provided in Appendix D.

Figure 2 plots optimal profiles of policy parameters by percentile of background skill ω_0: marginal and average tax rates in Panels A-B and education spending per child in Panel C. The figure shows optimal tax rates both when the planner can predistribute (our model in red) and when the planner cannot (standard Mirrleesian model in black). The education schedule is highly progressive: optimal

[27] This structural literature follows the Ben-Porath tradition where earnings are a function of educational dollars, educational time (sometimes called on-the-job training), and initial human capital. Browning, Hansen, and Heckman (1999) report a range of 0.5 to 0.99 for the returns to scale on educational dollars plus educational time. More recent papers report estimates of similar orders of magnitude: Manuelli and Seshadri (2014) estimate returns to scale of 0.89 for flow educational dollars, while Attanasio *et al.* (2020) estimate returns to scale of 0.52-0.59 for material investments but negligible returns to time.

spending is monotonically declining in background skill, falling from roughly $21,000 per child at the bottom to almost zero at the top. Optimal spending is about six times larger in the bottom percentiles than at the median. The equalizing effect of education provision relieves pressure on the tax-transfer system to equalize incomes. As a result, the optimal tax system becomes less progressive. With predistribution, marginal tax rates are everywhere lower, average tax rates are flatter, and the demogrant (not shown) is much smaller ($8,553 vs $23,541).[28]

Figure 3 illustrates the income-equalizing effects of predistribution and redistribution at the optimal policy. It shows density distributions of background skill ω_0, realized skill ω, and after-tax income $z - T(z)$ in our predistribution model (Panel A) and in the standard Mirrleesian model (Panel B). The difference between the densities of ω_0 and ω represents predistribution, while the difference between the densities of ω and $z - T(z)$ represents redistribution. As we can see, the effects of predistribution are substantial. Progressive education spending boosts realized skill (i.e., potential earnings) at the bottom of the distribution, raising the lower bound from roughly zero to $30,000. Progressive taxes further raise the lower bound (but by less) and reduce incomes at the top. By contrast, when traditional redistribution is the only option, taxes and transfers are left to do more income equalization on their own. They cannot compensate fully, however, and the distribution of after-tax income in Panel B is therefore more dispersed than in Panel A.

Table 2 examines the effects of predistributive and redistributive policies on different measures of inequality. It reports inequality in background skill ω_0, realized skill ω, and after-tax income $z - T(z)$ across a range of different indicators. As above, predistribution corresponds to the difference between ω_0 and ω, while redistribution corresponds to the difference between ω and $z - T(z)$. Considering the Gini coefficient, we find that predistribution reduces the coefficient from 0.559 to 0.477 and that redistribution further reduces it to 0.301. As a result, predistribution is responsible for 31% of the aggregate reduction in this inequality measure. Relative to the Mirrleesian benchmark, optimal policy reduces the Gini coefficient by an extra 12%, from 0.559 to 0.301 instead of only to 0.342.

The role of predistribution is larger for inequality indicators that emphasize changes at the bottom of the distribution. As shown in the table, the share of predistribution equals 44%, 39%, 56%, 72%, and 80% when considering the bottom 10% share, the bottom 25% share, the Palma ratio, the P90-P10 ratio, and the P50-P10 ratio, respectively.[29] The improvement in inequality relative to the Mirrleesian benchmark is also larger for these indicators. This pattern matches those shown in Figures 2-3: the progressivity of the

[28] Because the distribution of types is bounded, labor supply responses are at the intensive margin, and the bottom type has positive ω_0 and, therefore, works, we replicate the standard result in optimal labor income taxation that the bottom and top marginal tax rates equal zero. Replicating the standard U-shaped pattern (e.g., Saez, 2001) requires an unbounded distribution with a thick tail (e.g., Pareto) and a bottom type that does not work (Piketty and Saez, 2013).

[29] The Palma ratio is defined as the ratio of the top 10% share to the bottom 40% share.

education profile compresses the distribution of ω primarily by boosting the bottom types, with relatively small effects at the top of the distribution.

To gain a better understanding of the tradeoff between predistribution and redistribution, we analyze the comparative statics with respect to the different parameters that govern this tradeoff. To economize on exposition, we focus on the Gini coefficient. Figure 4 presents comparative statics with respect to the tax elasticity ε, which summarizes the efficiency costs of redistribution.[30] The figure shows that, when ε goes to zero, the planner implements perfect equality in after-tax incomes using taxes and transfers, as in the standard Mirrleesian model.[31] As ε increases, the efficiency costs of taxation constrain the optimal level of redistribution. However, the availability of a predistributive instrument relaxes this constraint, allowing the planner to substitute from tax-transfer policy towards education provision as ε increases. At $\varepsilon = 1$, more than half of the policy-induced reduction in the Gini coefficient is achieved by predistribution.

Figure 5 presents comparative statics with respect to the population-average education elasticity $\mathbb{E}[\zeta]$, which governs the productivity of public education. The analysis is implemented by recalibrating B in equation (46), holding the returns-to-scale parameter η constant, to match different values of $\mathbb{E}[\zeta]$ at the baseline policy environment. The relative importance of predistribution responds strongly to this parameter. As $\mathbb{E}[\zeta]$ falls to 0.1 or less, the role of predistribution disappears and the optimum converges to the Mirrleesian benchmark. Conversely, as $\mathbb{E}[\zeta]$ increases above its baseline value of 0.197, predistribution increases strongly and the optimal Gini coefficient falls substantially. At $\mathbb{E}[\zeta] = 0.3$, the income-equalizing effect of government policy is driven almost entirely by predistribution.

Finally, Figure 6 presents comparative statics with respect to the inequality-aversion parameter ρ. As inequality aversion increases, the planner implements a lower Gini coefficient through a combination of pre- and redistribution. In the limit of Rawlsian social preferences ($\rho \to \infty$), the planner reduces the Gini coefficient in after-tax incomes to 0.18, with almost 60% of the reduction driven by predistribution. This finding is consistent with Figures 2-3, in which we saw that predistribution is the preferred instrument for boosting the incomes of those from the lowest family backgrounds. At the opposite extreme, when preferences for redistribution vanish ($\rho \to 0$), the government provides a uniform amount of education, $E(\omega_0) = \overline{E}$, for efficiency reasons, funded by a regressive tax system. The resulting inequality in after-tax incomes is slightly larger than inequality in family-background skill.

[30] Changing ε has two effects: it affects the efficiency cost of redistribution—the channel we want to highlight in Figure 4—but it also affects the welfare benefit of predistribution by changing the marginal willingness to pay for higher skill, $\frac{du}{d\omega} = \frac{A(1-\tau(z))^{1+\varepsilon}}{1+\varepsilon}$. To isolate the redistribution-cost channel in the comparative statics, for each value of ε, we recalibrate A such that $\frac{du}{d\omega}$ is constant at the baseline tax system.

[31] We note that predistribution remains relevant even when redistribution is unrestricted ($\varepsilon = 0$): with no residual inequality, the planner provides a uniform level of education, $E(\omega_0) = \overline{E}$, for efficiency reasons. This mechanically decreases the Gini coefficient.

5.3 Results with Preference Heterogeneity

The baseline simulations abstract from preference heterogeneity, thereby neutralizing any implications of the planner's equal-opportunity objective. These simulations isolate the effect of one of our model's innovations: introducing a predistribution instrument into an otherwise Mirrleesian framework. As we have seen, allowing for predistribution can have dramatic effects on policy recommendations.

We now extend the simulations to allow for heterogeneity in preferences for working, A. Individuals are characterized by a two-dimensional vector of background types (A, ω_0), where we simplify by assuming that A and ω_0 are independently distributed. We denote earnings capacity (before education investments) by $\widetilde{\omega} = A\omega_0$. Our objective is to study, for a fixed distribution of $\widetilde{\omega}$, how optimal policy changes as a larger share of the variation in $\widetilde{\omega}$ is explained by variation in A. Our simulations in the previous section correspond to the limit case where none of the variation in $\widetilde{\omega}$ is explained by A. Another limit case is where ω_0 is constant across individuals, so that all of the variation in $\widetilde{\omega}$ is explained by A. Leveraging the independence between A and ω_0, we operationalize the intermediate cases of this decomposition with the R^2 of a (weighted) linear regression of $\widetilde{\omega}$ on A. When A is homogeneous, we have $R^2 = 0$. When ω_0 is homogeneous, we have $R^2 = 1$. Intermediate cases have $R^2 \in (0, 1)$, providing an interpretable measure of the relative importance of preferences in driving earnings capacity.[32]

The challenge is to compute different pairs of marginal distributions of A and ω_0 that yield the same distribution of $\widetilde{\omega} = A \cdot \omega_0$ but different values of the relative importance of preferences as measured by the R^2 statistic described above. We fix the distribution of $\widetilde{\omega}$ using the calibration of the baseline model without preference heterogeneity; that is, we assume that $\widetilde{\omega}$ is log-normal with parameters $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}}) = (3.95, 1.09)$.[33] Leveraging the independence of A and ω_0, we assume that each of the marginal distributions are also log-normal with parameters (μ_A, σ_A) and $(\mu_{\omega_0}, \sigma_{\omega_0})$. Based on known properties of log-normal distributions, this gives $\mu_{\widetilde{\omega}} = \mu_A + \mu_{\omega_0}$ and $\sigma_{\widetilde{\omega}} = \sqrt{\sigma_A^2 + \sigma_{\omega_0}^2}$. In Appendix D we use the expressions for the mean and variance of log-normal distributions to further show that the R^2 of the regression of $\widetilde{\omega}$ against A can be written as a function of the marginal log-normal parameters: $R^2 = \frac{\exp(\sigma_A^2)-1}{\exp(\sigma_A^2+\sigma_{\omega_0}^2)-1}$. This yields an identification result: given $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}})$, the values of $\mathbb{E}[\omega_0]$, $\mathbb{E}[A]$, and R^2 jointly determine (μ_A, σ_A) and $(\mu_{\omega_0}, \sigma_{\omega_0})$. That is, we can back out marginal log-normal distributions of A and ω_0 that induce the exact distribution of $\widetilde{\omega} = A \cdot \omega_0$ and have constant means $\mathbb{E}[\omega_0]$ and $\mathbb{E}[A]$, but whose variances vary to generate different values of the R^2. For our comparative statics, we compute marginal distributions that span a grid of R^2 that ranges from 0 to 1.

[32] Standard variance decompositions do not work in this case because we are concerned with the product, not the sum, of A and ω_0.

[33] More generally, we note that the calibration above remains valid for this exercise as the decomposition between A and ω_0 affects optimal policy through the definition of welfare weights, but not the labor supply decisions in the absence of policy.

Given the set of marginal distributions, we numerically solve for the optimal policy for each value of R^2 (see Appendix D for details). The Gini decomposition is presented in Panel A of Figure 7. Consistent with our theoretical analysis, the figure shows that, as the importance of preference heterogeneity increases, the planner equalizes incomes by less through either predistribution or redistribution. The reason is that the profile of welfare weights flattens and, therefore, the planner becomes less inequality averse, consistent with the arguments in Lockwood and Weinzierl (2015) and Berg, Håvarstein, and Stubhaug (2025). In the limit where $R^2 = 1$, the planner implements a regressive tax system to finance uniform schooling, $E(\omega_0) = \overline{E}$, motivated solely by an efficiency rationale.[34]

This aggregate pattern, however, hides an important aspect of predistribution and its interaction with preference heterogeneity. Panel B of Figure 7 plots a measure of the progressivity of education spending against the relative importance of preferences (as measured by the R^2). For each R^2, we compute the average elasticity of education spending $E(\omega_0)$ with respect to background skill ω_0, $-d\log E(\omega_0)/d\log\omega_0$, which coincides with the object $-\mathbb{C}(\log E(\omega_0), \log\omega_0)/\mathbb{V}(\log\omega_0)$.[35] This measure, which resembles standard scalar measures of the progressivity of a tax system (Musgrave and Thin 1948), provides insights into the compression effects of education per unit of remaining opportunity inequality. The figure shows that the progressivity of education increases with the importance of preference heterogeneity. Intuitively, as preference heterogeneity increases, earnings become less informative about opportunity and therefore a worse tag of opportunity. Instead, the planner chooses to address the welfare-relevant inequality with education, as that instrument targets the right source of variation. This increase in education progressivity is not visible in Panel A because a larger variance in preferences A mechanically lowers the variance in background skill ω_0 such that education has a smaller impact on aggregate inequality.

5.4 Empirical Decomposition of Choice vs Opportunity

The preceding section highlights the share of earnings variation explained by preferences relative to family background as an important statistic for policy design. In this section, we propose to measure this statistic based on a decomposition of U.S. earnings variation into choice vs opportunity.

Inspired by a long tradition built upon Roemer (1998), we implement the following empirical exercise. Using data from the 1979 cohort of the National Longitudinal Survey of Youth (NLSY79), we classify

[34] The constant education provision does not affect the Gini before taxes and transfers, which differs from the limit case with no inequality aversion in Figure 6, where the constant education spending induced predistribution. The reason relates to properties of the Gini coefficient. When the constant education spending is added to a non-constant background type ω_0, inequality falls because the constant is quantitatively more relevant to low types. However, when the constant education spending is added to a constant background type ω_0, this raises the constant factor multiplying the varying A, so the change is proportional and hence distributionally neutral. See Appendix E for a formal argument.

[35] We define this elasticity in absolute value (as $E(\omega_0)$ is non-increasing in ω_0 at the optimum).

covariates into vectors of opportunity O and choice C and conduct the following sequential procedure. First, we regress wage earnings on the opportunity vector O and denote by R^2_O the adjusted R^2 of that regression. We then regress wage earnings on both the opportunity O and choice C vectors and denote by R^2_{O+C} the adjusted R^2 of that regression. With these measures, we build two benchmarks for the relative importance of preferences. First, we consider the quantity $R^2_{O+C} - R^2_O$, which we note is a lower bound on the share of earnings variation explained by choices as it assigns all the unexplained variation to opportunity. Second, we consider the quantity $\left(R^2_{O+C} - R^2_O\right) / R^2_{O+C}$, which we call the "proportional benchmark" as it assigns the unexplained variation to opportunity and choice in proportion to their explained shares. In both benchmarks, our sequential procedure assigns all the covariance between O and C to opportunity given standard omitted variable bias mechanics.

When implementing this exercise, the researcher must take a stance on which variables belong to the vectors O and C, which remains contentious in the literature (Roemer and Trannoy 2016). Our preferred specification defines O to include sex, race, year of birth, country of birth, parent education, parent employment, parent income, parent age, and the Armed Forces Qualification Test score (an aptitude measure). We define C to include the child's own annual hours worked, four-year-lagged annual hours worked, own educational attainment, family choices like age at first birth and marital status, and crime and health behaviors like the number of delinquent acts and the number of cigarettes per day. All continuous variables are included as quartics.

Table 3 displays the results from opportunity-only regressions in Panel A and opportunity+choice regressions in Panel B. In our preferred specification of Panel A, opportunity covariates explain roughly 23% ($R^2_O = 0.229$) of the variation in the child's own wage earnings in 2013, when children were aged 48-56. The sequential addition of opportunity covariates in our baseline set incrementally increases the adjusted R^2. However, numerous other opportunity variables available in the NLSY79 have relatively little impact, including childhood family size, religion, newspaper subscription, library card holding, and psychological test scores for locus of control and self-esteem. Panel B shows that when adding hours worked—the first choice covariate—the adjusted R^2 rises substantially to 0.371. Adding the remaining choice variables yields $R^2_{O+C} = 0.414$: observed opportunity and choice variables together explain 41% of the variance in 2013 earnings. In the context of our metrics, the values $R^2_O = 0.229$ and $R^2_{O+C} = 0.414$ imply a lower bound on the share of earnings variation explained by choices equal to 18.5% and a benchmark share of 44.7%. Returning to Figure 7, we see that the x-axis values of 18.5% and 44.7% both imply a substantial role for predistribution. However, in the benchmark case, nearly half of earnings variation is due to choices that the planner respects and does not attempt to reduce, yielding a Gini coefficient that is over 0.1 higher than in the lower-bound case.

6 Conclusion

Governments around the world engage in both redistributive tax-transfer policies and predistributive policies such as early education provision. Predistribution is often motivated by concerns about equality of opportunity. Yet the canonical normative framework in economics abstracts from pre-market policy instruments and from social preferences for equal opportunity. This paper has generalized the Mirrleesian framework to jointly characterize optimal redistribution and optimal predistribution with a generalized Diamond-Saez rule and a generalized Samuelson rule. We find that predistribution is more attractive when redistribution is more distortionary, when predistributive spending has larger effects on earnings, and when earnings are a noisier signal of opportunity. Calibrating our model to U.S. data, we show that optimal policy is more progressive than in the Mirrleesian benchmark, but that much less of the progressivity is achieved through traditional taxes and transfers.

The analysis highlights three empirical parameters necessary to characterize optimal policy, only one of which—the elasticity of earnings with respect to the net-of-tax rate—has enjoyed decades of extensive research. The other two parameters would benefit from continued study: the elasticity of earnings with respect to education provision and other predistributive spending, and the share of earnings variation attributable to preferences vs opportunity. Greater empirical consensus on either would sharpen optimal policy recommendations.

References

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso. 2018. "Intergenerational Mobility and Preferences for Redistribution." *American Economic Review*, 108(2): 521–554.

Arrow, Kenneth. 1971. "A Utilitarian Approach to the Concept of Equality in Public Expenditures." *Quarterly Journal of Economics*, 85(3): 409–415.

Atkinson, Anthony, and Joseph Stiglitz. 1976. "The Design of Tax Structure: Direct versus Indirect Taxation." *Journal of Public Economics*, 6(1–2): 55–75.

Atkinson, Anthony, and Nicholas Stern. 1974. "Pigou, Taxation, and Public Goods." *Review of Economic Studies*, 41(1): 119–128.

Attanasio, Orazio, Sarah Cattan, Emla Fitzsimons, Costas Meghir, and Marta Rubio-Codina. 2020. "Estimating the Production Function for Human Capital: Results from a Randomized Controlled Trial in Colombia." *American Economic Review*, 110(1): 48–85.

Barr, Andrew, and Chloe Gibbs. 2022. "Breaking the Cycle? Intergenerational Effects of an Antipoverty Program in Early Childhood." *Journal of Political Economy*, 130(12): 3253–3285.

Berg, Kristoffer, Morten Håvarstein, and Magnus Stubhaug. 2025. "Meritocratic Labor Income Taxation." Working Paper.

Blanchet, Thomas, Lucas Chancel, and Amory Gethin. 2022. "Why Is Europe More Equal than the United States?" *American Economic Journal: Applied Economics*, 14(4): 480–518.

Blomquist, Sören, and Vidar Christiansen. 1995. "Public Provision of Private Goods as a Redistributive Device in an Optimum Income Tax Model." *Scandinavian Journal of Economics*, 97(4): 547–567.

Boadway, Robin, and Maurice Marchand. 1995. "The Use of Public Expenditures for Redistributive Purposes." *Oxford Economic Papers*, 47(1): 45–59.

Bovenberg, Lans, and Bas Jacobs. 2005. "Redistribution and Education Subsidies Are Siamese Twins." *Journal of Public Economics*, 89(11–12): 2005–2035.

Browning, Martin, Lars Peter Hansen, and James Heckman. 1999. "Micro Data and General Equilibrium Models." *Handbook of Macroeconomics*, 1: 543–633.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics*, 1: 451–488.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011*a*. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review*, 101(3): 471–475.

Chetty, Raj, John Friedman, and Jonah Rockoff. 2014. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review*, 104(9): 2633–2679.

Chetty, Raj, John Friedman, Nathaniel Hendren, Maggie Jones, and Sonya Porter. 2026. "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility." *American Economic Review*, 116(1): 1–51.

Chetty, Raj, John Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Schanzenbach, and Danny Yagan. 2011*b*. "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR." *Quarterly Journal of Economics*, 126(4): 1593–1660.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where Is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." *Quarterly Journal of Economics*, 129(4): 1553–1623.

Christiansen, Vidar. 1984. "Which Commodity Taxes Should Supplement the Income Tax?" *Journal of Public Economics*, 24(2): 195–220.

Diamond, Peter. 1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review*, 88(1): 83–95.

Diamond, Peter, and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives*, 25(4): 165–190.

Fleurbaey, Marc. 2008. *Fairness, Responsibility, and Welfare.* Oxford, UK:Oxford University Press.

Fleurbaey, Marc, and François Maniquet. 2011. "Compensation and Responsibility." In *Handbook of Social Choice and Welfare*. Vol. 2, ed. K. J. Arrow, A. Sen and K. Suzumura, Chapter 22. Elsevier B.V.

Fleurbaey, Marc, and François Maniquet. 2018. "Optimal Income Taxation Theory and Principles of Fairness." *Journal of Economic Literature*, 56(3): 1029–1079.

Gaubert, Cécile, Patrick Kline, Damian Vergara, and Danny Yagan. 2025. "Place-Based Redistribution." *American Economic Review*, 115(10): 3415–3450.

Gelber, Alexander, and Matthew Weinzierl. 2016. "Optimal Taxation When Children's Abilities Depend on Parents' Resources." *National Tax Journal*, 69(1): 11–40.

Gray-Lobe, Guthrie, Parag Pathak, and Christopher Walters. 2023. "The Long-Term Effects of Universal Preschool in Boston." *Quarterly Journal of Economics*, 138(1): 363–411.

Heckman, James, Rodrigo Pinto, and Peter Savelyev. 2013. "Understanding the Mechanisms Through Which an Influential Early Childhood Program Boosted Adult Outcomes." *American Economic Review*, 103(6): 2052–2086.

Heckman, James, Seong Moon, Rodrigo Pinto, Peter Savelyev, and Adam Yavitz. 2010. "The Rate of Return to the HighScope Perry Preschool Program." *Journal of Public Economics*, 94(1–2): 114–128.

Hufe, Paul, Andreas Peichl, John Roemer, and Martin Ungerer. 2017. "Inequality of Income Acquisition: The Role of Childhood Circumstances." *Social Choice and Welfare*, 49(3–4): 499–544.

Hufe, Paul, Ravi Kanbur, and Andreas Peichl. 2022. "Measuring Unfair Inequality: Reconciling Equality of Opportunity and Freedom from Poverty." *Review of Economic Studies*, 89(6): 3345–3380.

Hylland, Aanund, and Richard Zeckhauser. 1979. "Distributional Objectives Should Affect Taxes but Not Program Choice or Design." *Scandinavian Journal of Economics*, 81(2): 264–284.

Jackson, Kirabo, and Claire Mackevicius. 2024. "What Impacts Can We Expect from School Spending Policy? Evidence from Evaluations in the United States." *American Economic Journal: Applied Economics*, 16(1): 412–446.

Jacobs, Bas. 2013. "Optimal Redistributive Tax and Education Policies in General Equilibrium." *International Tax and Public Finance*, 20: 312–337.

Kaplow, Louis. 1996. "The Optimal Supply of Public Goods and the Distortionary Cost of Taxation." *National Tax Journal*, 49(4): 513–533.

Kaplow, Louis. 2004. "On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives*, 18(4): 159–175.

Kaplow, Louis, and Steven Shavell. 2001*a*. "Any Non-Welfarist Method of Policy Assessment Violates the Pareto Principle." *Journal of Political Economy*, 109(2): 281–286.

Kaplow, Louis, and Steven Shavell. 2001*b*. "Fairness versus Welfare." *Harvard Law Review*, 114(4): 961–1388.

Kleven, Henrik. 2021. "Sufficient Statistics Revisited." *Annual Review of Economics*, 13: 515–538.

Kleven, Henrik. 2026. "Externalities and the Taxation of Top Earners." *American Economic Journal: Economic Policy*, forthcoming.

Kleven, Henrik, Claus Kreiner, and Emmanuel Saez. 2009. "The Optimal Income Taxation of Couples." *Econometrica*, 77(2): 537–560.

Kleven, Henrik, Claus Kreiner, Kristian Larsen, and Jakob Søgaard. 2025. "Micro vs Macro Labor Supply Elasticities: The Role of Dynamic Returns to Effort." *American Economic Review*, 115(9): 2849–2890.

Kline, Patrick, and Christopher Walters. 2016. "Evaluating Public Programs with Close Substitutes: The Case of Head Start." *Quarterly Journal of Economics*, 131(4): 1795–1848.

Kreiner, Claus, and Nicolaj Verdelin. 2012. "Optimal Provision of Public Goods: A Synthesis." *Scandinavian Journal of Economics*, 114(2): 384–408.

Lockwood, Benjamin, and Matthew Weinzierl. 2015. "De Gustibus non est Taxandum: Heterogeneity in Preferences and Optimal Redistribution." *Journal of Public Economics*, 124: 74–80.

Mankiw, Gregory, Matthew Weinzierl, and Danny Yagan. 2009. "Optimal Taxation in Theory and Practice." *Journal of Economic Perspectives*, 23(4): 147–174.

Manuelli, Rodolfo, and Ananth Seshadri. 2014. "Human Capital and the Wealth of Nations." *American Economic Review*, 104(9): 2736–2762.

Mirrlees, James. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies*, 38(2): 175–208.

Musgrave, Richard, and Tun Thin. 1948. "Income Tax Progression, 1929-48." *Journal of Political Economy*, 56(6): 498–514.

Nichols, Albert, and Richard Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review*, 72(2): 372–377.

Piketty, Thomas, and Emmanuel Saez. 2013. "Optimal Labor Income Taxation." In *Handbook of Public Economics*. Vol. 5, ed. A. Auerbach, R. Chetty, M. Feldstein and E. Saez, Chapter 7. Elsevier B.V.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics*, 133(2): 553–609.

Pirttilä, Jukka, and Matti Tuomala. 2002. "Publicly Provided Private Goods and Redistribution: A General Equilibrium Analysis." *Scandinavian Journal of Economics*, 104(1): 173–188.

Roemer, John. 1998. *Equality of Opportunity.* Cambridge, MA:Harvard University Press.

Roemer, John, and Alain Trannoy. 2015. "Equality of Opportunity." In *Handbook of Income Distribution*. Vol. 2A, ed. A. B. Atkinson and F. Bourguignon, Chapter 4. Elsevier B.V.

Roemer, John, and Alain Trannoy. 2016. "Equality of Opportunity: Theory and Measurement." *Journal of Economic Literature*, 54(4): 1288–1332.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Tax Rates." *Review of Economic Studies*, 68(1): 205–229.

Saez, Emmanuel. 2002. "The Desirability of Commodity Taxation Under Non-linear Income Taxation and Heterogeneous Tastes." *Journal of Public Economics*, 83(2): 217–230.

Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2(3): 180–212.

Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." *American Economic Review*, 106(1): 24–45.

Saez, Emmanuel, Joel Slemrod, and Seth Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature*, 50(1): 3–50.

Samuelson, Paul. 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics*, 36(4): 387–389.

Stantcheva, Stefanie. 2017. "Optimal Taxation and Human Capital Policies Over the Life Cycle." *Journal of Political Economy*, 125(6): 1931–1990.

Stantcheva, Stefanie. 2021. "Understanding Tax Policy: How Do People Reason?" *Quarterly Journal of Economics*, 136(4): 2309–2369.

Stiglitz, Joseph. 1982. "Self-Selection and Pareto Efficient Taxation." *Journal of Public Economics*, 17(2): 213–240.

Stiglitz, Joseph, and Partha Dasgupta. 1971. "Differential Taxation, Public Goods, and Economic Efficiency." *Review of Economic Studies*, 38(2): 151–174.

U.S. Census Bureau. 2026. "Per Pupil Amounts of Public Elementary-Secondary School System Finances by Enrollment-Size Groups: Fiscal Year 2024." *Annual Survey of School System Finances, Summary Tables, Table 14*, Available at https://www.census.gov/data/tables/2024/econ/school-finances/secondary-education-finance.html. Accessed June 16, 2026.

TABLE 1: ESTIMATE OF ζ: THE ELASTICITY OF ADULT EARNINGS WITH RESPECT TO GOVERNMENT K-12 EDUCATION SPENDING

A. Causal impact of government K-12 education spending on test scores (Jackson and Mackevicius 2024)	
Causal impact of four years of $1,000 (in 2018$) of government K-12 education spending on standard deviations of contemporaneous test scores	0.0316 (0.0056)
B. Percent earnings increase per standard deviation impact on test scores (Chetty, Friedman, and Rockoff 2014)	
Percent earnings increase per standard deviation increase in contemporaneous test scores	12.7% (3.3%)
Inputs to calculation based on teacher-value-added impacts (12.7% = $350 / $21,256 / 0.13):	
Causal impact of one year of one-std.-dev.-higher-teacher-value-added teacher on age-28 wage earnings	$350 ($92)
Age-28 mean earnings within the sample population equals $21,256 (precisely estimated)	$21,256
Causal impact of one year of one-std.-dev.-higher-teacher-value-added on contemporaneous test scores (precisely estimated)	0.13
C. Other inputs	
Per-pupil K-12 annual expenditure in 2024 (Census)	$20,668
CPI-U inflation 2018-2024 (BLS)	24.9%
D. Implied estimate of education elasticity ζ	
Estimate of the elasticity ζ of adult earnings with respect to government K-12 education spending, equal to:	0.216
Percent effect on earnings dz/z (= 0.0316 x 0.127)	0.4%
Divided by: Percent dosage of government education spending dE/E (= 4 x $1,000 x 1.249 / (13 x 20,668))	1.9%
Memo: Implied population-average estimate of ζ, reflecting the CFR sample's below-average parent income (numerical)	0.196

Notes: Panel A reports the meta-analysis estimate of the impact of government K-12 spending on student test scores, based on twenty-six U.S. quasi-experiments (Jackson and Mackevicius 2024, Table 2 column 1). Panel B reports estimates of test-score impacts and adult earnings impacts of quasi-experimental variation in teacher exposure in a large U.S. school district (Chetty, Friedman, and Rockoff 2014, Table 3 column 1 and Section IV.C). Panel C lists other inputs. Panel D uses Panels A-C to compute an estimate of the elasticity ζ of adult earnings with respect to government K-12 education spending. The computation's assumptions are that the Panel A studies are representative of the test score impacts of government education spending generally and that the Panel B ratio of the percentage adult earnings impact per standard deviation of test score impact applies to Panel A's test score impact. ζ is heterogeneous and declines with background skill. In our numerical simulation that yields $\zeta = 0.216$ on average for the bottom 89% of types (matching CFR's mean parent income), we obtain $\zeta = 0.196$ on average across all types.

TABLE 2: THE EFFECTS OF POLICY ON DIFFERENT INEQUALITY INDICATORS

	Predistribution Model				Mirrleesian Model	**Improvement over Mirrlees**
	ω_0	ω	$z - T(z)$	**Share $\omega_0 - \omega$**		
Gini	0.559	0.477	0.301	**31.47%**	0.342	**12.02%**
Bottom 10%	0.92%	3.04%	5.72%	**44.25%**	4.51%	**26.80%**
Bottom 25%	4.09%	8.29%	14.86%	**39.00%**	12.53%	**18.63%**
Top 10%	42.91%	39.42%	28.87%	**24.85%**	30.46%	**5.23%**
Top 1%	11.34%	10.38%	7.33%	**24.03%**	7.80%	**5.95%**
Palma Ratio	4.669	2.738	1.193	**55.54%**	1.420	**16.00%**
P90/P10	16.436	6.656	2.894	**72.22%**	3.656	**20.86%**
P90/P50	4.131	3.699	2.293	**23.46%**	2.389	**4.00%**
P50/P10	3.979	1.799	1.262	**80.22%**	1.531	**17.55%**

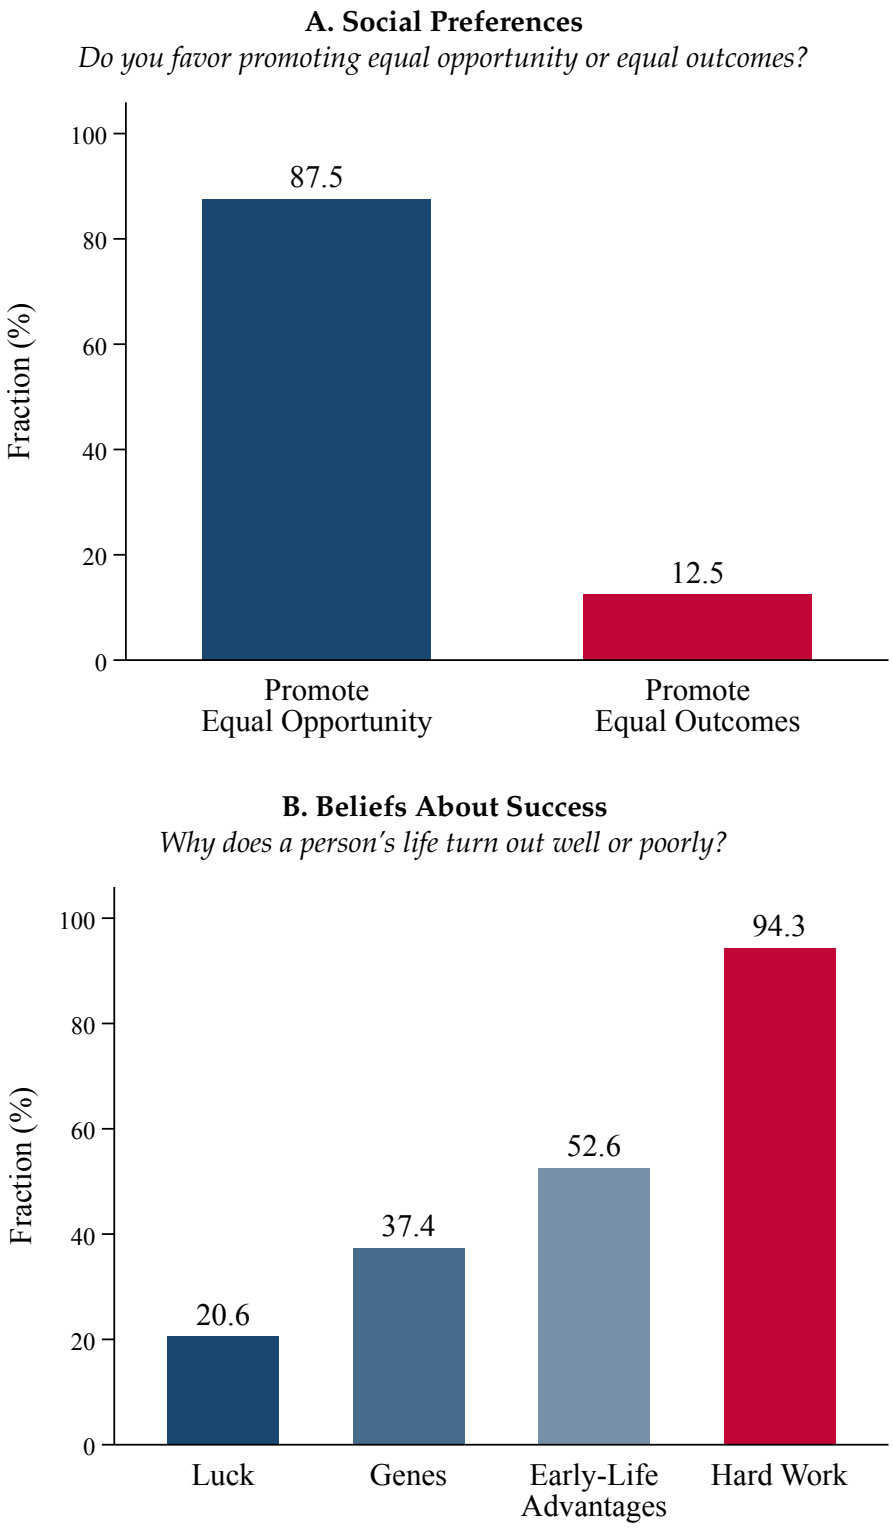
Notes: This table reports inequality in background skill ω_0, realized skill ω, and after-tax income $z - T(z)$ at the optimal policy. The difference between ω_0 and ω reflects predistribution, while the difference between ω and $z - T(z)$ reflects redistribution. The different rows consider different inequality indicators. The Palma ratio is defined as the ratio of the top-10% share to the bottom-40% share. The first four columns present results from our predistribution model, where the planner jointly optimizes tax-transfer policy and education provision. The column labeled "Share $\omega_0 - \omega$" shows the share of the total inequality reduction—from background skill to after-tax income—achieved through predistribution. That is, if $\Gamma(x)$ is the inequality indicator for the variable x, then "Share $\omega_0 - \omega$" is calculated as $(\Gamma(\omega_0) - \Gamma(\omega)) \,/\, (\Gamma(\omega_0) - \Gamma(z - T(z)))$. The table also reports inequality in after-tax income at the Mirrleesian benchmark (achieved through taxes and transfers alone), and the inequality improvement (in percentage terms) in our predistribution model relative to the Mirrleesian benchmark.

TABLE 3: DECOMPOSITION OF NLSY79 EARNINGS INTO OPPORTUNITY VS CHOICE

	Child's Own Adult Wage Earnings in 2013					
	(1)	(2)	(3)	(4)	(5)	(6)
A. Opportunity Covariates						
Demographics	X	X	X	X	X	X
Parent education		X	X	X	X	X
Parent income and age			X	X	X	X
AFQT score				X	X	X
Eight other childhood environment					X	X
Five other own endowment						X
R^2_O	0.086	0.149	0.177	0.229	0.234	0.240
B. Opportunity+Choice Covariates						
Opportunity covariates						
Demographics	X	X	X	X	X	X
Parent education	X	X	X	X	X	X
Parent income and age	X	X	X	X	X	X
AFQT score	X	X	X	X	X	
Choice covariates						
Hours worked	X	X	X	X	X	X
Lagged hours worked		X	X	X	X	X
Educational attainment			X	X	X	X
Family choices				X	X	X
Crime and health behaviors					X	X
R^2_O	0.229	0.229	0.229	0.229	0.229	0.177
R^2_{O+C}	0.371	0.384	0.404	0.414	0.414	0.404
$R^2_{O+C} - R^2_O$	14.2%	15.5%	17.5%	18.5%	18.5%	22.7%
$(R^2_{O+C} - R^2_O)$ / R^2_{O+C}	38.3%	40.3%	43.4%	44.6%	44.7%	56.2%

Notes: This table presents adjusted R^2 statistics from regressions of 2013 individual wage earnings in the 1979 National Longitudinal Survey of Youth. Each column of each panel represents a separate regression and includes indicators for missing values of all included covariates. All regressions are estimated on the same sample of 6,782 individuals. Demographics comprise indicators for race, sex, birth year, and country of birth. Parent education comprises indicators for mother years of education and father years of education. Parent income and age comprise quartics in family income in each year 1978-1981, mother age at birth, and father age at birth; indicators for poverty in each year 1978-1981; and indicators for mother's and father's employment. AFQT score comprises a quartic in the 1980-measured and 2006-re-scored Armed Forces Qualification Test score; we obtain a nearly identical increase in the adjusted R^2 between Panel A columns 3 and 4 when restricting the sample to birth cohorts 1962-1964, which were 18 years old or younger in 1980. Eight other childhood environment variables comprise a quartic in 1979 family size and indicators for religion raised in, newspaper subscription at age 14, magazine subscription at age 14, library card at age 14, interactions of parent types living with the respondent at age 14, and approximately one hundred occupational categories each for mother's occupation and father's occupation. Five other own endowment variables comprise quartics in height, 1979 Rotter's Internal-External Locus of Control score, and 1980 Rosenberg Self-Esteem Scale score and indicators for years of education top-coded at 12 (when the typical respondent was 18 years old) and 1979 health limitations. Hours worked comprises a quartic in hours worked in 2013. Lagged hours worked comprises a quartic in hours worked in 2009. Educational attainment comprises indicators for years of education. Family choices comprise quartics in age at first birth and spouse hours per week and indicators for being married in 2014 and number of children in the household in 2014. Crime and health behaviors comprise quartics in the number of delinquent acts and number of cigarettes per day and an indicator for ever having been incarcerated.

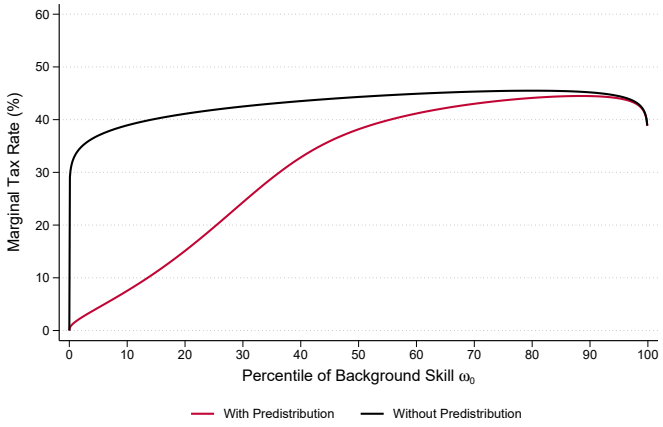
FIGURE 1: BELIEFS ABOUT EQUALITY OF OPPORTUNITY



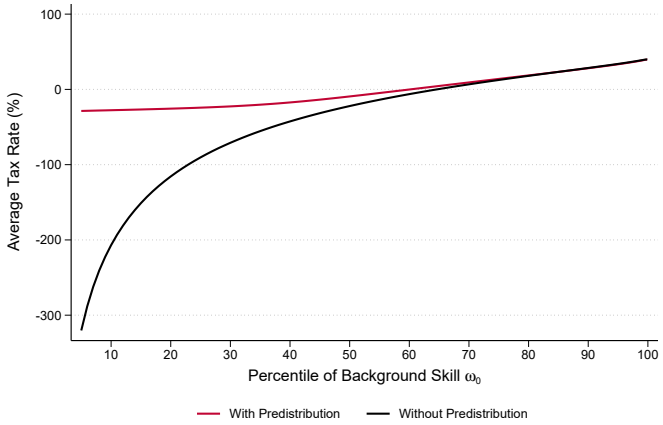
Notes: This figure reports social preferences and beliefs about success from the 1993 wave of the General Social Survey (GSS), the only wave that includes these questions. Panel A shows the share of respondents who favor promoting equal opportunity ("allowing everyone to compete for jobs and wealth on a fair and even basis") over promoting equal outcomes ("insuring that everyone has a decent standard of living and that there are only small differences in wealth and income between the top and bottom in society"). Panel B shows the share of respondents who rate different factors as either "important" or "very important" for why a person's life turns out well or poorly. A fifth surveyed factor ("such things are decided by God") is omitted in the figure. Shares are computed using GSS person post-stratification weights.

FIGURE 2: OPTIMAL POLICY PARAMETERS

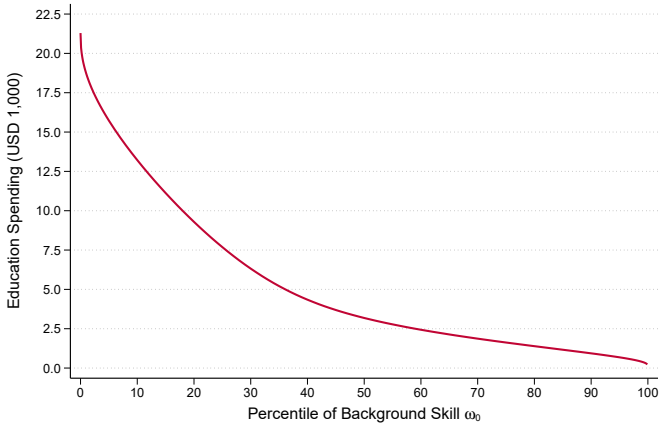
A. Redistribution
Marginal Tax Rate Profile



B. Redistribution
Average Tax Rate Profile



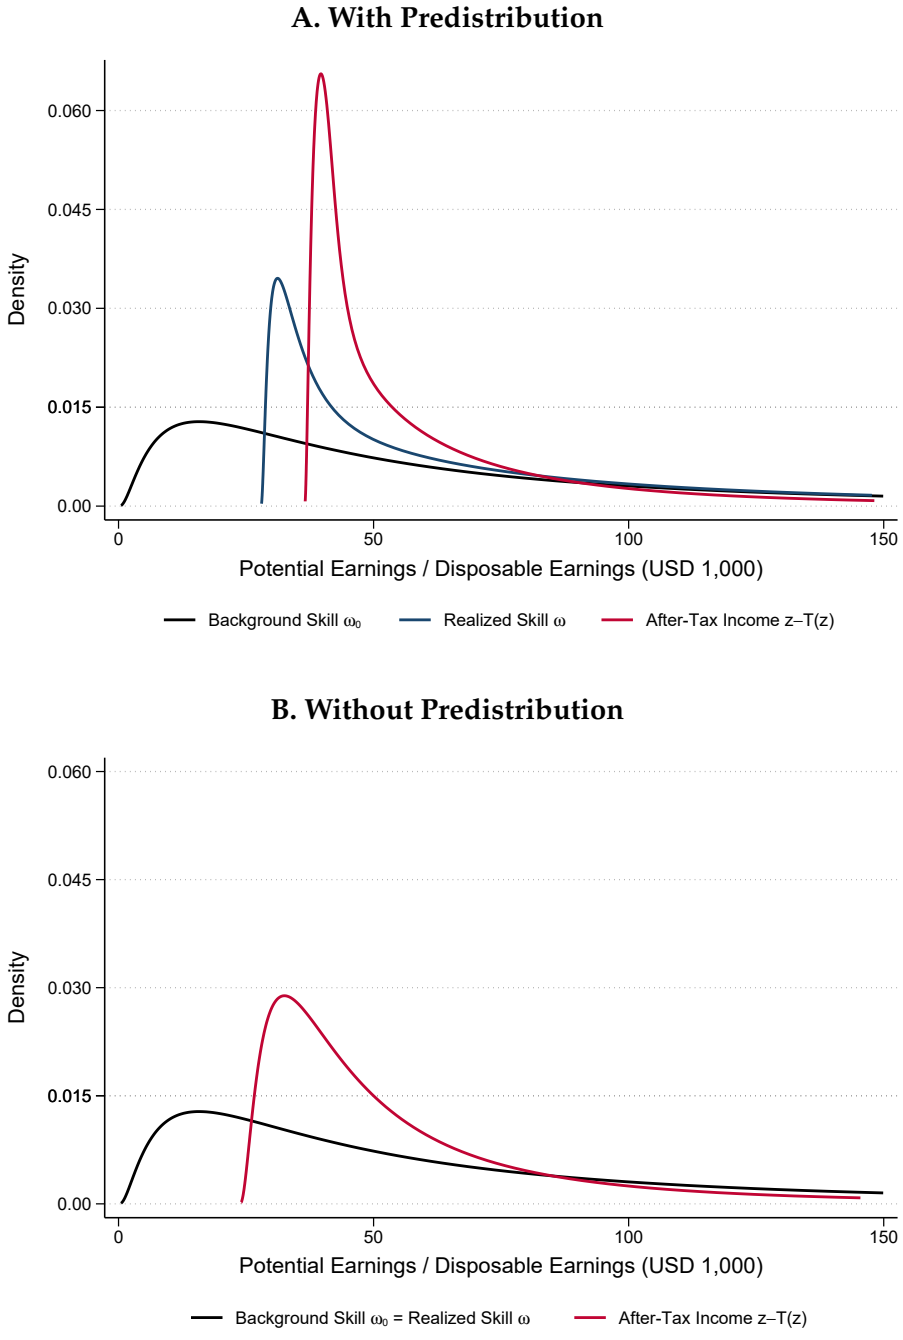
C. Predistribution
Education Profile



Notes: This figure plots the optimal policy functions by percentile of background skill ω_0. Panel A displays the profile of marginal tax rates. Panel B displays the profile of average tax rates, $T(z)/z$, truncating the x-axis at the 5th percentile for readability. In these panels, the red curves ("With Predistribution") depict optimal tax policy functions when the planner is also optimizing education provision, while the black curves ("Without Predistribution") depict optimal tax policy functions without any public education provision (Mirrleesian benchmark). Panel C displays the optimal profile of education spending, $E(\omega_0)$.

FIGURE 3: DISTRIBUTIONS OF POTENTIAL EARNINGS AND AFTER-TAX EARNINGS

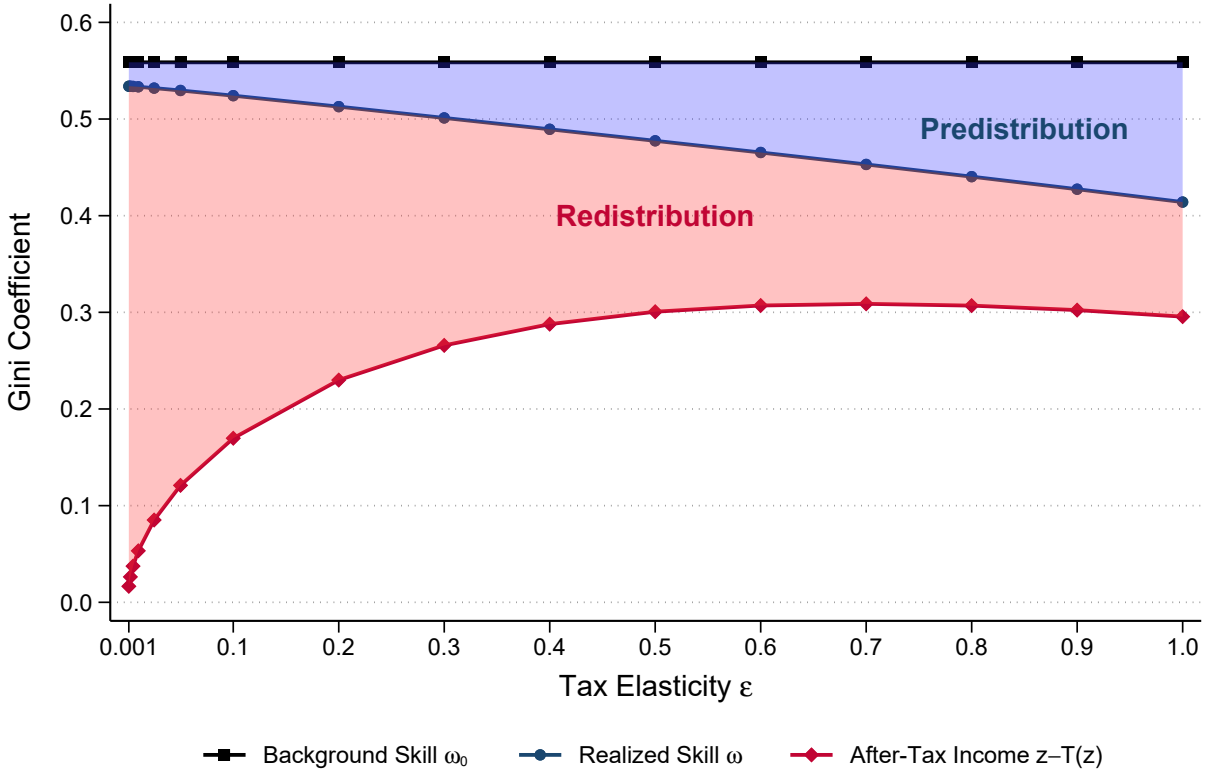
DISTRIBUTIONS OF ω_0, ω, AND $z - T(z)$ AT THE OPTIMAL POLICY



Notes: This figure plots the densities of background skill ω_0, realized skill ω, and after-tax income $z - T(z)$ at the optimal policy. Panel A ("With Predistribution") considers the case where the planner jointly optimizes taxes and education provision. Panel B ("Without Predistribution") considers the case where the planner only optimizes taxes (Mirrleesian benchmark without public education provision), in which case the distributions of background skill ω_0 and realized skill ω are identical.

FIGURE 4: EFFECT OF EARNINGS ELASTICITY WRT. $1-\tau$

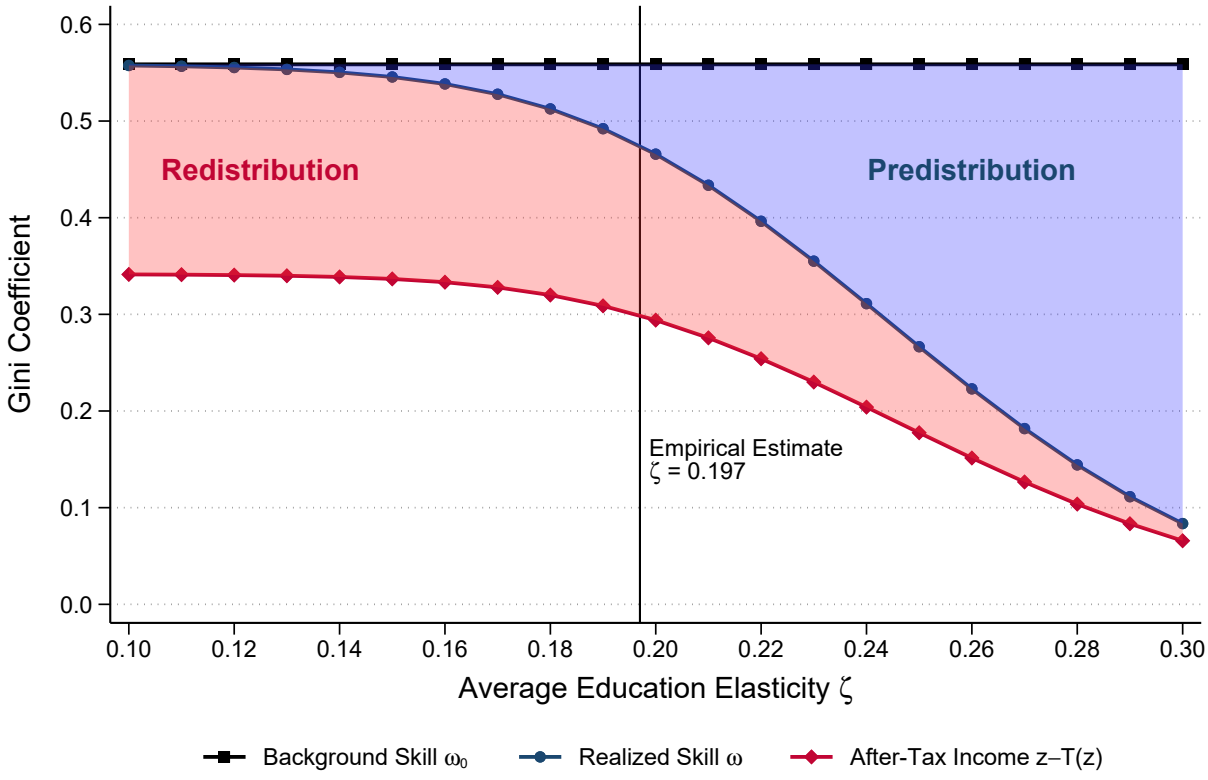
DECOMPOSITION OF THE GINI COEFFICIENT



Notes: This figure plots the Gini coefficient of background skill ω_0 (black squares), realized skill ω (blue circles), and after-tax income $z - T(z)$ (red diamonds) for different values of the tax elasticity ε. The difference between ω_0 and ω reflects predistribution, while the difference between ω and $z - T(z)$ reflects redistribution. When changing ε, the scalar parameter A is recalibrated to ensure that the marginal willingness to pay for higher skill, $du/d\omega = A\left(1-\tau(z)\right)^{1+\varepsilon} / (1+\varepsilon)$, is held constant at the baseline tax system.

FIGURE 5: EFFECT OF EARNINGS ELASTICITY WRT. EDUCATION

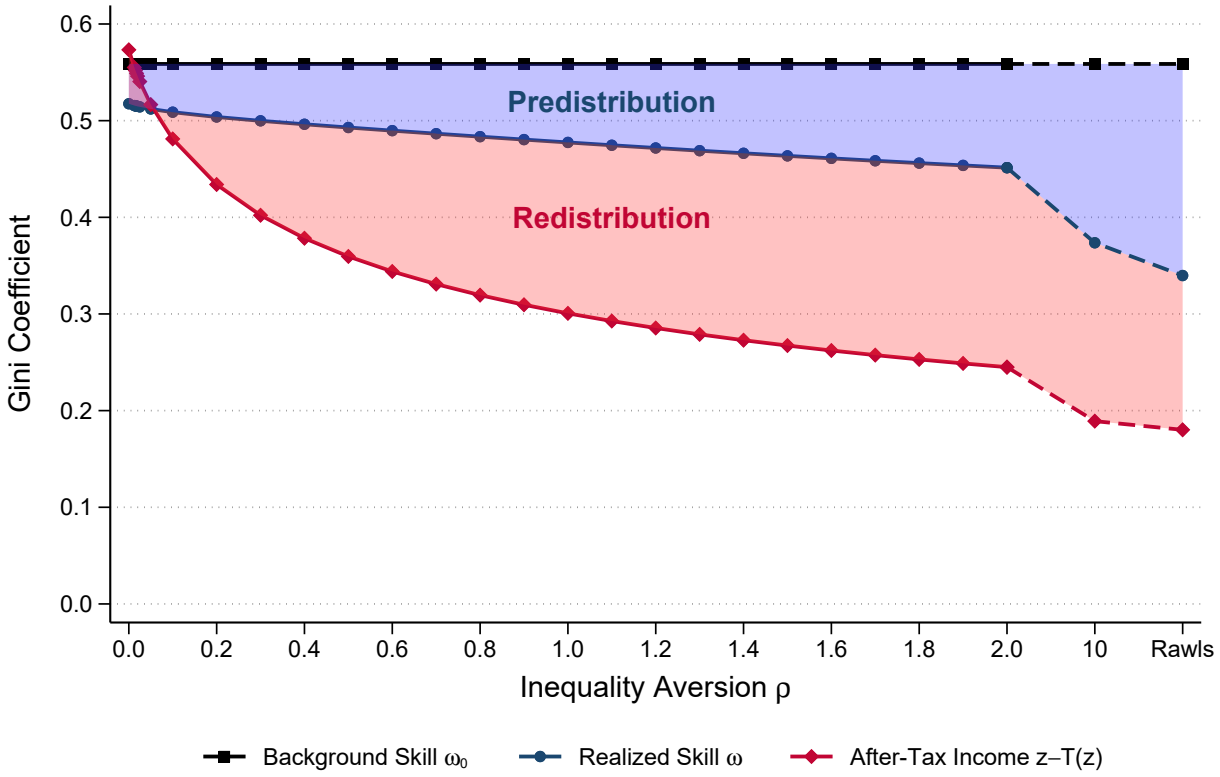
DECOMPOSITION OF THE GINI COEFFICIENT



Notes: This figure plots the Gini coefficient of background skill ω_0 (black squares), realized skill ω (blue circles), and after-tax income $z - T(z)$ (red diamonds) for different values of the average education elasticity $\mathbb{E}[\zeta]$. The difference between ω_0 and ω reflects predistribution, while the difference between ω and $z - T(z)$ reflects redistribution. The analysis is implemented by recalibrating B in equation (46), holding the returns-to-scale parameter η constant, to match different values of $\mathbb{E}[\zeta]$ at the baseline policy environment.

FIGURE 6: EFFECT OF INEQUALITY AVERSION

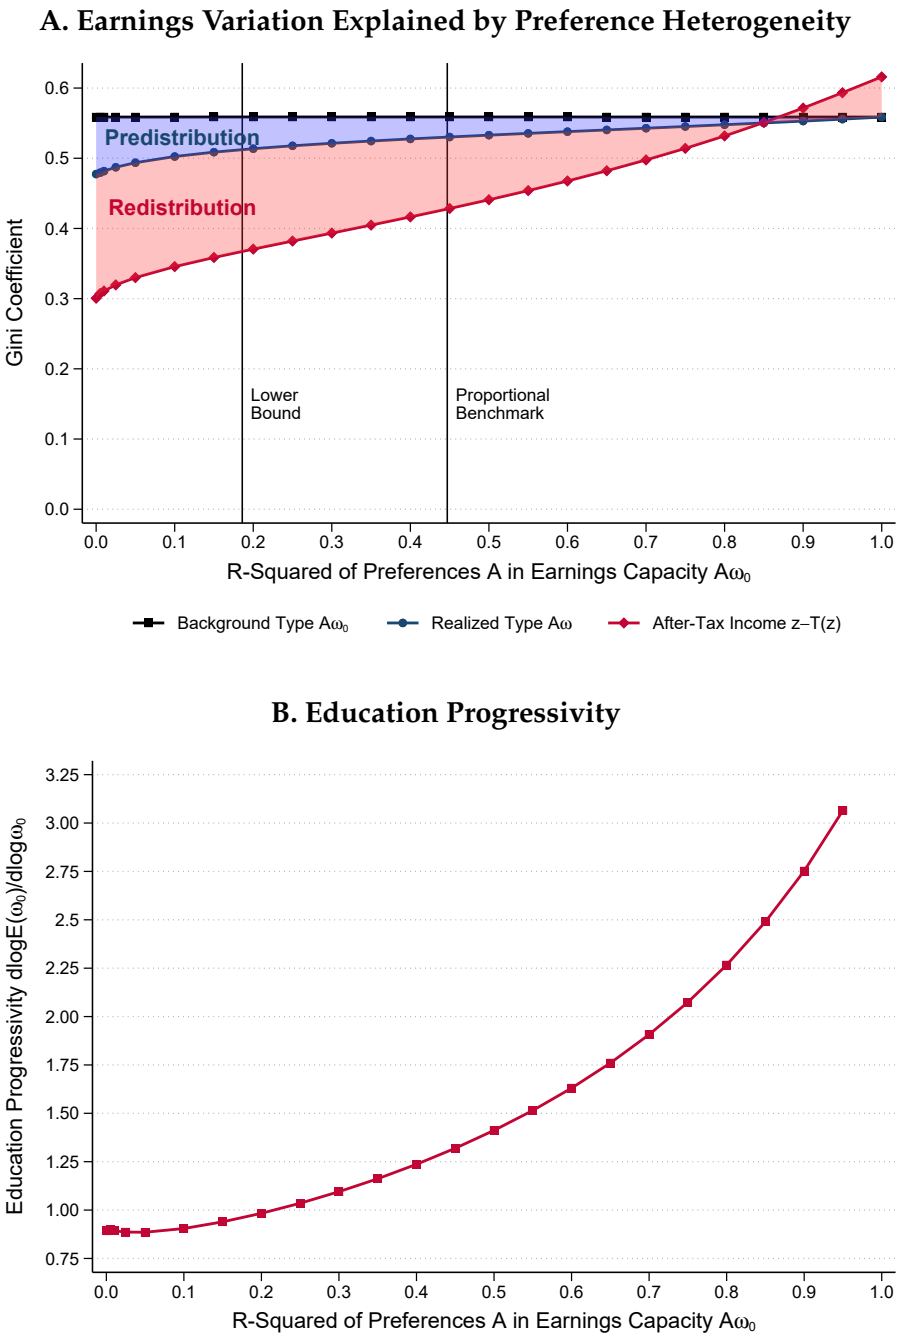
DECOMPOSITION OF THE GINI COEFFICIENT



Notes: This figure plots the Gini coefficient of background skill ω_0 (black squares), realized skill ω (blue circles), and after-tax income $z - T(z)$ (red diamonds) for different values of the inequality-aversion parameter ρ. The difference between ω_0 and ω reflects predistribution, while the difference between ω and $z - T(z)$ reflects redistribution. The point labeled "Rawls" on the x-axis corresponds to the limit where $\rho \to \infty$.

FIGURE 7: EFFECT OF PREFERENCE HETEROGENEITY

DECOMPOSITION OF THE GINI COEFFICIENT



Notes: This figure presents results from numerical simulations with preference heterogeneity. Panel A plots the Gini coefficient of background type $A\omega_0$ (black squares), realized type $A\omega$ (blue circles), and after-tax income $z - T(z)$ (red diamonds) for different values of the relative importance of preference heterogeneity as measured by the R^2 of a regression of $A\omega_0$ against A. The difference between $A\omega_0$ and $A\omega$ reflects predistribution, while the difference between $A\omega$ and $z - T(z)$ reflects redistribution. Panel B plots the progressivity of education spending as measured by the elasticity of education with respect to background skill in absolute value, i.e. $-d \log E(\omega_0) / d \log \omega_0 = -\mathbb{C}(\log E(\omega_0), \log \omega_0) / \mathbb{V}(\log \omega_0)$.

Online Appendix

Equality of Opportunity vs Equality of Outcomes: A Theory of Optimal Predistribution

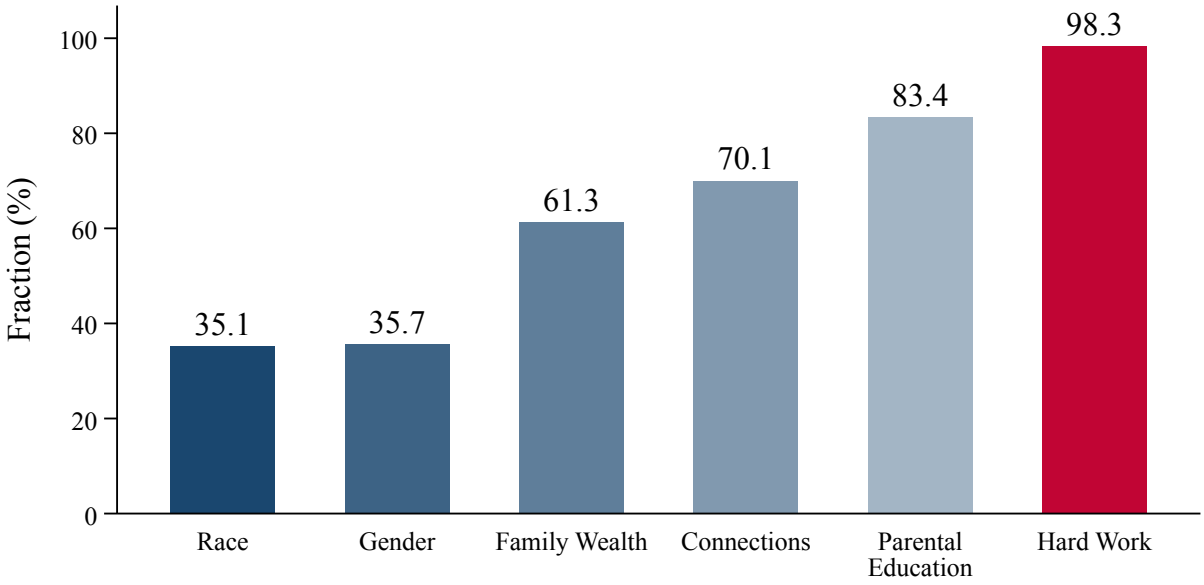
Raj Chetty
Henrik Kleven
Damián Vergara
Danny Yagan

August 2026

A Supplementary Exhibits

FIGURE A.1: BELIEFS ABOUT SUCCESS

ELICITED BELIEFS ABOUT WHY PEOPLE GET AHEAD IN LIFE



Notes: This figure reports beliefs about the determinants of success from the 2021 wave of the General Social Survey (GSS). Respondents were asked how important different factors are "for getting ahead in life," with answers on a five-point scale ranging from "essential" to "not important at all." The figure shows the share of respondents who rate each factor as at least "fairly important." Race, gender, family wealth, parental education, and hard work correspond to the survey items "a person's race," "being born a man or a woman," "coming from a wealthy family," "having well educated parents," and "hard work," respectively. The variable labeled "connections" combines two survey items, "knowing the right people" and "having political connections," and is coded as at least fairly important if either item is. Shares are computed using GSS person post-stratification weights.

B Proof of Remark 1

The *actual* indirect utility of an individual of type (A, ω_0) is given by

$$u(A, \omega_0) = z\left(A\widehat{\Omega}(\omega_0)\right) - T\left(z\left(A\widehat{\Omega}(\omega_0)\right)\right) - \frac{A\widehat{\Omega}(\omega_0)}{1+1/\varepsilon}\left(\frac{z\left(A\widehat{\Omega}(\omega_0)\right)}{A\widehat{\Omega}(\omega_0)}\right)^{1+1/\varepsilon}, \tag{B1}$$

where $z\left(A\widehat{\Omega}(\omega_0)\right) = A\widehat{\Omega}(\omega_0)\left(1-\tau\left(z\left(A\widehat{\Omega}(\omega_0)\right)\right)\right)^{\varepsilon}$.

The *imputed* utility used by the planner to evaluate social welfare is given by

$$u(A^*, \omega_0) = z\left(A^*\widehat{\Omega}(\omega_0)\right) - T\left(z\left(A^*\widehat{\Omega}(\omega_0)\right)\right) - \frac{A^*\widehat{\Omega}(\omega_0)}{1+1/\varepsilon}\left(\frac{z\left(A^*\widehat{\Omega}(\omega_0)\right)}{A^*\widehat{\Omega}(\omega_0)}\right)^{1+1/\varepsilon}, \tag{B2}$$

where $A^* = \mathbb{E}[A]$. Aggregate social welfare is defined as

$$W = \int_A \int_{\omega_0} \Psi\left[u(A^*, \omega_0)\right] dF_0(\omega_0|A)\, dF_A(A). \tag{B3}$$

In what follows, it will be convenient to rewrite W in terms of individuals indexed by i such that $u_i = u(A_i, \omega_{0i})$, $u_i^* = u(A^*, \omega_{0i})$, and $W = \int_i \Psi[u_i^*]\, di$.

To prove that W is Paretian, consider a set of individual-specific utility perturbations $\Delta = \{\delta_i\}_{\forall i}$ such that the actual utility of individual i equals $u(A_i, \omega_{0i}) + \delta_i$, and the planner's imputed utility of individual i equals $u(A^*, \omega_{0i}) + \delta_i$. Given quasi-linearity, the perturbations δ_i can be interpreted as individual-specific lump-sum transfers ($\delta_i > 0$) or lump-sum taxes ($\delta_i < 0$).

We may then show that, for any two sets $\Delta^j = \left\{\delta_i^j\right\}_{\forall i}$ such that $u(A_i, \omega_{0i}) + \delta_i^1 \geq u(A_i, \omega_{0i}) + \delta_i^2$ with strict inequality for at least one individual, we have $W\left(\Delta^1\right) > W\left(\Delta^2\right)$ where $W\left(\Delta^j\right) = \int_i \Psi\left[u_i^* + \delta_i^j\right] di$. The proof is straightforward:

$$\begin{aligned} u(A_i, \omega_{0i}) + \delta_i^1 \geq u(A_i, \omega_{0i}) + \delta_i^2 \quad &\Rightarrow \quad \delta_i^1 \geq \delta_i^2 \\ &\Rightarrow \quad u(A^*, \omega_{0i}) + \delta_i^1 \geq u(A^*, \omega_{0i}) + \delta_i^2 \\ &\Rightarrow \quad \Psi\left[u(A^*, \omega_{0i}) + \delta_i^1\right] \geq \Psi\left[u(A^*, \omega_{0i}) + \delta_i^2\right], \end{aligned}$$

given $\Psi[.]$ is a strictly increasing function. If this holds for every individual i, then $W\left(\Delta^1\right) = \int_i \Psi\left[u(A^*, \omega_{0i}) + \delta_i^1\right] di \geq \int_i \Psi\left[u(A^*, \omega_{0i}) + \delta_i^2\right] di = W\left(\Delta^2\right)$. Moreover, if $u(A_i, \omega_{0i}) + \delta_i^1 > u(A_i, \omega_{0i}) + \delta_i^2$ for at least one individual, then $W\left(\Delta^1\right) > W\left(\Delta^2\right)$ holds with strict inequality.

These derivations prove Remark 1: Given a distribution of types (A, ω_0), the social welfare function is Paretian with respect to any combination of individual-specific utility perturbations (thus allowing for

lump-sum taxes and transfers based on type). We may also prove the converse: any non-Paretian welfare effects must involve manipulating the structure of primitives.

To see this, suppose W is non-Paretian. Then there exist allocations $\left\{A_i^1, \omega_{0i}^1, \delta_i^1\right\}_{\forall i}$ and $\left\{A_i^2, \omega_{0i}^2, \delta_i^2\right\}_{\forall i}$ such that $u\left(A_i^1, \omega_{0i}^1\right) + \delta_i^1 \geq u\left(A_i^2, \omega_{0i}^2\right) + \delta_i^2$ for every individual i (Allocation 1 Pareto dominates), yet $\Psi\left[u\left(A^*, \omega_{0i}^1\right) + \delta_i^1\right] < \Psi\left[u\left(A^*, \omega_{0i}^2\right) + \delta_i^2\right]$ for at least one individual (such that Allocation 1 may generate a lower value of W). Hence,

$$\begin{aligned}
\Psi\left[u\left(A^*, \omega_{0i}^1\right) + \delta_i^1\right] < \Psi\left[u\left(A^*, \omega_{0i}^2\right) + \delta_i^2\right] \quad &\Rightarrow \quad u\left(A^*, \omega_{0i}^1\right) + \delta_i^1 < u\left(A^*, \omega_{0i}^2\right) + \delta_i^2 \\
&\Rightarrow \quad u\left(A^*, \omega_{0i}^1\right) + \delta_i^1 + u\left(A_i^1, \omega_{0i}^1\right) - u\left(A_i^1, \omega_{0i}^1\right) < u\left(A^*, \omega_{0i}^2\right) + \delta_i^2 \\
&\Rightarrow \quad u\left(A^*, \omega_{0i}^1\right) + u\left(A_i^2, \omega_{0i}^2\right) + \delta_i^2 - u\left(A_i^1, \omega_{0i}^1\right) < u\left(A^*, \omega_{0i}^2\right) + \delta_i^2 \\
&\Rightarrow \quad u\left(A_i^2, \omega_{0i}^2\right) - u\left(A_i^1, \omega_{0i}^1\right) < u\left(A^*, \omega_{0i}^2\right) - u\left(A^*, \omega_{0i}^1\right)
\end{aligned}$$

Now, assume that the allocations are constrained to be fixed in primitives, i.e., $A_i^1 = A_i^2$ and $\omega_{0i}^1 = \omega_{0i}^2$ for any individual i. From the inequality derived above, this constraint generates a contradiction: $0 < 0$. Therefore, any change in allocation that involves non-Paretian behavior of the social welfare function must be associated with a reallocation of the primitives of the model.

C General Dynamic Framework: Optimal Policy

Consider the dynamic framework laid out in Section 4.3. To solve for the optimal policies in this model, we start by defining social marginal welfare weights. Given tax and education schedules depend on income, it is convenient to express welfare weights in terms of income levels. Specifically, we define

$$g_t\left(z_t, A_t, A_t^*\right) \equiv \frac{\delta_t \Psi'\left[u_t\left(\frac{A_t^*}{A_t} z_t, A_t^* \omega_t\left(z_t, A_t\right)\right)\right]}{\beta^t \lambda}, \tag{C1}$$

as the social marginal welfare weight at earnings z_t and preferences A_t in generation t, evaluated at average preferences A_t^*. Correspondingly, we denote the average welfare weight at earnings level z_t by $g_t\left(z_t, A_t^*\right)$ and the average welfare weight *above* earnings level z_t by $G_t\left(z_t, A_t^*\right)$.

To characterize the optimal tax schedule in generation t, we consider a small reform which changes the marginal tax rate by $d\tau_t$ on the earnings interval $(z_t, z_t + dz_t)$. As in the static model, such a reform has three welfare effects: a mechanical revenue effect dM_{τ_t}, a behavioral revenue effect dB_{τ_t}, and a direct welfare effect dW_{τ_t}. The three effects may be derived as follows.

The mechanical revenue effect (evaluated at the discounted marginal value of public funds $\beta^t \lambda$) is

given by

$$dM_{\tau_t} = \beta^t \lambda \cdot d\tau_t \cdot dz_t \cdot (1 - H_t(z_t)), \tag{C2}$$

where $H_t(z_t)$ denotes the unconditional distribution function for earnings in generation t.

The behavioral revenue effect (also evaluated at the discounted marginal value of public funds $\beta^t \lambda$) is given by

$$\begin{aligned} dB_{\tau_t} &= \beta^t \lambda \cdot [\tau_t(z_t) - pe_t(z_t)] \cdot dz_t \cdot h_t(z_t)\, dz_t \\ &= -\beta^t \lambda \cdot \left[\frac{\tau_t(z_t) - pe_t(z_t)}{1 - \tau_t(z_t)}\right] \cdot \varepsilon \cdot z_t \cdot d\tau_t \cdot h_t(z_t)\, dz_t, \end{aligned} \tag{C3}$$

where $e_t(z_t) \equiv E_t'(z_t)$ denotes the marginal effect of earnings on education provision, and $h_t(z_t)$ denotes the unconditional density function for earnings in generation t.

Finally, the direct welfare effect can be written as

$$dW_{\tau_t} = -\int_{A_t} \int_{z_t}^{\overline{z}_t} \delta_t \Psi'[.] \cdot d\tau_t \cdot dz_t \cdot dF_{z_t}(z_t' | A_t)\, dF_{A_t}(A_t) = -G_t(z_t, A_t^*) \cdot dM_{\tau_t}. \tag{C4}$$

We may then characterize optimal redistribution as follows:

Proposition 6 (Optimal Redistribution: Dynamic Model). *In each generation t and at each income level z_t, the optimal marginal tax rate satisfies*

$$\frac{\tau_t(z_t) - pe_t(z_t)}{1 - \tau_t(z_t)} = \frac{1 - G_t(z_t, A_t^*)}{\varepsilon \cdot \alpha_t(z_t)}, \tag{C5}$$

where $e_t(z_t) \equiv E_t'(z_t)$ denotes marginal education spending in generation t, ε is the elasticity of earnings with respect to the marginal net-of-tax rate, $\alpha_t(z_t) \equiv z_t h_t(z_t) / (1 - H_t(z_t))$ is the local Pareto parameter of the earnings distribution in generation t, and $G_t(z_t, A_t^*)$ is the average social marginal welfare weight on workers above income level z_t in generation t.

Proof. The result follows from the condition that $dM_{\tau_t} + dB_{\tau_t} + dW_{\tau_t} = 0$ in each generation t. Inserting equations (C2)-(C4) and rearranging terms gives the optimal tax rule in equation (C5). □

The dynamic optimal tax rule in Proposition 6 is similar to the static rule in Proposition 1. The key difference is the term $pe_t(z_t)$ on the left-hand side of equation (C5), creating a direct link between education provision and tax policy. This term captures a fiscal externality of earnings responses operating through education spending on the next generation. If education provision is progressive ($e_t(z_t) < 0$), then lower earnings in generation t imply higher education spending on generation $t+1$, thereby increasing the fiscal externality of raising taxes. All else equal, the extra term in the optimal tax rule implies

that, within a generation, redistribution and predistribution are *policy substitutes*: greater progressivity in education spending implies lower progressivity in income taxation.

To characterize the optimal education schedule in generation t, we consider a small reform which changes education provision by dE_t on the earnings interval $(z_t, z_t + dz_t)$. This affects skill levels in generations $s = t+1, ..., T$ among those descending from individuals in the targeted income range in generation t. The mass of affected individuals in each generation s equals the mass of targeted individuals in generation t, i.e. $h_t(z_t)\, dz_t$. Again, we have three welfare effects of the reform: a mechanical revenue effect dM_{E_t}, a behavioral revenue effect dB_{E_t}, and a direct welfare effect dW_{E_t}.

The mechanical revenue effect equals

$$dM_{E_t} = -\beta^t \lambda \cdot p \cdot dE_t \cdot h_t(z_t)\, dz_t. \tag{C6}$$

The behavioral revenue effect can be written as

$$dB_{E_t} = \sum_{s=t+1}^{T} \beta^s \lambda \cdot \mathbb{E}\left[(\tau_s(z_s) - pe_s(z_s)) \cdot \frac{dz_s}{d\omega_s} \cdot \widehat{\Omega}'_{E_t} \mid z_t \right] \cdot dE_t \cdot h_t(z_t)\, dz_t, \tag{C7}$$

where the expectation $\mathbb{E}[.]$ in each generation s is conditional on descending from individuals at the targeted income level z_t in generation t.

Finally, the direct welfare effect can be written as

$$dW_{E_t} = \sum_{s=t+1}^{T} \delta_s \cdot \mathbb{E}\left[\Psi'[.] \cdot \frac{du_s}{d\omega_s} \cdot \widehat{\Omega}'_{E_t} \mid z_t \right] \cdot dE_t \cdot h_t(z_t)\, dz_t. \tag{C8}$$

These derivations allow us to characterize optimal predistribution:

Proposition 7 (Optimal Predistribution: Dynamic Model). *In each generation t and at each income level z_t, optimal education provision satisfies*

$$\sum_{s=t+1}^{T} \beta^{s-t} \mathbb{E}\left[\widehat{\Omega}'_{E_t} \cdot \left(g_s(z_s, A_s, A_s^*) \cdot \frac{du_s}{d\omega_s} + (\tau_s(z_s) - pe_s(z_s)) \cdot \frac{dz_s}{d\omega_s} \right) \mid z_t \right] = p, \tag{C9}$$

where $e_s(z_s) \equiv E'_s(z_s)$ denotes marginal education spending in generation s, $du_s/d\omega_s$ is the marginal willingness to pay for higher skill in generation s, and $g_s(z_s, A_s, A_s^*)$ is the social marginal welfare weight at earnings z_s and preferences A_s in generation s.

Proof. The result follows from the condition that $dM_{E_t} + dB_{E_t} + dW_{E_t} = 0$ in each generation t. Inserting equations (C6)-(C8) and rearranging terms gives the optimal education rule in equation (C9). □

Proposition 7 generalizes Proposition 2 to allow for intergenerational dynamics. The dynamic rule

differs from the static rule in two respects. First, the social marginal benefit of education provision in generation t (left-hand side of equation (C9)) is a present discounted value of effects in generations $t+1, ..., T$. The trajectory of effects over time depends on the strength of intergenerational skill transmission embodied in the function $\widehat{\Omega}(.)$. Second, the social marginal benefit of education provision accounts for the fact that earnings responses to higher skill in future generations create fiscal externalities through education spending. This effect is captured by the term $pe_s(z_s)$ on the left-hand side of equation (C9).

D Simulation Details

This appendix describes all simulation details. It first describes the baseline case with no preference heterogeneity, and then moves to the cases with preference heterogeneity and parental investments.

D.1 Simulations Without Preference Heterogeneity

D.1.1 Preliminaries and Main Equations

We start by assuming no preference heterogeneity (i.e., assuming a constant A). Note that abstracting from preference heterogeneity affects optimal policy but not the calibration procedure, so the same calibrated values are valid for the simulations with preference heterogeneity. For completeness, we present derivations for an arbitrary constant A, although the baseline simulations normalize, without loss of generality, A to 1.[36] Then, the utility function is given by

$$u(z;\omega_0) = z - T(z) - \frac{A\widehat{\Omega}(\omega_0)}{1+1/\varepsilon}\left(\frac{z}{A\widehat{\Omega}(\omega_0)}\right)^{1+1/\varepsilon}. \tag{D1}$$

Optimal earnings obey

$$z(\omega_0) = A\widehat{\Omega}(\omega_0)\left(1-\tau(z(\omega_0))\right)^{\varepsilon}, \tag{D2}$$

with $\tau(z) = T'(z)$. Indirect utility is given by $U(\omega_0) = \max_z u(z;\omega_0) = u(z(\omega_0);\omega_0)$.

Realized skill $\widehat{\Omega}(\omega_0)$ is given by

$$\widehat{\Omega}(\omega_0) = \Omega(\omega_0, E(\omega_0)) = \omega_0 + B \cdot E(\omega_0)^{\eta}, \tag{D3}$$

with ω_0 background skill. Background skill is distributed across individuals with PDF f. We assume ω_0 is distributed log-normal with parameters $(\mu_{\omega_0}, \sigma_{\omega_0})$, and assume that ω_0 has finite support $[\underline{\omega}_0, \overline{\omega}_0]$, with

[36] This normalization is without loss of generality in the constant-A case because any constant A will yield the same results after a corresponding recalibration of μ_{ω_0} and B, defined below.

$\underline{\omega}_0 > 0$ and $\overline{\omega}_0 < \infty$.[37] Under this utility function, ω_0 is interpreted as potential earnings in the absence of policy (taxes and education).

The government provides $E(\omega_0)$ at a revenue cost of p per unit, so the government budget constraint is given by

$$\int_{\omega_0} T(z(\omega_0)) f(\omega_0) d\omega_0 = R + p \int_{\omega_0} E(\omega_0) f(\omega_0) d\omega_0, \tag{D4}$$

where R is an exogenous revenue requirement. Implicit in this formulation is the normalization of the population size to 1. We denote the Lagrange multiplier associated with the government budget constraint by λ.

The social welfare function is given by

$$W = \int_{\omega_0} \Psi[U(\omega_0)] f(\omega_0) d\omega_0. \tag{D5}$$

In the simulations, we assume $\Psi[U] = U^{1-\rho}/(1-\rho)$, with ρ a measure of inequality aversion.

The solution to the optimal policy problem is characterized by the optimality conditions for education

$$A(1-\tau(z(\omega_0)))^{\varepsilon} \left(\tau(z(\omega_0)) + g(\omega_0) \frac{1-\tau(z(\omega_0))}{1+\varepsilon} \right) B\eta(E(\omega_0))^{\eta-1} = p, \tag{D6}$$

for all ω_0, where $g(\omega_0) = U(\omega_0)^{-\rho}/\lambda$ is the social marginal welfare weight of individuals of type ω_0; the optimality conditions for marginal tax rates

$$\frac{\tau(z(\omega_0))}{1-\tau(z(\omega_0))} = \frac{1}{\varepsilon} \cdot \frac{\widehat{\Omega}'(\omega_0)}{\widehat{\Omega}(\omega_0) f(\omega_0)} \cdot \int_{\omega_0}^{\infty} (1-g(s)) f(s) ds, \tag{D7}$$

for all ω_0, where $\widehat{\Omega}'(\omega_0) = 1 + B\eta(E(\omega_0))^{\eta-1} E'(\omega_0)$ is the Jacobian of the change of variables; and the budget constraint (equation D4) that pins down the intercept of the tax system $T(\underline{\omega}_0)$ given marginal tax rates because of the fundamental theorem of calculus (as is standard, we operationalize this as a universal lump-sum transfer plus a sequence of marginal tax rates). To understand why $\widehat{\Omega}'(\omega_0)$ shows up in equation (D7) and A does not, define $\hat{\omega} = A\omega$. Proposition 1 is written in terms of earnings. To express it in terms of types, the translation involves two changes of variables: from z to $\hat{\omega}$, and from $\hat{\omega}$ to ω_0. We use the local linearization approach of Saez (2001). Assuming $\tau'(z) = 0$, equation (D2) yields $\hat{\omega} \frac{dz}{d\hat{\omega}} = z$. Then, the identity $h(z) dz = f_{\hat{\omega}}(\hat{\omega}) d\hat{\omega}$ implies $\hat{\omega} f_{\hat{\omega}}(\hat{\omega}) = \hat{\omega} h(z) \frac{dz}{d\hat{\omega}} = z h(z)$. Then, when A is constant, we have that $\frac{d\hat{\omega}}{d\omega_0} = A\widehat{\Omega}'(\omega_0)$. Then, from the identity $f(\omega_0) d\omega_0 = f_{\hat{\omega}}(\hat{\omega}) d\hat{\omega}$, we get that $f(\omega_0) = f_{\hat{\omega}}(\hat{\omega}) A\widehat{\Omega}'(\omega_0)$, so that $z h(z) = \hat{\omega} \frac{f(\omega_0)}{A\omega'(\omega_0)} = \frac{\widehat{\Omega}(\omega_0) f(\omega_0)}{\widehat{\Omega}'(\omega_0)}$, which gives the result. Note that, when

[37] While the log-normal distribution is unbounded from above and defined over all positive numbers, we operationalize this by choosing, for given $(\mu_{\omega_0}, \sigma_{\omega_0})$, $\underline{\omega}_0$ and $\overline{\omega}_0$ so that $F(\underline{\omega}_0) = 0.00001$ and $F(\overline{\omega}_0) = 0.99999$.

education policy is non-existent, as in the standard Mirrleesian model, we get that $zh(z) = \omega_0 f(\omega_0)$, which would be the standard result given the assumed utility function.

D.1.2 Discretization

We work with a discretized type distribution. To tractably over-sample the bottom part of the distribution (where most of the log-normal mass is concentrated), we simulate an equally-spaced grid of a standard normal variable $y^i \in y^1, ..., y^N$, with bin width $\Delta_y = y^i - y^{i-1}$ and then recover the background skill grid using $\omega_0^i = \exp\left(\mu_{\omega_0} + y^i \cdot \sigma_{\omega_0}\right)$, with $\omega_0^i \in \left\{\omega_0^1, ..., \omega_0^N\right\}$. The PMF of each type is given by $\pi\left(\omega_0^i\right)$ for $i = 1, ..., N$ and computed as follows: $\Phi\left(y^i + \frac{\Delta_y}{2}\right) - \Phi\left(y^i - \frac{\Delta_y}{2}\right)$. We set $N = 250$, $\Phi\left(y^1\right) = 0.00001$ and $\Phi\left(y^N\right) = 0.99999$. The PMFs are normalized so that $\sum_{i=1}^{N} \pi\left(\omega_0^i\right) = 1$.

D.1.3 Calibration

We fix the values of the following parameters: $(\varepsilon, \rho, \eta) = (0.5, 1, 0.75)$. Then, we choose the parameters $(B, \mu_{\omega_0}, \sigma_{\omega_0})$ to match three empirical moments. Units are calibrated so that every dollar-valued object is measured in thousands of dollars per year. Because the model has an implicit within-individual lifetime steady-state interpretation, we abstract from population growth and assume that individuals engage in K-12 education for 13 years and then work and pay labor income taxes for 30 years (between ages 30 and 60), abstracting from transitional dynamics between education, work, and retirement. Then, we normalize $p = 1$ so that $E(\omega_0)$ measures per-taxpayer annual thousands of dollars spent on their own public K-12 investment, which equals $\frac{13}{30}$ times the annual per-pupil expenditure in public K-12.

To match empirical moments, we solve the model under a fixed policy environment that broadly approximates the actual policy with respect to the vectors T and E. To calibrate the fixed tax system, we use the Piketty, Saez, and Zucman (2018) US-DINA microfiles and fit a linear regression of net taxes paid on pre-tax income to back out an empirical linear tax system. More details on this procedure are found in Section D.3.1 below. This exercise yields $T^{fixed}(z) = 0.303 \cdot z - 12.249$, that is, a universal lump-sum transfer of approximately \$12,250 per capita and a linear tax rate of approximately 30%. To calibrate the fixed education vector, we assume that every type ω_0 receives the same education spending at baseline $E^{fixed}(\omega_0)$. From the US Census Bureau, we get an annual average public spending in K-12 per pupil of \$20,668, so we set $E^{fixed}(\omega_0) = \frac{13}{30} \cdot 20.668 = 8.956$, for all ω_0. This number corresponds to the total K-12 public expenditure, inclusive of current and capital spending, reported in Summary Table 14 of the 2024 Public Elementary-Secondary Education Finance Data (accessed in June 2026).

The three moments we match to calibrate $(B, \mu_{\omega_0}, \sigma_{\omega_0})$ are the following. First, we compute two earnings moments to identify μ_{ω_0} and σ_{ω_0}. We compute labor-income distributions (ages 30–60) in the

Panel Study of Income Dynamics (PSID) for the five most recent waves (2015, 2017, 2019, 2021, 2023). Restricting to individuals with strictly positive labor income, we compute the (properly weighted) mean and standard deviation of the pooled panel (after properly adjusting for inflation so that all incomes are measured in 2024 dollars). We compute a mean of \$90,165 and a standard deviation of \$116,575.

Second, we compute one education moment to identify B. This moment is computed as follows. From the Jackson and Mackevicius (2024) meta-analysis, we get that the average effect of spending an extra \$1,000 per pupil per year for 4 years is a 0.0316 standard-deviation increase in test scores. These numbers are in 2018 dollars. Then, after adjusting for inflation, and using the fixed education vector from above, we get a treatment of $(4 \cdot 1,000 \cdot (1 + 0.249)) / (13 \cdot 20,668) = 1.86\%$, where 0.249 is the inflation adjustment, and $13 \cdot 20,668$ is total expenditure per pupil in K-12 (20,668 annual expenditure per pupil times 13 years). Then, the Jackson and Mackevicius (2024) number in terms of our calibration numbers implies that a 1.86% increase in education expenditure increases test scores by 0.0316 standard deviations. To then translate this into an effect on earnings, we use Chetty, Friedman, and Rockoff (2014) to conclude that a one-standard-deviation increase in test scores increases earnings at age 28 by 12.66%. This calculation comes from two normalizations since that paper estimates the effect of a one-standard-deviation increase in teacher value-added (VA) on test scores to be 0.13, and the causal impact of one standard deviation of teacher VA on age-28 earnings to be \$349.84 (relative to an average of \$21,256), so $0.1266 = (349.84/21,256)/0.13$. Combining all the numbers, we conclude that an increase of 1.86% in $E^{fixed}(\omega_0)$ yields an increase of 0.4% in annual earnings ($0.0316 \cdot 12.66\% = 0.4\%$). Then, the moment we use comes from simulating this thought experiment (increasing education by 1.86%) and estimating the effect on earnings. When doing this, we account for the fact that the numbers from Chetty, Friedman, and Rockoff (2014) are computed in a selected sample with average earnings at 28 of the affected children of \$21,256 (in 2010 dollars) and average earnings of their parents of \$40,808 (in 2010 dollars). Given our steady-state approach, we consider the parents' income, adjusted for inflation, which yields \$58,695, and we further adjust this number by 4% so that it represents average earnings at 45 years old (the average age of the parents in Chetty, Friedman, and Rockoff, 2014 is 40, and DQYDJ's calculation using the Census ASEC/IPUMS CPS suggests a life-cycle adjustment of 4% to translate earnings from age 40 to age 45).[38] This yields average earnings for the population of interest of \$61,029 (68% of our calibration mean earnings).

Using these fixed policy parameters and the moments, we proceed as follows. Given $\left(T^{fixed}, E^{fixed}\right)$ and a guess of $(B, \mu_{\omega_0}, \sigma_{\omega_0})$, we: (1) simulate the ω_0 distribution, (2) apply the Ω function to recover ω, and (3) compute optimal earnings for each type. Then, we (1) compute average earnings, (2) compute the

[38] See DQYDJ's average-income-by-age tabulations of the Census ASEC/IPUMS CPS at https://dqydj.com/average-median-top-income-by-age-percentiles/ (accessed August 2026).

standard deviation of earnings, and (3) identify the population whose average earnings equal \$61,029, increase their fixed education policy by 1.86%, and check by how much their earnings change on average. We repeat this process until the earnings moments match the PSID moments and the education-driven earnings increase in the population of interest matches 0.4%.

The outcome of this procedure yields $(B, \mu_{\omega_0}, \sigma_{\sigma_0}) = (2.7787, 3.9462, 1.0896)$. Moments are matched exactly. With these values, the endpoints of the grid in the baseline simulation are $\omega_0^1 = 0.5$ and $\omega_0^N = 5,400$, so that the distribution of annual potential earnings ranges from \$500 to \$5,400,000.

Finally, we use the calibration outcome to recover the exogenous revenue requirement R. We set R as the residual in the budget constraint under the fixed tax system: $R = 0.303 \cdot \int_{\omega_0} z(\omega_0) f(\omega_0) d\omega_0 - 12.249 - 8.9561 = 6.1149$, so, on top of the lump-sum transfer and the education spending, we calibrate per-capita spending of \$6,115, which is held fixed across the optimal policy simulations.

D.1.4 Optimal Policy

Given calibrated parameters, we simulate our baseline distribution of ω_0 and proceed with the optimal policy problem. In our optimal policy simulations, we set $\rho = 1$, so that social welfare is logarithmic.

The optimal policy problem consists of a vector of unknowns of size $2 \times N + 1$: (1) N education parameters (one for each type, according to equation (D6)), (2) N marginal tax rates (one for each type, according to equation (D7)), and (3) 1 universal lump-sum transfer recovered using the budget constraint.

The routine works as follows. For a guess of the policy parameters, we (1) apply the Ω function to recover ω, (2) compute optimal earnings for each type (which, in turn, allows us to recover their indirect utilities), (3) compute the budget constraint multiplier $\lambda = \sum_{i=1}^{N} \pi\left(\omega_0^i\right) U(\omega_0)^{-1}$ to compute the welfare weights, and (4) evaluate the $2 \times N + 1$ social planner's first-order conditions. The routine iterates until all FOCs are satisfied. The no predistribution benchmark solves the FOCs of the N marginal tax rates and the universal lump-sum transfer after imposing $E(\omega_0) = 0$ for all ω_0.

In evaluating the planner's FOC, we need to adjust the marginal tax rate condition to suit the discrete approximation (see, e.g., Mankiw, Weinzierl, and Yagan (2009); this is not needed for the education FOCs, as they are pointwise in both cases). The adjusted MTR conditions are given by

$$\frac{\tau\left(z\left(\omega_0^i\right)\right)}{1-\tau\left(z\left(\omega_0^i\right)\right)} = \frac{1}{\varepsilon} \cdot \frac{\widehat{\Omega}'\left(\omega_0^i\right)}{\omega\left(\omega_0^i\right) \frac{\pi\left(\omega_0^i\right)}{\Delta_{\omega_0}^i}} \cdot \sum_{j=i}^{N} \left(1 - g\left(\omega_0^j\right)\right) \pi\left(\omega_0^j\right), \tag{D8}$$

where $\Delta_{\omega_0}^i$ is the relevant bin width to transform the PDF to a PMF. Note that the bin width is not constant because our grid is not equally spaced in the ω_0-space and because a non-constant education vector further stretches the distribution. Since $\Delta_{\omega}^i = \widehat{\Omega}'\left(\omega_0^i\right) \cdot \Delta_{\omega_0}^i$, we can implement this numerically by

computing the numerical gradient of $\widehat{\Omega}(\omega_0)$ at ω_0^i as $\Delta_\omega^i = \frac{1}{2}\left(\omega^{i+1} - \omega^{i-1}\right)$ (with one-sided differences at the endpoints).[39]

Likewise, the discrete version of the budget constraint is given by

$$\sum_{i=1}^{N} T\left(z\left(\omega_0^i\right)\right)\pi\left(\omega_0^i\right) = R + p\sum_{i=1}^{N} E\left(\omega_0^i\right)\pi\left(\omega_0^i\right), \tag{D9}$$

where $T\left(z\left(\omega_0^i\right)\right)$ is computed as follows. Let $\tau_i = \tau\left(z\left(\omega_0\right)^i\right)$ and let T_0 be the intercept (universal lump-sum). We define $T\left(z\left(\omega_0^1\right)\right) = \tau_1 \cdot z\left(\omega_0^1\right) - T_0$, and, for $i > 1$, we compute

$$T\left(z\left(\omega_0^i\right)\right) = T\left(z\left(\omega_0^{i-1}\right)\right) + \tau_i \cdot \left(z\left(\omega_0^i\right) - z\left(\omega_0^{i-1}\right)\right). \tag{D10}$$

D.1.5 Comparative Statics

We conduct numerical comparative statics exercises. In each comparative static exercise, we solve for the optimal policy for different parametric combinations and compute the Gini coefficient of ω_0 (constant across iterations), the Gini coefficient of ω (not constant across iterations, as it depends on the optimal education policy), and the Gini coefficient of after-tax income $z - T(z)$ (not constant across iterations, as it depends on both the optimal education policy and the optimal tax system).

Our first comparative static exercise concerns ε. We compute a grid of values for ε ranging from 0.001 to 1. For each value of ε, we solve the optimal policy problem holding the rest of the calibrated parameters fixed at their baseline values. However, we note that, under our utility function, the marginal willingness to pay for increased skill $du/d\omega$ is given by

$$\frac{du}{d\omega} = \frac{A\left(1 - \tau\left(z\left(\omega_0\right)\right)\right)^{1+\varepsilon}}{1+\varepsilon}. \tag{D11}$$

This implies that changing ε changes both the marginal efficiency cost of taxation (which is the comparative-static effect we are interested in) and the marginal welfare benefit of education. To isolate the comparative-static effect of the first force, we solve for the A parameter that, under the new ε, would leave the marginal willingness to pay for skill constant at the initial calibration level (that is, under the fixed tax system $T^{fixed}(z)$ described above), so that each iteration changes both ε and A.[40]

Our second comparative static exercise concerns the education elasticity moment. We proceed as follows. Integrating equation (17) across types, we compute the average elasticity of earnings with respect

[39] The discretization implies that the two transversality conditions cannot both be enforced simultaneously. The sum evaluated at $i = 1$ yields $\tau\left(z\left(\omega_0^1\right)\right) = 0$ exactly, but $\tau\left(z\left(\omega_0^N\right)\right)$ may be positive as $\pi\left(\omega_0^N\right)$ is not necessarily negligible, especially when considering the log-spacing of the grid. This nuance has zero quantitative bite for the aggregate results.

[40] Note that, at calibration, the marginal tax rate is constant for all ω_0, which implies that the baseline value of $\frac{du}{d\omega}$ is a constant, so solving for A given $\frac{du}{d\omega}$ and ε is straightforward.

to education at the calibration values. This gives us an average elasticity of 0.196. Based on this value, we use a grid of this elasticity ranging from 0.1 to 0.3, where for each elasticity value, we compute the B parameter that would give that average elasticity at calibration values, holding μ_{ω_0} and σ_{ω_0} fixed. So, essentially, this is a comparative static exercise with respect to B but expressed in units that are interpretable in elasticity terms.

Our third comparative static exercise concerns the inequality aversion parameter ρ. This comparative static does not consider any recalibration. We approximate the Rawlsian case by considering a long grid and numerically finding the values above which the Gini coefficients under optimal policy converge.

D.2 Simulations With Preference Heterogeneity

Simulations with preference heterogeneity inherit most of the structure and calibration from the baseline case by re-interpreting the former distribution of ω_0 as the distribution of $A\omega_0$ and then considering different combinations of marginal distributions of A and ω_0 that yield similar distributions of $A\omega_0$ but different relative importance of A and ω_0 in explaining the aggregate variance in $A\omega_0$.

In what follows, we define the notation $\widetilde{\omega} = A\omega_0$. We also make the following assumptions: (1) A and ω_0 are independent, and (2) A and ω_0 are distributed log-normal, with parameters (μ_A, σ_A) and $(\mu_{\omega_0}, \sigma_{\omega_0})$. Then, the distribution of $\widetilde{\omega}$ is also log-normal with parameters $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}})$, which we know from the baseline simulation with no preference heterogeneity. That is: $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}}) = (3.95, 1.09)$.

We first describe how we analytically characterize the marginal distributions of A and ω_0, and then describe how we simulate them. Finally, we describe the solution of the optimal policy problem.

D.2.1 Characterizing Marginal Distributions

Let A be log-normal with parameters (μ_A, σ_A) and ω_0 be log-normal with parameters $(\mu_{\omega_0}, \sigma_{\omega_0})$. Define $\widetilde{\omega} = A\omega_0$. If A and ω_0 are independent, known properties yield that $\widetilde{\omega}$ is log-normal with parameters $\left(\mu_A + \mu_{\omega_0}, \sqrt{\sigma_A^2 + \sigma_{\omega_0}^2}\right)$. We also know that, if X is log-normal with parameters (μ, σ), then

$$\mathbb{E}[X] = \exp\left(\mu + \frac{\sigma^2}{2}\right), \tag{D12}$$

$$\mathbb{V}[X] = \left(\exp\left(\sigma^2\right) - 1\right) \cdot \exp\left(2\mu + \sigma^2\right). \tag{D13}$$

The goal of the exercise is to find different combinations of $(\mu_A, \sigma_A, \mu_{\omega_0}, \sigma_{\omega_0})$ such that:

- $\mathbb{E}[\widetilde{\omega}]$ and $\mathbb{V}[\widetilde{\omega}]$ are constant across cases (specifically, equal to the values inherited from the case with no preference heterogeneity).

- $\mathbb{E}[A]$ and $\mathbb{E}[\omega_0]$ are constant across cases (also with values known from the baseline case with no preference heterogeneity).
- The variance of $\widetilde{\omega}$ explained by the variance of A is equal to some value γ. We propose using the R^2 of a (properly weighted) regression of $\widetilde{\omega}$ on A, because it (1) is between 0 and 1 (with 0 the limit equal to our baseline simulation with no preference heterogeneity and 1 the limit case of only preference heterogeneity), (2) is interpretable in levels and can potentially be estimated empirically, and (3) can be determined analytically as follows. The R^2 of the proposed regression is given by

$$R^2 = \frac{\left(\mathbb{C}\left(\widetilde{\omega}, A\right)\right)^2}{\mathbb{V}\left(\widetilde{\omega}\right)\mathbb{V}\left(A\right)}. \tag{D14}$$

  Also, we know that:

$$\mathbb{C}\left(\widetilde{\omega}, A\right) = \mathbb{E}\left[A^2\omega_0\right] - \mathbb{E}\left[A\right]^2\mathbb{E}\left[\omega_0\right] = \mathbb{E}\left[\omega_0\right]\left(\mathbb{E}\left[A^2\right] - \mathbb{E}\left[A\right]^2\right) = \mathbb{E}\left[\omega_0\right]\mathbb{V}\left[A\right], \tag{D15}$$

  which implies that:

$$R^2 = \frac{\left(\mathbb{E}\left[\omega_0\right]\mathbb{V}\left[A\right]\right)^2}{\mathbb{V}\left(\widetilde{\omega}\right)\mathbb{V}\left(A\right)} = \frac{\left(\mathbb{E}\left[\omega_0\right]\right)^2\mathbb{V}\left[A\right]}{\mathbb{V}\left(\widetilde{\omega}\right)} = \frac{\exp\left(\sigma_A^2\right) - 1}{\exp\left(\sigma_A^2 + \sigma_{\omega_0}^2\right) - 1}, \tag{D16}$$

  where the last equality comes from replacing the moment definitions and simplifying terms.

With this strategy, we have five moments and four parameters. However, some of these equations are redundant, yielding an exactly identified set of equations:

- We know, from the baseline case with no preference heterogeneity, that $\mathbb{E}[\widetilde{\omega}]$ and $\mathbb{V}[\widetilde{\omega}]$ determine $\mu_{\widetilde{\omega}} = \mu_A + \mu_{\omega_0}$ and $\sigma_{\widetilde{\omega}}^2 = \sigma_A^2 + \sigma_{\omega_0}^2$ using equations (D12) and (D13). In particular, $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}})$ are the log-normal parameters that we calibrate for ω_0 for the case with $\sigma_A = 0$ and $\mu_A = 0$ (which yield $\mathbb{E}[\omega_0] = \mathbb{E}[\widetilde{\omega}]$, $\mathbb{V}[\omega_0] = \mathbb{V}[\widetilde{\omega}]$, $\mathbb{E}[A] = 1$, and $\mathbb{V}[A] = 0$).
- Then, given $\sigma_{\widetilde{\omega}}^2$ and R^2, we can identify σ_A using equation (D16) (as we can write σ_{ω_0} as a function of $\sigma_{\widetilde{\omega}}^2$ and σ_A).
- Then, given σ_A and $\mathbb{E}[A]$, we can identify μ_A using equation (D12).
- Finally, given (μ_A, σ_A) and $(\mu_{\widetilde{\omega}}, \sigma_{\widetilde{\omega}})$, we can identify $(\mu_{\omega_0}, \sigma_{\omega_0})$ with the identities $\mu_{\widetilde{\omega}} = \mu_A + \mu_{\omega_0}$ and $\sigma_{\widetilde{\omega}}^2 = \sigma_A^2 + \sigma_{\omega_0}^2$.

This procedure provides us with the exact combinations of $(\mu_A, \sigma_A, \mu_{\omega_0}, \sigma_{\omega_0})$ for each R^2 that generate constant distributions of $A\omega_0$ but with controlled relative variances of A and ω_0 to match different values of R^2.

D.2.2 Simulation of Marginal Distributions

The simulations of the log-normal marginal distributions inherit the same procedure as in the baseline case. That is, for each set of parameters $(\mu_A, \sigma_A, \mu_{\omega_0}, \sigma_{\omega_0})$, we generate grids $A^i \in \{A^1, ..., A^N\}$ and $\omega_0^i \in \{\omega_0^1, ..., \omega_0^N\}$, where $A^i = \exp(\mu_A + y^i \cdot \sigma_A)$ and $\omega_0^i = \exp(\mu_{\omega_0} + y^i \cdot \sigma_{\omega_0})$, with y^i and its corresponding density as defined in Section D.1.2 above. We consider $N = 250$ for each marginal distribution, matching the grid size of the baseline case. To double-check our discretization procedure does not compromise the exact analytical result described above, we confirm that the R^2 of the regressions of $A\omega_0$ against A using the discrete grids and their PMFs match the R^2 values with substantial precision.

D.2.3 Optimal Policy

In this case with preference heterogeneity, the solution to the optimal policy problem consists of the education vector of constant size N, a vector of marginal tax rates, and the universal lump-sum transfer. A challenge in this simulation is that earnings in terms of primitives are a function of combinations of A and ω_0 (through the object $A\hat{\Omega}(\omega_0)$) which, in terms of the discrete approximation, implies that the tax schedule should be defined over $N \times N$ points, which is not tractable ($250^2 = 62,500$). In addition to this dimensionality problem, also because of the discrete approximation, the raw PMF of $A\omega_0$ built from the marginal distributions is non-smooth, which generates issues in the optimization of marginal tax rates.

To deal with these challenges, we restrict the profile of marginal tax rates to be defined over a grid of N points (instead of N^2) and proceed as follows:

- For a guess of the policy parameters, we apply the education vector to the marginal distribution of ω_0 to recover the ω vector.
- Then, we recover the raw distribution of $A\omega$ by considering all the unique combinations between the grid points in A and ω, with their corresponding PMFs (given the independence assumption, the PMF of $A\omega(\omega_0)$ is just the PMF of A times the PMF of ω_0).
- Then, we define a grid of $A\omega$ of size N which is log-spaced between the percentiles 0.0001 and 0.9999 of the raw discrete distribution of $A\omega$ (of dimension N^2). These are the N points for which the profile of marginal tax rates is defined. We assign mass to each bin by smoothly grouping the raw bins that correspond to each grid point bin (smoothly means that the mass is distributed between contiguous grid points to avoid discrete jumps in the numerical solver when changes in the education vector reclassify types across grid points).
- Then, for each $A\omega$ combination we recover the marginal tax rate they face by using linear interpolation of the marginal tax rate schedule defined over the coarser grid of size N. Assigning

a marginal tax rate for each type is needed for computing the fiscal externalities relevant to the education FOC and the budget constraint.

- Having recovered the marginal tax rates for all $A\omega$ types, we can compute the actual taxes paid by each type, using a similar sequential logic as in the case with no preference heterogeneity.

- To compute each $A\omega$ welfare weight, we compute a hypothetical distribution of $A\omega$ where we replace A with $A^* = \mathbb{E}[A]$ (since $\mathbb{E}[A] = 1$, this hypothetical distribution coincides with the marginal distribution of ω_0). For this hypothetical distribution, we can recover the marginal tax rate by linear interpolation, compute optimal earnings, compute taxes paid by linear interpolation of the taxes paid schedule on the full $A\omega$ raw grid, and then compute the hypothetical indirect utility. We normalize this indirect utility by the average social marginal utility of this hypothetical distribution to compute the welfare weights. This gives us a vector of ω_0-specific welfare weights.

- Then, we compute the planner's FOCs as follows. For the education FOCs, we recover the average fiscal externality for each type by taking the average across A for a fixed ω_0, using our previously computed matrix of marginal tax rates for all $A\hat{\Omega}(\omega_0)$ combinations. The marginal willingness to pay for increased skill is recovered using the previously computed ω_0-specific welfare weight and computing the hypothetical marginal willingness to pay for skill using the marginal tax rate that applies to each hypothetical case.

- For the marginal tax rate FOCs, we evaluate the FOC in each of the N points of the log-spaced grid of $A\omega$ of size N, considering the average welfare weight within each bin, which is computed numerically given the types that are allocated to the bin.

- Finally, for the budget constraint, we proceed with the exact solution as we know the taxes paid by each $A\omega$ type of the raw distribution.

This procedure is repeated until all FOCs are satisfied. This is conducted separately for each R^2.

D.3 Further Details

D.3.1 Computing the Empirical Tax System Using PSZ

To estimate the empirical tax system $T^{fixed}(z)$ used in the calibration procedure described above, we pool the US-DINA microfiles of Piketty, Saez, and Zucman (2018) between 2016 and 2020 and proceed as follows. First, we keep tax units where all members are between 30 and 60 years old and collapse the data to the tax unit level (single or married filing). Then, we restrict to wage earners by only keeping tax units

where wage income as a share of total fiscal income is between 95% and 105%. Then, we compute income, taxes paid, and transfers received (excluding education expenditures), and then adjust by inflation so values are measured in 2024 dollars. Finally, we run a regression of net taxes paid (taxes minus transfers) on income, controlling for number of children, marital status, and flexible interactions. The coefficient on income is the linear marginal tax rate (30.3%), and the intercept is the lump-sum transfer ($12,249).

E Gini Coefficient and Constant Education

Consider a vector$\{x_i\}_1^n$. The Gini coefficient is given by:

$$G = \frac{\sum_{i=1}^n \sum_{j=1}^n |x_i - x_j|}{2n^2 \bar{x}} = \frac{\sum_{i=1}^n \sum_{j=1}^n |x_i - x_j|}{2n \sum_{i=1}^n x_i}.$$

Consider the following application. Our random variable is $x = \widetilde{\omega} = A\omega_0$. Assume the provision of education is constant (e.g., because there is no variation in welfare weights). So $y = A\omega = A\left(\omega_0 + \overline{E}\right)$, where $\overline{E} = BE\left(\omega_0\right)^\eta > 0$ is the constant provision of education ($E\left(\omega_0\right)$ is the same for all ω_0).

Consider two polar cases. First, there is no preference heterogeneity, so A is constant but ω_0 follows some distribution. Then, we have that:

$$\begin{aligned} G_x = \frac{\sum_{i=1}^n \sum_{j=1}^n |A\omega_{0i} - A\omega_{0j}|}{2n \sum_{i=1}^n A\omega_{0i}} &= \frac{\sum_{i=1}^n \sum_{j=1}^n |\omega_{0i} - \omega_{0j}|}{2n \sum_{i=1}^n \omega_{0i}} \\ &= \frac{\sum_{i=1}^n \sum_{j=1}^n \left|\omega_{0i} + \overline{E} - \omega_{0j} - \overline{E}\right|}{2n \sum_{i=1}^n \omega_{0i}} \\ &> \frac{\sum_{i=1}^n \sum_{j=1}^n \left|\omega_{0i} + \overline{E} - \omega_{0j} - E\right|}{2n \sum_{i=1}^n (\omega_{0i} + \overline{E})} \\ &= \frac{\sum_{i=1}^n \sum_{j=1}^n \left|A(\omega_{0i} + \overline{E}) - A(\omega_{0j} + E)\right|}{2n \sum_{i=1}^n A(\omega_{0i} + \overline{E})} = G_y. \end{aligned} \tag{E1}$$

So, when there is dispersion in ω_0 but not in A, a uniform education policy decreases the Gini. This is what we observe in the limit point of Figure 6.

Now, consider the opposite case where there is full preference heterogeneity, so ω_0 is constant but A follows some distribution. Then, we have that:

$$
\begin{aligned}
G_x = \frac{\sum_{i=1}^{n}\sum_{j=1}^{n}|A_i\omega_0 - A_j\omega_0|}{2n\sum_{i=1}^{n}A_i\omega_0} &= \frac{\sum_{i=1}^{n}\sum_{j=1}^{n}|A_i - A_j|}{2n\sum_{i=1}^{n}A_i} \\
&= \frac{(\omega_0+\overline{E})\sum_{i=1}^{n}\sum_{j=1}^{n}|A_i - A_j|}{(\omega_0+\overline{E})2n\sum_{i=1}^{n}A_i} \\
&= \frac{\sum_{i=1}^{n}\sum_{j=1}^{n}\left|A_i(\omega_0+\overline{E}) - A_j(\omega_0+\overline{E})\right|}{2n\sum_{i=1}^{n}A_i(\omega_0+\overline{E})} = G_y.
\end{aligned}
$$

So, when there is dispersion in A but not in ω_0, a uniform education policy does not affect the Gini. This is what we observe in the limit point of Figure 7.