Optimal income taxation with labor supply responses at two margins: When is an Earned Income Tax Credit optimal?

Emanuel Hansen*

December 20, 2020

Abstract

This paper studies optimal non-linear income taxation in a model with labor supply responses at the intensive (hours, effort) and extensive (participation) margins. It shows that an *Earned Income Tax Credit (EITC)* with negative marginal taxes and negative participation taxes at the bottom is optimal if, first, semi-elasticities of participation are decreasing along the income distribution and, second, social concerns for redistribution from the poor to the very poor are sufficiently weak. This result is driven by a previously neglected trade-off between distortions at the intensive margin and distortions at the extensive margin, i.e., between two aspects of efficiency. Numerical simulations suggest that a strong expansion of the *EITC* for childless singles in the US could be welfare-increasing.

Keywords: Optimal income taxation, Extensive margin, Intensive margin, Earned Income Tax Credit
JEL classification: H21; H23; D82

*University of Cologne, Center for Macroeconomic Research, Albertus-Magnus-Platz, D-50923 Köln, Germany. Email: hansen@wiso.uni-koeln.de. I am very grateful to the Co-Editor, Johannes Spinnewijn, and three anonymous referees for extremely helpful comments and suggestions, and to Felix Bierbrauer, Ulrich Glogowsky, and Dominik Sachs for many fruitful discussions. I also thank Christian Hellwig, Martin Hellwig, Albert Jan Hummel, Sebastian Koehne, Wojciech Kopczuk, Etienne Lehmann, Thorsten Louis, Johannes Münster, Nico Pestel, Paul Schempp, Joel Slemrod, Dezsö Szalay, John Weymark, and numerous participants at conferences and seminars for valuable comments.

1 Introduction

The Earned Income Tax Credit (EITC) is a refundable tax credit that has been introduced in 1975 and extended in several steps over the following decades. Today, it represents one of the largest programs transfering resources towards the poor in the US (Nichols & Rothstein 2015). The EITC differs from traditional welfare programs in that it entails negative participation taxes and negative marginal taxes for low-income earners.[1] Hence, it provides incentives to take up work for non-working persons and to increase working hours for low-income earners. Among political practitioners, there seems to be a broad consensus that the EITC is an effective instrument for fighting poverty. Just a few years ago, President Obama and the Republican Chairman of the House of Representatives Budget Committee, Paul Ryan, even proposed to double the EITC payments for childless workers (see Executive Office 2014, House Budget Committee 2014). However, economists have struggled to rationalize tax-transfer schemes with the properties of the EITC. In particular, a standard result in optimal tax theory is that the marginal tax should never be negative if labor supply responds only at the intensive margin as in Mirrlees (1971).

The main contribution of this paper is to clarify under which conditions an EITC with negative marginal taxes and negative participation taxes at the bottom is optimal. In particular, I derive necessary and sufficient conditions for the optimality of an EITC in a model where labor supply responds at the intensive (hours, effort) and the extensive (participation) margin.[2] Since Saez (2002), it has been known that the sign of the optimal marginal tax is ambiguous and that an EITC *might* be optimal in models with both margins. However, the previous literature has not identified any case in which the optimality of a negative marginal tax is actually *ensured*. Jacquet et al. (2013) even argue that the marginal tax should generally be positive at all income levels. By contrast, I show both analytically and numerically that the marginal tax at the bottom should be negative in empirically relevant cases.

More specifically, I derive three sets of results. First, I study the effects of a small reform that introduces an EITC with negative participation taxes and negative marginal taxes up to some income threshold $\hat{y}$. I show that this reform gives rise to a positive fiscal externality if (i) the semi-elasticity of participation is decreasing along the income distribution. The existing empirical evidence provides support for this condition.[3] The reform also increases welfare if (ii) the income gradient of welfare weights is close enough to zero at the bottom of the income distribution, i.e., if society considers all agents with

[1] The participation tax $T^P : y \mapsto T^P(y)$ measures the difference between the net taxes to be paid at any income level y and 0. Hence, the participation tax is negative at income y if the transfer to a worker with income y is larger than the transfer to an individual with zero income.

[2] The empirical literature provides abundant evidence for the relevance of both margins: "the world is obviously a mix of the two [intensive-margin and extensive-margin] models" (Saez 2002: p. 1054). See also, amongst others, Chetty, Guren, Manoli & Weber (2013).

[3] All available empirical studies find that participation elasticities are either constant or decreasing along the income distribution (see, e.g., Juhn et al. 1991, 2002, Meghir & Phillips 2010, Bargain et al. 2014, Bastani et al. 2020). In both cases, the semi-elasticities of participation are strictly decreasing.

incomes below the threshold $\hat{y}$ similarly deserving. Second, I use a structural model to show that the conditions (i) and (ii) are necessary but not sufficient for an EITC with negative marginal and participation taxes to be optimal. I also show that the optimality of an EITC is *ensured* if, additionally, there is a sufficiently large share of workers with high incomes above $\hat{y}$. As will become clear, this sufficient condition (only) imposes a restriction on the income threshold $\hat{y}$. Third, I calibrate the model to provide numerical simulations of the optimal income tax for childless singles in the US. The simulation results substantiate the quantitative relevance of my previous results. They show that even the recent proposals to double the EITC for childless singles can be rationalized. In the following, I provide a more detailed account of these results.

Tax reform analysis: The fiscal externality of an EITC. In the first step, I use a tax perturbation approach to study the effects of introducing a small EITC. This analysis provides a novel intuition for why and when the EITC can be an attractive policy instrument. For this purpose, I consider a status quo tax that provides equal transfers to non-working agents and low-income workers. I start by studying a reform that increases the transfers to low-income workers and decreases those to non-working agents, thereby introducing negative participation taxes at the bottom. I show that this reform gives rise to a positive fiscal externality if labor supply responds at the extensive margin. I then demonstrate that, given negative participation taxes, a further reform that introduces negative marginal taxes at the bottom yields a positive fiscal externality as well.

For an intuitive understanding of the latter key insight, consider the reform illustrated in Figure 1. In this figure, the solid blue line depicts a pre-reform tax schedule such that all workers with positive incomes below some threshold y_p receive identical work subsidies, i.e., face zero marginal taxes and strictly negative participation taxes. For example, threshold y_p can be thought of as the poverty line. Then, consider a small tax reform that introduces a negative marginal tax in such a way that the work subsidies for workers with incomes close to zero are reduced, while the work subsidies for workers with slightly higher incomes below y_p are increased (see dashed red line in Figure 1). Transfers to non-working agents and taxes for higher-income workers remain constant.

Figure 1: A reform that introduces negative marginal taxes.

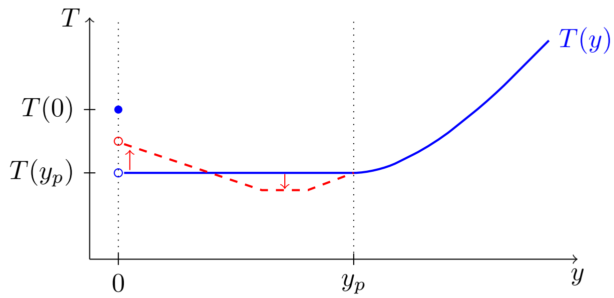

Notes: Figure 1 illustrates a reform that introduces negative marginal taxes at low incomes. The solid blue line represents the pre-reform tax schedule, which involves zero marginal taxes and negative participation taxes at all incomes below y_p. The dashed red line represents the post-reform tax schedule.

The reform induces labor supply responses at both margins. First, the changes in work subsidies lead to responses at the extensive margin: Some workers with very low incomes close to zero leave the labor market, and some non-working agents enter the labor market with higher incomes below y_p. Given a work subsidy prior to the reform, the former responses yield a gain in tax revenue, while the latter responses yield a revenue loss. The net effect on tax revenue is strictly positive if the participation responses of the former group are larger than those of the latter group (in line with the empirical evidence). Second, the reduction in marginal taxes induces workers to increase their working hours. With a negative marginal tax, this intensive-margin response yields a loss in tax revenue. Thus, there is a previously neglected tradeoff: tax revenue is positively affected by the extensive-margin responses and negatively affected by the intensive-margin responses. Starting from a tax schedule with zero marginal taxes, however, the revenue loss due to intensive-margin responses is initially negligible. Consequently, the introduction of a small negative marginal tax yields a strictly positive fiscal externality: a revenue gain due to behavioral responses that can be redistributed back to the population.

The welfare effect of this reform also depends on the government's redistributive preferences, as the reform redistributes resources from workers with very low incomes to workers with slightly higher incomes. If the income gradient of welfare weights is large, i.e., society considers the workers with very low incomes much more deserving than workers with incomes close to the poverty line y_p, the redistributive loss outweighs the gains from the positive fiscal externality of the reform. If the income gradient of welfare weights is close to zero, by contrast, the redistributive loss is small and the fiscal externality dominates: The introduction of a negative marginal tax increases welfare. In this case, it appears plausible that negative marginal taxes for low-income workers are also part of the optimal tax policy. However, there might exist other tax reforms that increase welfare even more, e.g., reforms that introduce a small *positive* marginal tax at the bottom and reduce marginal taxes for high-income workers. Hence, by construction, the tax reform analysis in this step provides a novel intuition but does not allow to identify the properties of the optimal tax.

Optimal tax analysis: Necessary and sufficient conditions for an EITC. In the second step, therefore, I take a mechanism design approach to solve for the welfare-maximizing income tax. For this purpose, I use a structural model in which the agents differ both in marginal costs of providing effort as in Mirrlees (1971) and in fixed costs of working as in Diamond (1980). This two-dimensional heterogeneity gives rise to labor supply responses at both margins. To simplify the exposition, I initially focus on a stylized version of this model with three skill groups and particularly simple functional forms. Then, I generalize my results to a model with a finite but arbitrarily large number of skill groups and without assumptions on functional forms, which is flexible enough to be calibrated to empirical moments.

In both model versions, I derive necessary conditions for the optimality of negative participation taxes and negative marginal taxes up to some income threshold $\hat{y}$, which mirror those from the previous tax reform analysis: An EITC *can only* be optimal if (i) the semi-elasticity of participation is decreasing over the skill distribution, and (ii) the skill gradient of welfare weights is close enough to zero at the bottom. I also identify a sufficient condition that did not appear before: The optimality of an EITC is *ensured* if, on top of the two previous conditions, the population share of workers with earnings above $\hat{y}$ is large enough. Intuitively, this implies that the tax base for taxes on high-income earners is large so that an EITC is easy to finance. In the model with three skill groups, this sufficient condition is satisfied if both the population share and the productivity of the highest-skilled group are large enough. In the general model, the condition is always satisfied for an EITC with a small income threshold $\hat{y}$, but not for an EITC with a large threshold. Hence, the sufficient condition turns into a restriction on the optimal size of the EITC. A supplementary analysis shows that an EITC remains optimal in the limit case in which the discrete skill set converges to a continuous set (see Online Appendix D.1).

Numerical analysis: Optimal income taxes for childless singles in the US. In the third step, I provide numerical simulations of the optimal income tax for a calibrated version of the general model. These simulations allow me, first, to assess the quantitative relevance of my analytical results and, second, to show that an EITC can even be optimal if only the necessary conditions discussed above are met. Specifically, I calibrate the model to the subgroup of childless singles in the US, using (a) data from the March 2016 Current Population Survey and (b) estimates of labor supply elasticities from the empirical literature. Under 2015 US tax rules, childless singles face a small EITC with negative marginal taxes for incomes below $\$6,580$, negative participation taxes for incomes below $\$14,820$ and a maximum tax credit of \$503. According to my benchmark simulations, by contrast, marginal taxes should be negative for annual incomes up to $\$15,000$ and participation taxes should be negative for incomes up to $\$34,000$. The maximum tax credit should even amount to levels around $\$2,000$. Hence, the recent proposals to strongly expand the EITC can indeed be rationalized (Executive Office 2014, House Budget Committee 2014).

A comprehensive sensitivity analysis shows that variations in intensive-margin and extensive-margin elasticities and in the discretization of the skill set (i.e., the number of skill types) have only limited effects on the quantitative properties of the optimal tax schedule. By contrast, the assumed preferences for redistribution have a crucial impact: An EITC is optimal if the welfare weights are flat at the bottom, but not if the welfare weights are steeply decreasing at the bottom (as in the simulations by Saez 2002 and Jacquet et al. 2013). Additionally, I provide the results of numerical simulations for the subgroup of single parents in the US.[4]

[4] In a nutshell, I find that the current EITC for single parents – which is much more generous than for childless singles – can be rationalized, but should not be expanded further (see Online Appendix D.5).

How relevant are these results? A main take-way from this paper is that, if labor supply responds at both margins, the optimal design of taxes for low-income earners depends crucially on the income gradient of welfare weights. This observation gives rise to the question what a convincing assumption on the income gradient of welfare weights is. Is it possible to rationalize welfare weights that are more or less flat at the bottom of the income distribution? The answer to this question depends on the type of social objective considered. A Rawlsian planner cares only about the utility of the non-working, i.e., the implied welfare weights drop discontinuously at the bottom of the income distribution. Similarly, in a model with only one-dimensional heterogeneity in skills, the assumptions of a utilitarian welfare function and a standard utility function imply that marginal welfare weights are convexly decreasing, i.e., particularly steep at the bottom (as the ones considered by Saez 2002 and Jacquet et al. 2013). In a two-dimensional model as the one I study, this is less obvious: If agents with higher skill types are disadvantaged in another dimension (e.g., because they face on average higher fixed costs than low-skill workers), marginal welfare weights may be flat or even increasing at the bottom (see Choné & Laroque 2010). Besides, if the social objective is to alleviate poverty, the social planner cares similarly for everyone below the poverty line. Correspondingly, Saez & Stantcheva (2016) suggest to model this objective by means of generalized welfare weights that are flat at the bottom. It has repeatedly been argued that a poverty alleviation objective appears more descriptive for real-world policy choice than standard welfarist objectives (Besley & Coate 1992, 1995, Kanbur et al. 1994, Saez & Stantcheva 2016). The recent political debate supports this view: Proponents of an EITC expansion mainly emphasize its potential to lift workers above the poverty line and reduce the depth of poverty for others (e.g., Executive Office 2014, House Budget Committee 2014). Summing up, my results suggest that an EITC is an attractive instrument for real-world policy makers that are interested in reducing poverty, but less so for the utilitarian social planner that is typically considered in economics textbooks.

Outline. The paper proceeds as follows. Section 2 briefly reviews the related literature. Section 3 introduces the basic framework. Section 4 studies whether the introduction of a small EITC is welfare-increasing, while Section 5 provides my results on the optimal income tax. Section 6 provides numerical simulations for a calibrated version of the model. Section 7 concludes. Appendices A and B contain the formal proofs for the tax reform analysis and the optimal tax analysis in the stylized model. Online Appendices C and D provide the formal proofs for the general model and supplementary material, respectively.

2 Related literature

This paper contributes to two strands of the literature. First and foremost, it contributes to the small literature on optimal income taxation in settings with labor supply responses at the intensive and extensive margins. The most closely related papers are given by Saez (2002) and Jacquet et al. (2013). Saez (2002) studies a model with a finite number of available income levels interpreted as occupations. He uses a perturbation approach to derive an optimal tax formula. Jacquet et al. (2013) study a two-dimensional model with continuous sets of skill types and fixed cost types. My model differs from theirs mainly in that I consider a discrete skill set. They use a mechanism design approach to derive first-order conditions that characterize the optimal allocation. Neither the optimal tax formula nor the first-order conditions pin down the optimal sign of the marginal tax, however.[5] As a result, both papers can only note that the optimal marginal tax *might be negative*, but neither of them provides examples in which this is indeed the case. Instead, Jacquet et al. (2013) provide a sufficient condition for optimal marginal taxes to be *positive* everywhere below the top.[6] Both papers also perform numerical simulations in which the optimal tax is always monotonically increasing over the entire range of positive income levels. Based on these simulations, Jacquet et al. (2013) argue that marginal taxes should even be positive if their sufficient condition is not met.[7] By contrast, I show that an EITC with negative marginal taxes is optimal under empirically plausible conditions.

Second, a few previous papers provide alternative rationalizations of an EITC with negative marginal taxes. In Beaudry et al. (2009), the social planner uses work subsidies to redistribute from agents working in an informal (black) labor market to formally employed workers with low wages. In Choné & Laroque (2010), the planner wants to transfer resources from low-skill, low-income workers to high-skill, high-income workers because the latter group is disadvantaged in some other dimension. To achieve this redistribution from the poor to the rich, she has to use negative marginal taxes. In Lockwood (2020), the agents suffer from a present bias that makes them underestimate some delayed benefits of their work effort (e.g., future promotions). As a result, they provide inefficiently little effort compared to non-biased agents. In this model, negative marginal taxes can increase welfare because they help to correct for the individual optimization errors. In my model, there is no informal labor market, welfare weights are monotonically decreasing along the skill distribution, and agents behave individually rational. Instead, I demonstrate that negative marginal taxes can be optimal as soon as the standard Mirrlees (1971) model is augmented to account for labor supply responses at the extensive margin.

[5]In Online Appendices D.7 and D.8, I present versions of the optimal tax formula and the first-order conditions, and explain why they do not allow to sign the optimal marginal tax.

[6]Specifically, Jacquet et al. (2013) identify a monotonicity condition under which optimal marginal taxes are ensured to be positive. To verify this condition, one has to evaluate the skill-specific marginal welfare weights and semi-elasticities at the second-best allocation. As they acknowledge, this is not trivial because welfare weights and semi-elasticities are in general endogenous objects.

[7]Besides, both papers find that the optimal transfer to low-income workers is sometimes higher than the transfer to non-working agents. Saez (2002) refers to the difference between the transfers to non-working agents and the lowest-earning workers as the *marginal tax at the bottom*. In contrast, Jacquet et al. (2013) and the subsequent literature refer to this difference as the *participation tax at the bottom*. I stick to the latter terminology.

3 The basic framework

This section introduces a basic framework for the normative analysis of income taxes with labor supply responses at both margins. I start by defining the economic environment. Then, I derive individually optimal earnings choices. Based on this framework, I study the welfare and revenue effects of introducing a small EITC in Section 4 and the properties of the optimal (welfare-maximizing) tax schedule in Section 5.

3.1 Environment

There is a continuum of agents of mass one. Agent i enjoys consumption c^i and suffers from a cost of providing output y^i. This cost can be separated into a variable effort cost that is continuously increasing in y^i and a fixed cost that the agent bears if and only if she provides strictly positive output. Formally, the agent's preferences can be represented by the utility function

$$u(c^i, y^i; \omega^i, \delta^i) = c^i - h(y^i, \omega^i) - \mathbb{1}_{y^i>0}\ \delta^i\ . \tag{1}$$

Function h captures the variable effort cost, which depends on the output level y^i and the agent's *skill type* $\omega^i \in \Omega$. I impose the standard assumptions that h satisfies $h(0, \omega) = 0$, $h_y(y, \omega) > 0$, $h_{yy}(y, \omega) > 0$, $h_\omega(y, \omega) < 0$ and $h_{y\omega}(y, \omega) < 0$ for all $y > 0$ and $\omega \in \Omega$. Parameter $\delta^i \in \Delta$ specifies the agent's *fixed cost type*.

Each agent is privately informed about her skill type ω^i and her fixed cost type δ^i. The type set $\Omega \times \Delta$ is a subset of $\mathbb{R}^2$. Specifically, there is a continuous set of fixed costs $\Delta = \left[\underline{\delta}, \bar{\delta}\right]$. The skill set Ω may be discrete or continuous. In any case, I denote by $\underline{\omega}$ and $\bar{\omega}$ the lowest and the highest skill type, respectively. The joint cross-section distribution $G : \Omega \times \Delta \mapsto [0, 1]$ of two-dimensional types in the population is commonly known and has full support on $\Omega \times \Delta$. To ensure that labor supply responds at the extensive margin, I assume that

$$\begin{aligned} \underline{\delta} &< \max\nolimits_{y>0}\ y - h\left(y, \underline{\omega}\right)\ , \\ \bar{\delta} &> \max\nolimits_{y>0}\ y - h\left(y, \bar{\omega}\right)\ . \end{aligned} \tag{2}$$

This assumption and $h_\omega < 0$ jointly ensure that the extensive margin is relevant in each skill group $\omega \in \Omega$: Under laissez-faire, some agents with skill type ω provide strictly positive output, while other agents with the same skill type provide zero output. As will become clear, the same is true for the optimal allocation.

The government redistributes resources by setting a non-linear income tax schedule $T : y \mapsto T(y)$ and letting each agent choose her individually optimal income $y^*(\omega, \delta \mid T) := \arg\max_y u(y - T(y), y; \omega, \delta)$. The resulting allocation is evaluated based on a social welfare function. Specifically, I assume that social welfare is given by

$$E_{\Omega\times\Delta}\left[\alpha(\omega)\ V(\omega, \delta \mid T)\right]\ , \tag{3}$$

where $V(\omega, \delta \mid T)$ denotes the indirect utility of an agent with type (ω, δ) facing tax schedule T, and $\alpha(\omega)$ denotes an exogenous welfare weight associated to every agent with skill type ω. I assume that welfare weights are strictly positive, decreasing over the skill distribution and normalized to have an average value of 1. The government's budget constraint is given by

$$E_{\Omega\times\Delta}\left[T\left(y^*(\omega, \delta \mid T)\right)\right] \geq B \tag{4}$$

where B is an exogenous revenue requirement. For the remainder of this paper, I set $B = 0$.

Note that I have imposed three simplifying assumptions. First, I have assumed that fixed costs enter the utility function in an additively separable way. This implies that the fixed cost δ only affects an agent's decision whether or not to work at all, but not her income choice conditional on working. Second, I have imposed quasi-linearity of utility in consumption, thereby assuming away income effects in labor supply. Third, I consider a welfare function with exogenous welfare weights that depend only on the agents' skill types, but neither on their fixed cost types nor on the bundles they are assigned.[8] All three assumptions simplify the analysis, but do not eliminate the problems in signing the optimal marginal tax. The first assumption follows the random participation approach by Rochet & Stole (2002) and is standard in the literature on optimal taxation with two margins. The other two assumptions have been used in many other optimal tax papers, including Diamond (1998) and Saez (2001).[9]

3.2 Individual labor supply choices

Fix a tax schedule $T : \mathbb{R}_0^+ \to \mathbb{R}$ that is continuously differentiable and weakly convex for all $y > 0$.[10] An agent with type (ω, δ) maximizes her utility (1) by choosing the output level

$$y^*\left(\omega, \delta \mid T\right) = \begin{cases} y_T(\omega) & \text{if} \quad \delta \leq \delta_T(\omega) \\ 0 & \text{if} \quad \delta > \delta_T(\omega), \end{cases} \tag{5}$$

where the conditional optimum $y_T(\omega)$ is implicitly defined by the first-order condition

$$1 - T'\left[y_T(\omega)\right] = h_y\left(y_T(\omega), \omega\right) \ , \tag{6}$$

and the participation threshold $\delta_T(\omega)$ is defined by

$$\delta_T(\omega) \coloneqq y_T(\omega) - h\left(y_T(\omega), \omega\right) - T^P\left(y_T(\omega)\right) \ , \tag{7}$$

[8] In the working paper version (Hansen 2018), I relax this assumption to allow for marginal welfare weights that are endogenous or depend on both type dimensions.

[9] Jacquet et al. (2013) allow for income effects as well as endogenous welfare weights.

[10] The assumptions of continuous differentiability and weak convexity simplify the exposition. I maintain them in the tax reform analysis in Section 4. In the optimal tax analysis in Sections 5 and 6, by contrast, I allow for concavity and for kinks in the tax schedule.

with $T^P(y) := T(y) - T(0)$ denoting the participation tax at income $y > 0$. Hence, conditional on providing positive output, an agent's optimal intensive-margin choice depends only on her skill type ω: All working agents with skill type ω choose the same income level $y_T(\omega)$. The extensive-margin decision whether to provide positive output depends on whether her fixed cost type δ is below or above the skill-specific participation threshold $\delta_T(\omega)$. As a consequence, the skill-specific participation rate is given by $p_T(\omega) = G_\delta(\delta_T(\omega) \mid \omega)$. Correspondingly, the share of non-working agents is given by $E_\Omega\left[1 - p_T(\omega)\right]$.

Tax reforms can induce labor supply responses at both margins in each skill group. As usual, I measure intensive-margin responses by the skill-specific elasticity of income with respect to the retention rate,

$$\varepsilon_T(\omega) := \frac{\partial y_T(\omega)}{\partial\left[1 - T'(y_T(\omega))\right]} \, \frac{1 - T'(y_T(\omega))}{y_T(\omega)} \, . \tag{8}$$

Following Jacquet et al. (2013), I measure extensive-margin responses by the skill-specific semi-elasticity of participation,

$$\eta_T(\omega) := \frac{\partial p_T(\omega)}{\partial\left[y_T(\omega) - T^P(y_T(\omega))\right]} \, \frac{1}{p_T(\omega)} \, , \tag{9}$$

which gives the percentage change in the participation rate $p_T(\omega)$ that results from a one-unit increase in the net-of-tax labor income $y - T^P(y)$.[11] In general, both ε_T and η_T vary with skill type ω and tax schedule T.

4 The welfare effects of introducing an EITC

In this section, I study the welfare effects of introducing a small EITC. For this purpose, I assume that there is a status-quo tax schedule with zero marginal and participation taxes at the bottom. While this tax schedule is generically not optimal, it allows me to investigate the effects of a small EITC in two steps. In the first step, I clarify the conditions under which the introduction of negative participation taxes for low-income workers increases tax revenue and welfare. In the second step, I show that, given a negative participation tax, the introduction of negative marginal taxes is revenue- and welfare-increasing under plausible conditions. The analysis is based on the variational approach introduced by Piketty (1997) and Saez (2001) and its extension to multi-bracket tax reforms by Bierbrauer et al. (2020).

The results of this section are expressed in terms of *sufficient statistics*: They only depend on objects that describe redistributive preferences and labor supply responses for small perturbations of the tax system (Kleven 2020*b*). Correspondingly, they could be derived based on any model of individual behavior that gives rise to intensive-margin and

[11] Alternatively, extensive-margin responses could be measured by the (standard) elasticity of participation, $\pi_T(\omega) = \eta_T(\omega)\left[y_T(\omega) - T^P(y_t(\omega))\right]$.

extensive-margin responses, and for any social objective (see Saez & Stantcheva 2016).[12] To maximize the clarity of the argument, I nevertheless use the model introduced in Section 3, which I subsequently also use to study optimal taxes. The results do not depend on whether the skill set Ω is discrete or continuous, either. Following most of the literature, I focus on the case of a continuous skill set and denote by $\underline{y} \coloneqq y_T(\underline{\omega})$ and $\bar{y} \coloneqq y_T(\bar{\omega})$ the incomes of the lowest-skilled and the highest-skilled workers, respectively.

To study the introduction of a small EITC, I assume that there is a pre-reform tax schedule T with the following properties.

Assumption 1. *The pre-reform tax schedule T is continuously differentiable and weakly convex for all $y > 0$. Moreover, it satisfies $T(y) - T(0) = t_p$ for all $y \in (0, y_p)$, where y_p is an income threshold in $\left(\underline{y}, \bar{y}\right)$.*

Assumption 1 implies that, prior to the reform, the marginal tax $T'(y)$ is zero and the participation tax is identical for all incomes below y_p. Below, I will focus on cases where the participation tax t_p at the bottom is either zero or negative. In the latter case, there is a work subsidy for low-income workers. With a continuous skill set Ω, Assumption 1 implies that T implements an income distribution F_y with density f_y on the interval $\left[\underline{y}, \bar{y}\right]$ and a mass point $F_y(0)$ at zero income. The following subsections make use of two additional pieces of notation. First, I denote by $\hat{\alpha}(y', y'')$ the average welfare weight of workers with incomes between y' and y'', and by $\hat{\alpha}_0$ the average welfare weight of non-working agents under tax schedule T. Second, I denote by $\eta(y')$ the semi-elasticity of participation for workers with income y', and by $\hat{\eta}(y', y'')$ the average semi-elasticity of participation for workers with incomes between y' and y''.[13]

4.1 Introducing negative participation taxes

I start by investigating the welfare effects of a reform that introduces negative participation taxes for low-income workers. For this purpose, I consider a reform that redistributes resources from non-working agents with zero income to low-income workers with strictly positive incomes below some threshold $y_a < y_p$. Figure 2a illustrates this reform by depicting a pre-reform tax schedule (solid blue line) and a post-reform tax schedule (dashed red line). Specifically, the reform increases the tax of non-working agents by an amount $\tau\ell$. It decreases the tax liability for all incomes in the range $(0, y_a - \ell)$ by the amount $\tau\ell\phi_a$ and increases the marginal tax in the income range $(y_a - \ell, y_a)$ by $\tau\phi_a$, where the weighting parameter ϕ_a is set to $F_0(0)/[F_y(y_a) - F_y(\underline{y})]$. The reform does not affect tax liabilities at incomes above y_a. A potential revenue gain (loss) due to the reform is compensated by a lump-sum transfer to (tax on) the entire population. Applied to a pre-reform tax that

[12] More specifically, the following results hold for any model such that there are no income effects in labor supply. Income effects would not alter the key insights: The size of the fiscal externality of a small EITC changes, but it remains strictly positive.

[13] Formal definitions of $\hat{\alpha}(y', y'')$, $\hat{\alpha}_0$, $\eta(y)$ and $\hat{\eta}(y', y'')$ are provided in Appendix A.

Figure 2: The introduction of a small EITC.

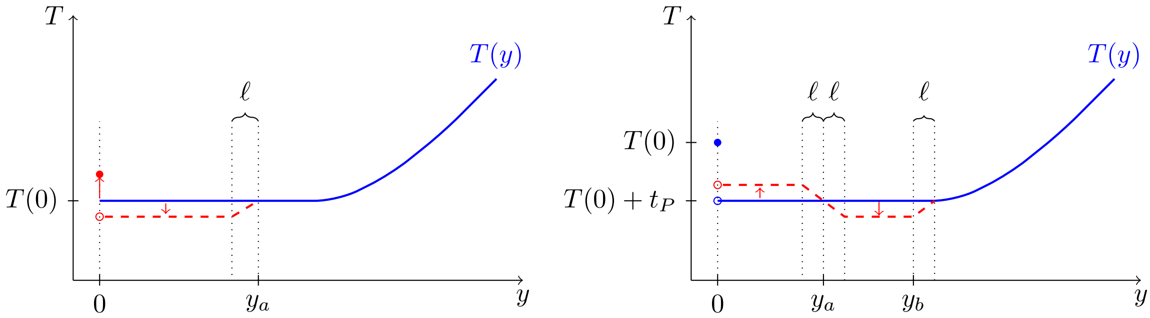

(a) Negative participation taxes. (b) Negative marginal taxes.

Notes: Figure 2a in the left panel illustrates a reform that introduces negative participation taxes for positive incomes below y_a. Figure 2b in the right panel illustrates a reform that introduces negative marginal taxes around income level y_a. In both figures, the solid blue line represents the pre-reform tax schedule T that satisfies Assumption 1, while the dashed red line represents the post-reform tax schedule.

satisfies Assumption 1 with $t_P = 0$, this reform introduces a negative participation tax for low-income earners (see dashed red line in Figure 2a).

The following lemma studies the welfare effect of this tax reform, focusing on the limit case where both τ and ℓ converge to zero. The details of the formalization and the proof are relegated to Appendix A.

Lemma 1. *Let the pre-reform tax T satisfy Assumption 1 with a zero participation tax $t_P = 0$ for all incomes below y_p. The introduction of a negative participation tax for incomes below $y_a < y_p$ is welfare-increasing if and only if*

$$\int_{y_p}^{\bar{y}} f_y(y)\eta(y)T^P(y)dy > \left[F_y(y_a) - F_y(\underline{y})\right] \; \left(\hat{\alpha}_0 - \hat{\alpha}(\underline{y}, y_a)\right) \; . \tag{10}$$

By Lemma 1, the welfare effect of introducing negative participation taxes at the bottom depends on the relative size of two terms. The term on the left-hand side of (10) captures a positive fiscal externality that is driven by previously non-working agents who enter the labor market with incomes above y_p. As the participation tax was positive above y_p before the reform, these extensive-margin responses yield a revenue gain that can be redistributed back as a lump-sum transfer. The size of this fiscal externality depends on the levels of the semi-elasticity of participation η and the participation tax T^P for workers with incomes above y_p. It would vanish if labor supply would only respond at the intensive margin, i.e., if $\eta(y) = 0$ for all incomes above y_p. With zero participation taxes at the bottom, $t_p = 0$, there are initially no further fiscal externalities. The term on the right-hand side captures the direct (mechanical) welfare loss of redistributing resources from non-working agents with welfare weight $\hat{\alpha}_0$ to low-income workers with average welfare weight $\hat{\alpha}(\underline{y}, y_a)$. Its size depends on the income gradient of the welfare weights and the length of the interval $(\underline{y}, y_a)$. In particular, the mechanical welfare loss is zero if society associates the same welfare weight $\alpha_p > 1$ to all agents with incomes below y_p. In this case, the reform is unambiguously welfare-increasing whenever labor supply responds at the extensive margin.

For a more complete understanding of Lemma 1, note that a reform with the properties depicted in Figure 2a in general has mechanical effects and behavioral effects due to both intensive-margin and extensive-margin responses on tax revenue. First, the mechanical effects are given by the revenue gain from the tax hike for non-working agents and the revenue loss from the tax cut for low-income workers. The population share of the former group is given by $F_y(0)$, the share of the latter group by $F_y(y_a) - F_y(\underline{y})$. With the weighting parameter ϕ_a set to $F_y(0)/[F_y(y_a) - F_y(\underline{y})]$, both mechanical effects offset each other. Second, the reform has a behavioral effect that results from the intensive-margin responses of workers with incomes close to y_a, for whom the marginal tax is increased. Under Assumption 1, however, these responses initially do not affect the workers' tax payments because the marginal tax $T'(y_a)$ was equal to zero before the reform. Third, the reform has a behavioral effect that results from extensive-margin responses. As the reform has reduced the participation tax at all income levels, taking up work at any positive income has become more attractive relative to non-working. As a result, some previously non-working agents enter the labor market with earnings levels between $\underline{y}$ and $\bar{y}$. For agents entering the labor market with an income below y_p, these extensive-margin responses initially have no effect on tax revenue as the pre-reform participation tax was given by $t_P = 0$. For agents entering with an income above y_p, by contrast, tax revenue is increased because the pre-reform participation tax was positive. As these revenue gains are transferred uniformly to the entire population, the reform induces the positive fiscal externality captured by the left-hand side of (10). The total effect on welfare is positive if this fiscal externality exceeds the redistributive loss on the right-hand side of (10).[14]

4.2 Introducing negative marginal taxes

In the next step, I show that the introduction of negative marginal taxes can be welfare-increasing if the participation tax at the bottom is negative (e.g., as a result of the reform studied above). For this purpose, consider a tax reform that redistributes resources from workers with incomes below some threshold y_a to workers with slightly higher incomes between y_a and another threshold y_b, implemented by a reduction in the marginal tax around y_a. Figure 2b illustrates this reform by depicting both the pre-reform tax schedule (solid blue line) and the post-reform tax (dashed red line). Specifically, the reform reduces the marginal tax in the interval $(y_a - \ell, y_a)$ by τ and increases the tax liability for incomes in $(\underline{y}, y_a - \ell)$ by the amount $\tau\ell$. Moreover, the reform decreases the marginal tax in the interval $(y_a, y_a + \ell)$ by $\tau\phi_b$ and increases the marginal tax in the interval $(y_b, y_b + \ell)$ by the same amount. As a result, tax levels for incomes between $y_a + \ell$ and y_b are reduced by the amount $\tau\ell\phi_b$, while taxes above $y_b + \ell$ remain constant constant. The weighting

[14]If the participation tax at the bottom t_p is different from zero, the reform gives rise to an additional fiscal externality due to the extensive-margin responses of low-income workers. For t_p above zero, this externality is positive and the welfare effect of the reform is further increased. For t_p below zero, the externality is negative and the welfare effect is subsequently decreased.

parameter ϕ_b is set to $\phi_b = [F_y(y_a) - F_y(\underline{y})]/[F_y(y_b) - F_y(y_a)]$. Again, I assume that a potential revenue gain (loss) of the reform is compensated by a lump-sum transfer to (tax on) the entire population.

Importantly, the solid blue line in Figure 2b depicts a pre-reform tax schedule that involves zero marginal taxes and negative participation taxes at the bottom, $t_p < 0$. Applied to this pre-reform tax, the reform introduces negative marginal taxes in the interval $(y_a - \ell, y_a + \ell)$. For the case with negative participation taxes $t_p < 0$ in the status quo, moreover, the reform reduces the work subsidy for the lowest-income workers and increases the work subsidy for workers with slightly higher incomes. The resulting tax schedule is a stylized version of the EITC in the US (see Figure 2b).

The following lemma provides a condition under which this reform increases welfare, focusing again on the limit case of a small reform with $\tau \to 0$ and $\ell \to 0$. The details of the formalization and the proof are relegated to Appendix A.

Lemma 2. *Let the pre-reform tax T satisfy Assumption 1 and let $y_a < y_b < y_p$. The introduction of a negative marginal tax around a low-income level y_a is welfare-increasing if and only if*

$$- t_p \left[\hat{\eta}(\underline{y}, y_a) - \hat{\eta}(y_a, y_b)\right] > \hat{\alpha}(\underline{y}, y_a) - \hat{\alpha}(y_a, y_b) \ . \tag{11}$$

Lemma 2 implies that the welfare effect of introducing negative marginal taxes depends on the relative sizes of two terms as well: a fiscal externality on the left-hand side, and a redistributive loss on the right-hand side of (11). The fiscal externality results because the reform makes it less attractive to take up work with an income below y_a rather than to remain non-working, but more attractive to take up work with an income above y_a. Assume that, prior to the reform, the participation tax t_p at the bottom was negative as in Figure 2b. Then, the fiscal externality on the left-hand side of (11) is positive if the semi-elasticity of participation is decreasing along the income distribution, in line with the empirical evidence.[15] The size of the externality depends on the level of the work subsidy $-t_p$ and on the income gradient of the semi-elasticity of participation. It would vanish completely if labor supply would not respond at the extensive margin, $\hat{\eta}(\underline{y}, y_a) = \hat{\eta}(y_a, y_b) = 0$. The right-hand side of (11) captures the (mechanical) welfare loss of redistributing resources from workers with incomes below y_a to workers with incomes between y_a and y_b. Again, the size of this effect depends on the income gradient of welfare weights and the length of both intervals. If society associates the same welfare weight α_p to all workers with incomes below y_b, the welfare loss vanishes and the overall welfare effect is unambiguously positive for any $t_p < 0$.

For a deeper understanding of Lemma 2, note that the reform depicted in Figure 2b as well affects tax revenue due to mechanical effects and behavioral effects that result from labor supply responses at both margins. As in the previous subsection, the choice

[15]The fiscal externality is also positive if the participation tax at the bottom is positive with $t_p > 0$ and the semi-elasticity of participation is increasing along the income distribution. The latter condition is not consistent with the available empirical evidence, however.

of the weighting parameter ϕ_b ensures that the mechanical revenue effects from the tax hike at incomes below y_a and the tax cut at incomes above y_a just offset each other. In general, the reform also induces behavioral effects due to the intensive-margin responses of workers with pre-reform incomes close to y_a and y_b, where the marginal tax is reduced and increased, respectively. Under Assumption 1, however, the marginal tax equals zero both at y_a and y_b before the reform; initially, these intensive-margin responses have hence no effect on tax revenue. Finally, the reform increases the participation tax below y_a and reduces the participation tax above y_a, thereby inducing extensive-margin responses in the entire income range between $\underline{y}$ and y_b. With $t_p < 0$, an agent receives larger transfers in-work than off-work. Hence, the labor market exits of workers with pre-reform incomes below y_a yield a gain in tax revenue, while the labor market entries of workers with incomes in (y_a, y_b) yield a revenue loss. If the former group responds more strongly as measured by the semi-elasticity of participation, the net effect on tax revenue due to extensive-margin responses is positive. As this revenue gain is redistributed to the population, the reform gives rise to a positive fiscal externality that raises welfare.

The previous arguments explain why the introduction of slightly negative marginal taxes can be welfare-increasing. Can we also rationalize a reform that makes negative marginal taxes *more negative*? Assume that, prior to the reform, the marginal tax $T'(y)$ is negative in some interval $(y_a - \varphi, y_a + \varphi)$. Furthermore, assume that the average participation tax for workers with incomes between $\underline{y}$ and y_a is given by $t_p^a < 0$, while the average participation tax for workers with incomes between y_a and y_b is given by $t_p^b < t_p^a$. Then, the fiscal externality due to extensive-margin responses depends on the difference $t_p^b\ \hat{\eta}(y_a, y_b) - t_p^a\ \hat{\eta}(\underline{y}, y_a)$. As long as t_p^b is close to t_p^a, the externality remains positive. Additionally, however, the reform affects revenue due to the intensive-margin responses of workers with incomes close to y_a for whom the marginal tax has become more negative. Specifically, the increase in these agents' incomes reduces their tax payments according to the negative marginal tax $T'(y_a)$. The implied revenue loss induces a *negative* fiscal externality. Hence, the reform gives rise to countervailing fiscal externalities due to labor supply responses at both margins or, put differently, to a trade-off between intensive-margin efficiency and extensive-margin efficiency. Starting from a pre-reform schedule with $T'(y_a)$ close to zero and t_p^a close to t_p^b, the sum of both fiscal externalities will initially be positive and decrease subsequently.

5 Optimal taxation in the discrete model

In the previous section, I have identified conditions under which the introduction of an EITC with negative participation and marginal taxes is welfare-increasing. Even if these conditions hold, however, there might exist other tax reforms that lead to larger welfare gains. Hence, the previous section does not clarify whether the welfare-maximizing tax policy involves an EITC.[16] The following section fills this gap by deriving the optimal income tax for two model versions with a discrete skill set $\Omega = \{\omega_1, \omega_2, \ldots, \omega_n\}$, with $\omega_{j+1} > \omega_j$ for all j. To simplify the notation, I denote by α_j the welfare weight associated to agents with skill type ω_j, by f_j the population share of these agents and by G_j the distribution of fixed costs among them, with corresponding density function g_j.

5.1 The optimal tax problem

I use a mechanism design approach to solve for the optimal income tax. This approach requires, in the first step, to identify the optimal allocation and, in the second step, to find a tax schedule that decentralizes this optimal allocation.

More specifically, the first step requires to solve for the allocation $(c, y) : \Omega \times \Delta \to \mathbb{R}^2$ that maximizes welfare over the set of allocations that are implementable, i.e., both feasible such that overall consumption does not exceed overall output and incentive-compatible such that each type (ω', δ') is assigned a bundle of consumption and income that makes her weakly better off than the bundle assigned to any other type (ω'', δ''). The two-dimensional type set $\Omega \times \Delta$ makes this maximization problem more challenging.

I can simplify the problem, however, by focusing on the set of allocations such that (i) all working agents with skill type ω_j receive the same bundle (c_j, y_j), (ii) all non-working agents receive the same consumption level c_0 and, (iii) an agent with skill type ω_j works if and only if her fixed cost type is below the participation threshold $\delta_j = c_j - h(y_j, \omega_j) - c_0$.[17] By (iii), the participation share in skill group j is given by $G_j(\delta_j)$. Under these restrictions, welfare can be rewritten as

$$\sum_{j=1}^{n} f_j \alpha_j \left\{ \int_{\underline{\delta}}^{\delta_j} [c_j - h(y_j, \omega_j) - \delta] dG_j(\delta) + [1 - G_j(\delta_j)]\, c_0 \right\} , \tag{12}$$

where, for each skill group j, the first term in the bracket gives the utilities of working agents and the second term gives the utilities of non-working agents. The feasibility condition simplifies to

$$\sum_{j=1}^{n} f_j G_j(\delta_j)\, [y_j - c_j] - \sum_{j=1}^{n} f_j\, [1 - G_j(\delta_j)]\, c_0 \geq 0 \, . \tag{13}$$

Moreover, an allocation with properties (i) to (iii) is incentive-compatible if and only if,

[16] While the perturbation approach can also be used to derive an optimal tax formula, this formula does not allow to easily sign the optimal marginal and participation taxes (see Online Appendices D.7 and D.8).

[17] Lemma B.1 in Appendix B shows that these restrictions are without loss of generality: An allocation that violates one of the conditions (i) to (iii) cannot be optimal because it is either not incentive-compatible or not second-best Pareto-efficient. As shown in Subsection 3.2, properties (i) to (iii) also hold in the allocation that follows from the agents' earnings choices given any weakly convex tax function.

for any pair of skill groups j and k in $J \coloneqq \{1, 2, \ldots, n\}$,

$$c_j - h(y_j, \omega_j) \geq c_k - h(y_k, \omega_j) \; . \tag{14}$$

Hence, I only have to take into account the one-dimensional incentive compatibility (IC) constraints among workers with different skill types, which take the same form as in one-dimensional models á la Mirrlees (1971). Incentive compatibility along the fixed cost dimension is already ensured by the threshold condition $\delta_j = c_j - h(y_j, \omega_j) - c_0$.

The second step requires to find a tax schedule that decentralizes the optimal allocation, invoking the taxation principle. In particular, I can infer the optimal signs of marginal taxes and participation taxes based on the intensive-margin and extensive-margin distortions in the optimal allocation. Following Mirrlees (1971), labor supply is said to be *downwards distorted at the intensive margin* for workers with skill type ω_j if the marginal rate of substitution between income and consumption is strictly smaller than the economy's marginal rate of transformation, $h_y\left(y_j, \omega_j\right) < 1$. Labor supply is said to be *upwards distorted at the intensive margin* if $h_y\left(y_j, \omega_j\right)$ is strictly larger than 1. As shown in Subsection 3.2, conditional on working, an agent's optimal income choice is characterized by the first-order condition $1 - T'(y_j) = h_y(y_j, \omega_j)$. Hence, to decentralize an allocation in which labor supply in skill group j is downwards (upwards) distorted at the intensive margin, the marginal tax at income level y_j has to be positive (negative).

In the model studied here, there is a second type of distortions: labor supply is said to be *downwards distorted at the extensive margin* for workers with skill type ω_j if the fixed cost of marginal workers – the participation threshold δ_j – is strictly smaller than $\max_{y>0} y - h(y, \omega_j)$, the maximal surplus that a worker with skill ω_j can generate. As shown in Subsection 3.2, the agents' optimizing behavior gives rise to the threshold condition $\delta_j = y_j - h(y_j, \omega_j) - T^P(y_j)$. Hence, to decentralize an allocation in which labor supply in skill group j is downwards distorted at the extensive margin, the participation tax at the income level that maximizes $y - h(y, \omega_j)$ must be positive. Labor supply is said to be *upwards distorted at the extensive margin* if the participation threshold δ_j is strictly larger than $y_j - h(y_j, \omega_j)$, the surplus a worker with skill ω_j generates under tax T. Hence, to decentralize an allocation in which labor supply in skill group j is upwards distorted at the extensive margin, the participation tax at income level y_j must be negative.[18]

To sum up, I can sign the optimal marginal and participation taxes by determining the labor supply distortions in the allocation that maximizes welfare (12) over the vectors $(c_j)_{j=0}^n$ and $(y_j)_{j=1}^n$, subject to feasibility (13), the threshold conditions $\delta_j = c_j - h(y_j, \omega_j) - c_0$ for all skill groups $j \in J$ and the reduced set of IC constraints (14).

As shown by Jacquet et al. (2013), it is sometimes possible to sign the optimal marginal tax based on a relaxed problem. Specifically, they suggest to solve for the *first-and-half-best allocation* that maximizes welfare (12) subject to feasibility (13) and the threshold

[18]Appendix D.10 discusses the definitions of labor supply distortions at both margins in more detail and illustrates them graphically.

conditions, but ignoring the IC constraints (14) along the skill dimension. They show that, if the tax function implied by the *first-and-half-best allocation* is monotonically increasing, the optimal *second-best* marginal tax is non-negative everywhere. In the following, I refine this method to identify conditions under which the optimal (second-best) marginal taxes and participation taxes are negative at the bottom.

5.2 A stylized model with three skill types

This subsection focuses on a model version that is simplified in three ways. First, there are only three skill types, $\Omega = \{\omega_1, \omega_2, \omega_3\}$. As will become clear, this is the minimal number of skill types such that the optimal marginal tax can be negative given decreasing welfare weights. I refer to the first skill group as the very poor, to the second as the poor and to the third as the rich.

Second, the effort cost function is isoelastic and has the commonly used form

$$h(y, \omega_j) = \frac{\sigma}{1+\sigma} \left(\frac{y}{\omega_j} \right)^{1+\frac{1}{\sigma}} . \tag{15}$$

This implies that the (intensive-margin) elasticity $\varepsilon_T(\omega_j)$ of income $y_T(\omega_j)$ with respect to the retention rate $1 - T'(y)$ is equal to parameter $\sigma > 0$ for each skill group j.

Third, fixed costs are uniformly distributed on the interval $\Delta = \left[0, \bar{\delta}\right]$ in all skill groups.[19] This implies that the skill-specific participation share is given by $G_j(\delta_j) = \delta_j/\bar{\delta}$ and the semi-elasticity of participation is given by $\eta_j = g_j(\delta_j)/G_j(\delta_j) = 1/\delta_j$ for each skill group j. The downward IC constraints along the skill dimension require that $\delta_3 > \delta_2 > \delta_1$. Hence, the semi-elasticities of participation are decreasing over the skill distribution in every implementable allocation, in line with the empirical evidence. The participation elasticity, by contrast, is equal to $1+\sigma$ for each skill group in the laissez-faire allocation.[20]

The following assumption summarizes these properties.

Assumption 2. *The skill set is* $\Omega = \{\omega_1, \omega_2, \omega_3\}$, *the effort cost function is given by* (15) *with* $\sigma > 0$ *and fixed costs are uniformly distributed on* $\left[0, \bar{\delta}\right]$ *in each skill group.*

With respect to the planner's redistributive concerns, I maintain the assumptions that welfare weights are decreasing with $\alpha_1 \geq \alpha_2 \geq \alpha_3$ and have an average value of 1. As a result, the welfare weight α_3 of high-skill workers satisfies

$$\alpha_3 = \frac{1 - f_1\alpha_1 - f_2\alpha_2}{f_3} . \tag{16}$$

Hence, the pair (α_1, α_2) completely characterizes the planner's redistributional concerns and, thereby, determines the welfare ranking among implementable allocations. Addi-

[19] In this stylized model, condition (2) on the type set is equivalent to $\bar{\delta} > \omega_3^{1+\sigma}/(1+\sigma)$.

[20] Under laissez-faire, workers in skill group j choose income $y_j = \omega_j^{1+\sigma}$, and the participation threshold follows as $\delta_j = \omega_j^{1+\sigma}/(1+\sigma)$. The participation elasticity results as $\eta_j \left[y_j - T^P(y_j)\right] = 1 + \sigma$.

tionally, I assume that $\alpha_1 < 2$. Under Assumption 2, this ensures that there is a unique welfare-maximizing allocation that can be identified using the first-order approach.[21]

In the following, I fix the collection of parameters $(\sigma, \bar{\delta}, \Omega, f_1, f_2)$, which jointly determine the set of implementable allocations. I vary only the planner's redistributive concerns as captured by the pair (α_1, α_2) to investigate the conditions under which the optimal allocation involves upwards distortions at both margins. The tax reform analysis in Section 4 suggests that upwards distortions are most likely to be optimal if society considers the poor and the very poor almost equally deserving. Therefore, I start by focusing on the limit case where α_1 and α_2 are identical. Subsequently, I generalize my results to cases where α_1 is close to, but larger than α_2.

No desire for redistribution at the bottom. Consider the limit case where the welfare weights α_1 and α_2 are equal to the same number $\alpha_p \in (1, 2)$. This implies that there is a social desire to redistribute resources from the rich to the poor and the very poor, but no desire for redistribution between the poor and the very poor. Put differently, the planner wants to redistribute resources towards the poor and the very poor in the way that induces as little distortions as possible. The following proposition provides necessary and sufficient conditions under which the optimal redistribution program gives rise to upwards distortions at both margins. The formal proof is provided in Appendix B.

Proposition 1. *Let Assumption 2 be satisfied and let* $\alpha_1 = \alpha_2 = \alpha_p \in (1, 2)$.

(i) Necessary condition: There is a number $\bar{\beta} \in (1, 2)$ *such that optimal labor supply in skill group* 2 *can only be upwards distorted at the intensive margin if* $\alpha_p > \bar{\beta}$.

(ii) Sufficient conditions: There are two functions $\phi_\omega(\alpha_p) > \omega_2$ *and* $\phi_f(\alpha_p) > 0$ *such that optimal labor supply in skill group* 2 *is upwards distorted at both margins if* $\alpha_p > \bar{\beta}$, $\omega_3 > \phi_\omega(\alpha_p)$ *and* $f_3 > \phi_f(\alpha_p)$.

The first part of Proposition 1 provides a necessary condition: Negative marginal taxes can only be optimal if the desire for redistribution from the rich to the poor and the very poor is sufficiently strong (as measured by α_p). The second part of the proposition provides jointly sufficient conditions: Negative marginal taxes are optimal if the desire for redistribution towards the poor is strong enough and, additionally, both the productivity ω_3 and the population share f_3 of high-skilled agents are large enough. The two latter conditions imply that the income share of high-skill workers is sufficiently large. Explicit characterizations of the critical values $\bar{\beta}$, ϕ_ω and ϕ_f are provided in the appendix.[22] Note

[21] Choné & Laroque (2011) explain in detail why, in some cases, the optimal allocation cannot be identified using the first-order approach in random participation models such as mine. In particular, there may exist local welfare minima, or multiple maxima and mimina. Under Assumption 2, the existence of a unique, well-behaved optimum is ensured if the welfare weights of all three skill groups are below 2. I provide a generalization of this condition for the model with n skill types in Hansen (2018). Note that an EITC always turns out to be optimal for a subset of the parameter constellations that satisfy this generalized condition.

[22] Note that the conditions in Proposition 1 (ii) do not define the empty set: The threshold $\phi_f(\alpha_p)$ is in $(0, 1)$ whenever $\omega_3 > \phi_\omega(\alpha_p)$.

that Proposition 1 focuses on the intensive-margin distortions in skill group 2. The reason is that, as I show in the formal proof, labor supply in skill groups 1 and 3 is never upwards distorted at the intensive margin under the assumptions imposed above. If the conditions in part (ii) are satisfied, labor supply in skill group 1 is upwards distorted at the extensive margin and undistorted at the intensive margin.

To understand the mechanism at work, consider in a first step the *first-and-half-best* problem suggested by Jacquet et al. (2013). In the solution to this problem, labor supply is undistorted at the intensive margin because the IC constraints (14) between workers with different skill types are not taken into account. Thus, the *first-and-half-best allocation* $(\tilde{c}, \tilde{y})$ provides no immediate information about marginal taxes. By contrast, labor supply is in general distorted at the extensive margin, and participation taxes differ from zero. Specifically, the *first-and-half-best* participation tax is given by

$$T^P_{fhb}(\tilde{y}_j) = \frac{1 - \alpha_j}{\tilde{\eta}_j} \ , \tag{17}$$

where $\tilde{\eta}_j$ is the semi-elasticity of participation of skill group j in the *first-and-half-best* allocation. Hence, the participation tax for workers in skill group j is negative if and only if the welfare weight α_j for these workers is above the population average of 1. This is independent of whether α_j is lower or higher than the welfare weight of non-working agents. For the intuition behind this result, recall that negative participation taxes give rise to a positive fiscal externality (see Lemma 1).

In the case where α_1 and α_2 are identical and above the average weight of 1, the *first-and-half-best* participation tax is thus negative for workers in skill groups 1 and 2, while it is positive for the workers in skill group 3. More precisely, the participation tax for the two lowest skill groups satisfies

$$T^P_{fhb}(\tilde{y}_j) = -\frac{\alpha_p - 1}{\tilde{\eta}_j} < 0 \text{ for each } j \in \{1, 2\} \ . \tag{18}$$

This is an *inverse elasticity rule*: The participation tax should be more negative for the skill group that responds less elastically at the extensive margin. Under Assumption 2, the poor respond less elastically than the very poor with $0 < \tilde{\eta}_2 < \tilde{\eta}_1$. Hence, the *first-and-half-best* participation tax is decreasing between $\tilde{y}_1$ and $\tilde{y}_2$. This confirms the intuition from the tax reform analysis that negative marginal taxes are beneficial due to a positive fiscal externality (see Lemma 2). Altogether, T^P_{fhb} is non-monotonic: decreasing from $\tilde{y}_1$ to $\tilde{y}_2$, but increasing from $\tilde{y}_2$ to $\tilde{y}_3$. Thus, the results of Jacquet et al. (2013) do not pin down the optimal sign of the *second-best* marginal tax.

To sign the optimal marginal tax, I have to turn to the non-relaxed (second-best) problem that takes into account the IC constraints (14) along the skill dimension. Specifically, I can proceed by checking whether the *first-and-half-best* allocation $(\tilde{c}, \tilde{y})$ satisfies the local IC constraints of the workers in all three skill groups. As the *first-and-half-best* tax is lower (more negative) for workers in skill group 2 than for workers in skill groups 1 and 3, it is clear that workers with skill ω_2 prefer their bundle to the bundles meant for the other workers: Both the downward IC and the upward IC constraint of ω_2 workers are satisfied. For the two remaining IC constraints, this is a priori unclear.

Consider the upward IC constraint of ω_1 workers first. If the difference between the tax levels $T^P_{fhb}(\tilde{y}_1)$ and $T^P_{fhb}(\tilde{y}_2)$ is small, the constraint is satisfied: ω_1 workers prefer their bundle despite receiving somewhat lower transfers. It is violated, by contrast, if the difference is large enough. I find that the tax differential is large enough to violate the upward IC constraint if and only if α_p exceeds a critical value $\bar{\beta} > 1$, i.e., if the welfare weight associated to very poor and poor workers is large enough. For the intuition behind this result, recall again the fiscal externality identified in Section 4. The more the social planner cares for low-income workers as measured by α_p, the larger should their work subsidies be, as can be seen in (18). The larger these work subsidies are, however, the larger is the fiscal externality that results from increasing the subsidy for ω_2 workers relative to the subsidy for ω_1 workers (see Lemma 2). Correspondingly, the optimal difference between $T^P_{fhb}(\tilde{y}_1)$ and $T^P_{fhb}(\tilde{y}_2)$ is increasing in α_p until, eventually, the upward IC of ω_1 workers is violated.

For $\alpha_p > \bar{\beta}$, hence, the social planner cannot follow the inverse elasticity rule (18) unless she slackens the upward IC of ω_1 workers by distorting labor supply of ω_2 workers upwards at the intensive margin. Hence, she faces a trade-off between two aspects of efficiency: She can only reduce extensive-margin distortions if she increases intensive-margin distortions and vice versa. As argued in Section 4, the introduction of a slight upward distortion at the intensive margin only leads to a negligible efficiency loss. Hence, the extensive-margin benefits of introducing small upward distortions are initially larger than the intensive-margin costs whenever $\alpha_p > \bar{\beta}$. For $\alpha_p \leq \bar{\beta}$, by contrast, the social planner can stick to the inverse elasticity rule (18) without violating an upward IC constraint. In this case, neither upwards distortions at the intensive margin nor negative marginal taxes bring an benefit. This explains why negative marginal taxes can only be optimal if $\alpha_p > \bar{\beta}$.

By part (ii) of Proposition 1, however, the previous condition is not sufficient: For negative marginal taxes to be ensured, the income share of the rich must be large enough as measured by the population share f_3 and the productivity ω_3. To understand this additional requirement, consider the downward IC constraint of workers with skill ω_3. As the *first-and-half-best* tax on the rich is positive while the tax on the poor in skill group 2 is negative, this IC constraint may be violated. If the welfare weight α_3 of the rich is above a critical value $\beta_D(\alpha_p)$, however, the difference between $T^P_{fhb}(\tilde{y}_3)$ and $T^P_{fhb}(\tilde{y}_2)$ is small enough so that the downward IC of ω_3 workers is satisfied.[23] Proposition 1 rephrases the condition $\alpha_3 > \beta_D(\alpha_p)$ as a condition on the productivity level ω_3 and the population share f_3. For an intuitive understanding, note that the taxes on ω_3 workers have to be large enough to finance the work subsidies to the lower-skilled workers. The larger the

[23] In the Mirrlees (1971) model without extensive-margin responses, by contrast, the downward IC of ω_3 workers is binding in the optimal allocation whenever $\alpha_3 < \alpha_p$.

share f_3 of high-skilled worker is, the smaller is the required tax payment $T(y_3)$ for each single high-skill worker. The larger the productivity level ω_3 is, moreover, the higher are both the gross income y_3 and the implied net income $y_3 - T(y_3)$ relative to the net income of ω_2 workers, $y_2 - T(y_2)$. Hence, the larger the share and the productivity of rich workers, the less increasing tax function T must be between y_2 and y_3 to finance the work subsidies to low-skill workers. Correspondingly, the induced intensive-margin costs become smaller until, eventually, they are dominated by the extensive-margin benefits of an EITC with negative marginal taxes for ω_2 workers.

If instead only the necessary condition holds but the conditions on f_3 and ω_3 fail, the *first-and-half-best* allocation violates both the upward IC of ω_1 workers and the downward IC of ω_3 workers. To slacken the former constraint, labor supply of ω_2 workers would have to be distorted *upwards* at the intensive margin. To slacken the latter constraint, in contrast, labor supply of ω_2 workers would have to be distorted *downwards*. But then, the planner would have to deviate even further from the inverse elasticity rule (18). Hence, it remains unclear whether optimal labor supply in skill group 2 is upwards or downwards distorted at the intensive margin.[24]

Weak desire for redistribution at the bottom. In the next step, I generalize the previous results to cases in which the welfare weight of the very poor is strictly higher than the welfare weight on the poor, $\alpha_1 > \alpha_2 > 1$. Thereby, I consider cases in which the social planner has a desire for redistribution from the poor to the very poor. The next proposition clarifies that upwards distortions at the intensive margin remain optimal as long as this desire for redistribution at the bottom is weak enough. The formal proof is relegated to Appendix B.

Proposition 2. *Let Assumption 2 be satisfied and let $\bar{\beta}$, ϕ_ω and ϕ_f be the thresholds defined in Proposition 1.*

(i) Necessary condition: For any $\alpha_1 \in (\bar{\beta}, 2)$, there is a threshold $\beta_U(\alpha_1) \in \left(\bar{\beta}, \alpha_1\right)$ such that optimal labor supply in skill group 2 can only be upwards distorted at the intensive margin if $\alpha_2 > \beta_U(\alpha_1)$.

(ii) Sufficient conditions: There is a function $\hat{\phi}_\omega(\alpha_1, \alpha_2) \geq \phi_\omega(\alpha_2)$ such that optimal labor supply is upwards distorted at both margins in skill group 2 if $\alpha_2 > \beta_U(\alpha_1)$, $\omega_3 > \hat{\phi}_\omega(\alpha_1, \alpha_2)$ and $f_3 > \phi_f(\alpha_2)\left(1 + f_1 \frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right)$.

Part (i) of Proposition 2 shows that the optimal allocation can only involve upwards distortions at the intensive margin if the welfare weights α_1 and α_2 are, first, larger than $\bar{\beta}$ – the critical value from Proposition 1 above – and, second, close enough to each other. Specifically, I find that the optimal marginal tax can be negative even if α_2 is strictly smaller than α_1 but larger than the threshold $\beta_U(\alpha_1)$.

[24] In this case, the optimal sign of the marginal tax for ω_2 workers depends on the complete set of parameters $(f_1, f_2, \bar{\delta}, \sigma, \omega_1, \omega_2, \omega_3, \alpha_p)$.

Part (ii) of Proposition 2 generalizes the sufficient conditions for the optimality of negative marginal taxes to the case where $\alpha_1 > \alpha_2$. On top of the necessary condition from part (i), the productivity level ω_3 and the population share f_3 of high-skill workers must be large enough again. The condition on the high-skill productivity ω_3 is the same as in Proposition 1 for $\alpha_1 = \alpha_2$, but somewhat tighter if $\alpha_1 > \alpha_2$. The condition on the high-skill population share f_3 uses the same threshold function ϕ_f, buts also gets harder to satisfy if $\alpha_1 > \alpha_2$. Intuitively, a higher value of α_1 implies that the planner prefers a larger transfer to the very poor than with $\alpha_1 = \alpha_2$ and, hence, needs to collect somewhat higher taxes from the high-skilled workers. As a result, the optimal tax schedule gets steeper between y_2 and y_3, which tends to increase the labor supply distortions at the intensive margin. The basic insight remains unchanged, however: The larger f_3 and ω_3 are, the easier it is to finance work subsidies for lower-skilled workers and the more likely it is than the extensive-margin benefits of an EITC dominate.

5.3 A general model with many skill types

In this subsection, I extend the previous results to a model version that is more general in two ways. First, I allow for a finite but arbitrarily large number $n \geq 3$ of skill groups. I assume that the log difference between adjacent skill types is constant, i.e., that the ratio ω_{j+1}/ω_j is equal to the same number $a > 1$ for all $j \in \{1, \ldots, n-1\}$.

Second, I do not specify the functional forms of the utility function (1) and the type distribution, and allow the fixed cost distributions to vary across skill groups. I only impose three assumptions on the implied labor supply elasticities. I start with a regularity condition on the labor supply responses at the intensive margin.

Assumption 3. *In any implementable allocation, the elasticity of income in each skill group $j \in J$ with respect to*

(i) the retention rate $1 - T'(y)$ is bounded from above by some number $\nu_1 \in (0, \infty)$;

(ii) the skill level ω is bounded from below by some number $\nu_2 \in (0, \infty)$.

The first part of Assumption 3 rules out cases in which labor supply explodes in response to a small reduction in the retention rate. The second part implies that, given a constant marginal tax rate, agents with higher skill types choose higher incomes. Hence, it ensures that the skill type is a meaningful measure of productivity.

Next, I impose an assumption on the relative strength of labor supply responses at the extensive margin, which ensures consistency with the empirical evidence.

Assumption 4. *In any implementable allocation, the semi-elasticity of participation η_j is strictly decreasing in ω_j and weakly decreasing in $c_j - c_0$ for any $j \in J$.*

Assumption 4 is satisfied if participation elasticities are decreasing along the income distribution, in line with empirical findings (see, e.g., Juhn et al. 1991, 2002 for the

US and Meghir & Phillips 2010 for the UK).[25] Even if the participation elasticity were constant along the skill distribution, however, the corresponding semi-elasticity would still be strictly decreasing.[26]

Finally, I impose an assumption that rules out extreme fluctuations in the skill gradient of the semi-elasticity of participation. For this purpose, consider two adjacent skill groups j and $j+1$ with participation semi-elasticities η_j and η_{j+1}, respectively. By Assumption 4, the ratio η_j/η_{j+1} is larger than 1 given any status-quo tax schedule T. Consider now a tax reform that reduces the participation taxes for the workers in both skill groups by some amount $\lambda > 0$. This is just the type of reform considered in Subsection 4.1 and illustrated in Figure 2a. In general, the reform can lead to variations in both semi-elasticities and in their ratio. I measure the latter variations by the semi-elasticity of the ratio η_j/η_{j+1} with respect to λ, i.e., with respect to the net-of-tax labor income $y - T^P(y)$. The following assumption puts bounds on these variations, i.e., on the effect of the reform on the relative participation responses in skill groups j and $j+1$.

Assumption 5. *In any implementable allocation, the semi-elasticity of the relative participation responses η_j/η_{j+1} with respect to the net-of-tax labor income $y - T^P(y)$ is between $\eta_{j+1} - \eta_j < 0$ and $\eta_j - \eta_{j+1} > 0$.*

Under Assumption 5, a uniform reduction of the participation tax may increase or decrease the ratio of the semi-elasticities η_j and η_{j+1}, but needs to have sufficiently small effects in absolute terms. While the semi-elasticity of the ratio η_j/η_{j+1} is in principle an observable object, it has never been estimated in the empirical literature to my knowledge. Most relatedly, Juhn et al. (1991, 2002) find that relative participation elasticities for different percentiles of the wage distribution have hardly changed between the 70s and the late 80s. Besides, a back-of-the-envelope calculation suggests that Assumption 5 is reasonably weak.[27]

Results I proceed by extending my previous results to the general model with an arbitrary number of skill groups. I assume that the optimal allocation can be identified using the first-order approach. It turns out that, under Assumptions 3 to 5, the results from the stylized three-type model remain qualitatively unchanged.

For this purpose, recall Proposition 2 for the stylized model. It provides necessary and sufficient conditions for the optimal allocation to involve upwards distortions at both

[25] According to Bastani et al. (2020) for Sweden and Miller et al. (2018) for the US, participation elasticities are also decreasing within the subset of low-skilled workers. In contrast, Bargain et al. (2014) estimate participation elasticities to be flat at the bottom and only decreasing at higher income levels.

[26] The participation elasticity π_T and the semi-elasticity of participation η_T are related by $\eta_T(\omega) = \pi_T(\omega)/\left[y_T(\omega) - T^P(y_T(\omega))\right]$. Hence, with π_T constant in ω, η_T is strictly decreasing as long as the net labor income $y_T(\omega) - T^P(y_T(\omega))$ is increasing in ω, i.e., the marginal tax T' is below 1.

[27] If the participation elasticities for low-skill workers with earnings $y_j = \$1,000$ and $y_{j+1} = \$1,500$ are both given by 0.5 as assumed by Saez (2002) and the US income tax is approximated as in Section 6, the ratio η_j/η_{j+1} equals 1.5. Hence, the semi-elasticity of participation in the first group is 50% higher than in the second group. Assumption 5 requires that, if the participation tax is reduced by \$500 for both groups, the ratio η_j/η_{j+1} remains between 32.3% and 67.7%.

margins in skill group 2: The welfare weights of the low-skill ω_1 and ω_2 workers need to be large enough – above some threshold $\bar{\beta} > 1$ – and close enough to each other such that $\alpha_2 > \beta_U(\alpha_1)$. Moreover, the welfare weight of the high-skill ω_3 workers needs to be large enough to satisfy $\alpha_3 > \beta_D(\alpha_2)$. The latter condition is satisfied if both the productivity level ω_3 and the share f_3 of high-skill workers are large enough. The following proposition generalizes the conditions on the welfare weights to a model with n skill groups. The formal proof is provided in Online Appendix C.

Proposition 3. *Let Assumptions 3, 4 and 5 be satisfied. Fix some skill group* $k \in \{2, \ldots, n-1\}$. *If* $a = \omega_{j+1}/\omega_j$ *is sufficiently small, there are two sets of functions* $(\beta_{Uj})_{j=1}^{n-1}$ *and* $(\beta_{Dj})_{j=1}^{n-1}$ *such that optimal labor supply is upwards distorted at both margins in the skill groups* $\{2, \ldots, k\}$ *if the welfare weights satisfy*

(i) $\alpha_{j+1} \geq \beta_{Uj}(\alpha_j)$ *for all* $j \in \{1, \ldots, k-1\}$ *with at least one strict inequality, and*

(ii) $\alpha_{j+1} \geq \beta_{Dj}(\alpha_j)$ *for all* $j \in \{k, \ldots, n-1\}$.

If the sequence of welfare weights $(\alpha_1, \alpha_2, \ldots, \alpha_n)$ satisfy the conditions in Proposition 3, labor supply is upwards distorted at both margins in the skill groups 2 to k. Labor supply in skill group 1, moreover, is upwards distorted at the extensive margin and undistorted at the intensive margin. Part (i) corresponds to the necessary condition in Proposition 2: It puts restrictions on the welfare weights of those workers for whom optimal labor supply is upwards distorted. Part (ii) corresponds to the additional sufficient conditions on the weights of the higher-skill workers: It puts restrictions on the welfare weights of the remaining skill groups $k+1$ to n. The proposition is valid whenever the relative distance $a = \omega_{j+1}/\omega_j$ between adjacent skill types is small enough, i.e., the skill set Ω is sufficiently dense. In particular, this restriction rules out degenerate cases in which optimal labor supply in some skill groups is undistorted at the intensive margin for all welfare weights.

The economic implications of Proposition 3 depend on the properties of the threshold functions β_{Uj} and β_{Dj}. To illustrate these properties, Figure 3 on the next page depicts all possible combinations of the welfare weights for an arbitrary pair of adjacent skill groups j and $j+1$. It focuses on the combinations with decreasing welfare weights $\alpha_{j+1} \leq \alpha_j$, i.e., in the shaded area below the 45° line. The purpose of this figure is to clarify which of these combinations satisfy the conditions (i) $\alpha_{j+1} \geq \beta_{Uj}(\alpha_j)$ and (ii) $\alpha_{j+1} \geq \beta_{Dj}(\alpha_j)$. Under Assumptions 3 to 5, the functions β_{Uj} and β_{Dj} always have the shape depicted in Figure 3 (see also Lemma C.1 in the appendix).[28]

First, function β_{Uj} is located below the 45° line for all values of α_j above a critical value $\bar{\beta}_j > 1$. Hence, the condition $\alpha_{j+1} \geq \beta_{Uj}(\alpha_j)$ is satisfied for all combinations in the red-shaded area, where both welfare weights α_j and α_{j+1} are large enough (above $\bar{\beta}_j$) and close enough to each other. As the critical value $\bar{\beta}_j$ is larger than the average weight of

[28]Note that the functions β_{Uj} and β_{Dj} are only defined if labor supply in skill groups j and $j+1$ responds at the extensive margin (i.e., if $\eta_j, \eta_{j+1} > 0$).

Figure 3: Local IC constraints in the first-and-half-best allocation.

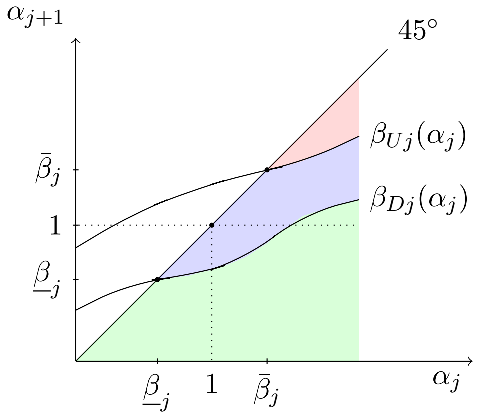

Notes: Figure 3 illustrates the properties of the threshold functions β_{Uj} and β_{Dj} under Assumptions 3 to 5. The red-shaded area contains the combinations of α_j and α_{j+1} such that $\alpha_{j+1} > \beta_{Uj}(\alpha_j)$. The blue-shaded area contains the combinations such that α_{j+1} is between $\beta_{Dj}(\alpha_j)$ and $\beta_{Uj}(\alpha_j)$. The green-shaded area contains the combinations such that $\alpha_{j+1} < \beta_{Dj}(\alpha_j)$. The optimal allocation involves upwards distortions at both margins if the combination (α_j, α_{j+1}) is located in the red-shaded area for each j below some $k \in \{2, \ldots, n-1\}$, and in the blue-shaded or red-shaded area for each $j \geq k$.

1, this case can only appear for pairs of low-skill workers (i.e., at the bottom of the skill distribution). Second, function β_{Dj} is located below β_{Uj} for all values of α_j, and below the 45° line for all values of α_j above another critical value $\underline{\beta}_j < 1$. Hence, the condition $\alpha_{j+1} \geq \beta_{Dj}(\alpha_j)$ is satisfied for all combinations in the red-shaded area and in the blue-shaded area, where both welfare weights are neither too low nor too far apart. As the critical value $\underline{\beta}_j$ is below the average weight of 1, this case can even appear for pairs of high-skill workers (i.e., at the top of the skill distribution). The condition $\alpha_{j+1} < \beta_{Dj}(\alpha_j)$ is violated, finally, for combinations in the green-shaded area where both welfare weights are either pretty low or pretty distant.

Summing up, Proposition 3 implies that the optimal allocation involves upwards distortions at both margins for the lowest k skill groups if, first, the welfare weights of all these skill groups are high enough and close enough to each other, i.e., if the skill gradient of the welfare weights at the bottom is small enough and, second, if the welfare weights of the remaining skill groups are close enough to the population average, i.e., the skill gradient at the top is moderate.

For the stylized model, the sufficient condition on the welfare weight α_3 of high-skill workers can be rewritten in terms of the type distribution: The condition $\alpha_3 \geq \beta_D(\alpha_2)$ is satisfied so that an EITC is optimal if both the productivity level ω_3 and the population share f_3 of high-skilled workers are large enough (see Propositions 1 and 2). The following proposition extends this insight to the general model. The formal proof is provided in Online Appendix C.

Proposition 4. *Let Assumptions 4, 5 and 3 be satisfied. Fix some skill group* $k \in \{2, \ldots, n-1\}$*. If* $a = \omega_{j+1}/\omega_j$ *is sufficiently small, there are two numbers* $\omega_m \in \Omega : \omega_m \geq \omega_{k+1}$ *and* $\bar{z} \in (0, 1)$ *such that, if*

(a) $\omega_n \geq \omega_m$ *and*

(b) $\sum_{j=m}^{n} f_j > \bar{z}$,

there exists a strictly decreasing sequence $(\alpha_1, \ldots, \alpha_n)$ *of welfare weights for which optimal labor supply in skill groups* $\{2, \ldots, k\}$ *is upwards distorted at both margins.*

Proposition 4 clarifies that, in the general model as well, the sufficient conditions on the welfare weights can be satisfied if the share of highly skilled agents is large enough. More specifically, condition (a) requires the productivity of the highest-skilled workers to be sufficiently large, relative to the productivity of the EITC recipients with skills below ω_k. In particular, the productivity of the highest-skilled workers must be above some threshold ω_m. Condition (b) requires the population share of the highly productive workers with skill types ω_m to ω_n to be large enough. Both conditions are natural generalizations of the sufficient conditions on the productivity ω_3 and the share f_3 of high-skill workers in the stylized model. Again, the proposition is valid whenever the relative distance $a = \omega_{j+1}/\omega_j$ between adjacent skill types is small enough.

An important difference is, however, that Proposition 4 effectively puts an upper bound on the number k of skill groups with upward distortions at both margins, i.e., on the income range with negative marginal taxes. To see this, note that the proposition provides sufficient conditions for the optimality of different allocations: some with upward distortions for only a few skill groups (k small, close to 2) and others with upward distortions for the majority of skill groups (k large, close to n). In the former case, an EITC with negative marginal taxes is concentrated on the working poor. In the latter case, an EITC even benefits workers in the (upper) middle class. As one would expect, the conditions in Proposition 4 are easier to satisfy in the former case than in the latter one. Specifically, the larger the productivity ω_k of the highest-skilled EITC recipient, the larger is the threshold skill type ω_m referred to in Proposition 4. For empirically reasonable calibrations such as those in the next section, both conditions (a) and (b) are always satisfied for number k being as small as 2, but violated for k being as large as $n - 1$.

The intuition for these sufficient conditions, and for the implied bound on the size of the EITC, is the same as in the stylized model with three skill types. An EITC with work subsidies and negative marginal taxes for the workers in the lowest k skill groups has to be financed by tax payments from the workers in higher skill groups. The more workers with high skill types and high incomes there are, the more easily the social planner can raise the required tax revenue. In particular, a larger share of high-skilled workers allows the planner to rely on a relatively flat tax schedule that only induces small distortions at the intensive margin. In this case, as discussed above, the extensive-margin benefits of an EITC with negative marginal taxes dominate the intensive-margin costs.

Finally, I explain the relevance of Assumptions 3 to 5. For this purpose, recall that the introduction of negative marginal taxes leads to a positive fiscal externality, which is proportional to $-t_P \left[\hat{\eta}(\underline{y}, y_a) - \hat{\eta}(y_a, y_b)\right]$ by Lemma 2. The term t_P is the participation tax for low-income workers; the term in brackets is a measure of the skill gradient of the semi-elasticity of participation among low-income workers. Assumption 4 implies that the semi-elasticity of participation is decreasing, i.e., the term in brackets is positive. Hence, it ensures that the fiscal externality is positive whenever the participation tax t_P is negative. Assumption 5 implies that changes in the participation tax have only limited effects on the skill gradient of the semi-elasticities, i.e., on the term in brackets. In particular, it ensures that the fiscal externality becomes larger when the participation tax t_P gets more negative. Together, Assumptions 4 and 5 ensure that the introduction of an EITC with negative marginal taxes leads to a positive fiscal externality and to welfare gains that are larger, the higher the welfare weights of low-income workers are. Finally, Assumption 3 implies that the introduction of a small negative marginal tax for low-income workers does not lead to huge changes in their income choices. This ensures that, initially, the intensive-margin costs of an EITC are negligible and, in particular, smaller than the extensive-margin benefits.

6 Numerical simulations

The theoretical analysis above has shown that an EITC can be optimal, but has not provided information about the optimal size of an EITC (i.e., eligibility threshold, length of phase-in range, optimal levels of marginal and participation taxes). To answer these quantitative questions, the following section performs numerical simulations for a model version that is calibrated to the subgroup of childless singles in the US. The focus on singles ensures consistency with the theoretical model studied above, which does not account for joint labor supply decisions within families. Moreover, childless singles are in the spotlight of a recent policy debate: prominent politicians from both political camps have proposed to strongly expand the EITC for this group. At the same time, the previous literature has not provided any support for negative marginal taxes in this group. In Online Appendix D.5, I additionally provide a calibration for single parents, who benefit from a more generous EITC under the current US tax system.

6.1 Calibration

I calibrate the model by imposing assumptions on the labor supply elasticities at both margins, the two-dimensional distribution of skill types and fixed cost types and the redistributive preferences of the social planner. To specify the labor supply elasticities, I mainly follow the survey by Saez et al. (2012) and the meta-study by Chetty, Guren, Manoli & Weber (2013), who consider estimates from fifteen studies providing quasi-experimental estimates of extensive-margin elasticities. Additionally, I consider the studies by Chetty, Friedman & Saez (2013) on the elasticities on EITC recipients and by Bargain et al. (2014) on the elasticities of childless singles (see Online Appendix D.9 for a discussion of empirical estimates). For completeness, note that I have assumed away income effects in labor supply by the quasi-linearity of the utility function (1). Saez (2002) and Jacquet et al. (2013) impose the same simplifying assumption in their calibrated models.[29]

First, with respect to the intensive-margin elasticity of income with respect to the retention rate, the best available estimates are in the range between 0.12 and 0.4 according to Saez et al. (2012). The preferred estimate by Chetty, Guren, Manoli & Weber (2013) is given by 0.33, while Bargain et al. (2014) estimate an elasticity of 0.18. I calibrate the model by assuming that the effort cost function h has the same functional form (15) as in the stylized model. Thus, the intensive-margin elasticity ε equals the parameter σ, which I set to 0.3 for the benchmark calibration. In the sensitivity analysis, I consider alternative values of 0.1 and 0.5, respectively.

Second, with respect to the extensive-margin elasticity of participation, the preferred estimate of Chetty, Guren, Manoli & Weber (2013) for the entire population is given by 0.25. Bargain et al. (2014) estimate an average elasticity of 0.28 for childless singles in the US. Moreover, a number of studies consistently find that participation elasticities are decreasing along the skill distribution (e.g., Juhn et al. 1991, 2002, Meghir & Phillips 2010, Bargain et al. 2014 and Bastani et al. 2020). In my model, the participation elasticities depend on the distributions of fixed costs in each skill skill group. Unfortunately, there is no empirical evidence on these fixed cost distributions to the best of my knowledge. I therefore follow Jacquet et al. (2013) by assuming a logistic distribution of the form

$$G_j(\delta) = \frac{\exp(-\psi_j + \rho_j \delta)}{1 + \exp(-\psi_j + \rho_j \delta)} \ . \tag{19}$$

This assumption ensures that the participation share is between 0 and 1 and that labor supply responds at the extensive margin in each skill group for any admissible tax function T. Moreover, it allows me to set the skill-specific parameters ψ_j and ρ_j to match empirically plausible values of the participation elasticity π_j and the employment share L_j in each skill group.

Following Jacquet et al. (2013) once more, I assume that participation elasticities decrease gradually along the skill distribution according to

$$\pi_j = \pi_1 + (\pi_n - \pi_1) \left(\frac{\omega_j - \omega_1}{\omega_n - \omega_1} \right)^{1/3} \ , \tag{20}$$

where π_1 and π_n are the participation elasticities in the lowest and the highest skill group, respectively. For my benchmark calibration, I assume that participation elasticities decrease from $\pi_1 = 0.4$ to $\pi_n = 0.18$ in the group of childless singles. These parameter choices ensure that the average participation elasticity $\bar{\pi} = \sum_{j=1}^{n} f_j \pi_j$ equals the preferred estimate of 0.25 in Chetty, Guren, Manoli & Weber (2013). In the sensitivity analysis, I consider cases with an average participation elasticity of 0.12 and 0.43, respectively.[30]

[29] The existing evidence on income effects is scarce and inconclusive. For example, Holtz-Eakin et al. (1993) and Imbens et al. (2001) find only small and sometimes insignificant income effects.

[30] The case $\bar{\pi} = 0.12$ corresponds to the point estimate for the preferred model in Kleven (2020*a*), see Tables 3 and 6. The case $\bar{\pi} = 0.43$ is in line with the calibration in Jacquet et al. (2013).

Similarly, I target skill-specific employment shares given by

$$L_j = L_1 + (L_n - L_1) \left(\frac{\omega_j - \omega_1}{\omega_n - \omega_1} \right)^{1/3} , \tag{21}$$

where L_1 and L_n are the participation elasticities in the lowest and the highest skill group, respectively. I assume that the participation share is increasing from 0.7 in the lowest skill group to 0.85 in the highest skill group. This implies an average participation share close to 0.8, in line with CPS data.

Third, I calibrate the unconditional skill distribution to match the observed income distribution in the US economy. Specifically, I estimate the latter distribution based on income data for childless singles at ages 25 to 60 in the March 2016 CPS. I restrict the sample to respondents that are neither living with an unmarried spouse nor with any family members in the same household. I then calculate an agent's earned income as the sum of, first, wage and salary income and, second, business and self-employed income.[31] Using the OECD tax database, I approximate the US income tax in 2015 for both population groups by a linear tax function with a marginal tax rate of 29.3% (OECD 2017, p. 54).[32] Based on this approximation, I can use the first-order condition of the individual optimization program to back out the skill types of all CPS respondents with strictly positive earned incomes.

Fourth, I consider a discrete skill set with $n = 96$ skill types, where the relative distance between each pair of adjacent skill types is equal to $\omega_{j+1}/\omega_j = 1.05$.[33] Compared to most previous papers, this represents a relatively fine skill set. For example, Saez (2002) uses a model with only 17 income groups. The workers with the lowest and highest skill types receive annual incomes of \$500 and \$206,942, respectively. In the March 2016 CPS, 98.8% of the employed childless singles have incomes in this range. To obtain a smooth distribution, I estimate the share of workers in each skill group j with a kernel density approximation of the distribution of computed skill types. This procedure gives the conditional skill distribution among employed workers under the 2015 US tax regime. In the last step, I use the skill-specific employment rates specified by (21) to compute the unconditional skill distribution (including workers and non-workers).

Finally, I need to calibrate the redistributive preferences of the social planner. For my benchmark analysis, I consider a sequence α^A of exogenous welfare weights that depend only on the agents' skill types and have an average value of 1. Specifically, sequence α^A assigns a welfare weight of 1.04 to all workers with annual incomes below the 2015 poverty line for childless singles in the US, which was given by \$11,770. This includes the agents

[31] In particular, I compute each worker's average income per week in employment according to the CPS data. To calculate an agent's skill, I then multiply the weekly income by 52 to get individually optimal incomes conditional on working the entire year.

[32] This approximation accounts for federal and (average) state income taxes as well as the employee share of payroll taxes (OECD 2017, see also http://www.oecd.org/tax/tax-policy/tax-database.htm).

[33] In Online Appendix D.3, I additionally provide simulation results for models with a larger and a smaller number of skill groups (n equal to 192 and 48, respectively).

in the lowest 49 skill groups or, put differently, the bottom 5.6% of the skill distribution. The welfare weights for higher-skilled workers with incomes above the poverty line are gradually decreasing. In particular, sequence α^A is constructed to satisfy the sufficient conditions for an EITC with negative marginal taxes for incomes below $\$11,770$ (see Proposition 3). Hence, I already know that the optimal income tax is given by an EITC. The numerical simulation allows to assess the quantitative properties of this EITC, e.g., the optimal levels of marginal taxes and the maximum tax credit. In the sensitivity analysis, I also consider alternative welfare weights that violate the sufficient conditions for an EITC provided in Proposition 3. Further details on the construction of these sequences of welfare weights and graphical illustrations are provided in Online Appendix D.2.

6.2 Benchmark results

Figure 4 illustrates the results of my benchmark simulations. In particular, the solid blue line depicts the optimal participation tax schedule $T_A^P(y) = T_A(y) - T_A(0)$ for annual incomes below $\$60,000$. To make the discreteness of the skill set transparent, I mark the optimal income-tax bundle for the workers in each skill group by a dot. As a reference, the red dotted line in Figure 4 depicts the actual EITC for childless singles in the US. In 2015, childless singles were eligible for the EITC if their earned income was below $\$14,820$. The marginal income tax for this group was -7.65% for incomes below \$6.580 (phase-in range) and $+7.65\%$ for incomes between \$8.240 and \$14.820 (phase-out range). Workers with incomes between \$6.580 and \$8.240 received the maximum tax credit of \$503.[34]

Figure 4: Optimal participation taxes for benchmark case.

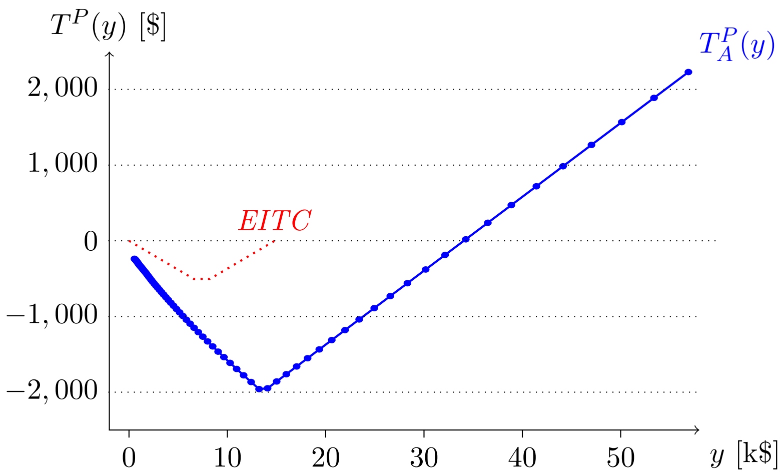

Notes: Figure 4 illustrates the simulation results by depicting the optimal participation tax schedule T_A^P for welfare weights α_A as a function of gross labor income y. Each dot represents the tuple $(y_j, T^P(y_j))$ for the workers in one skill group. For comparison, the red dotted line depicts the actual 2015 EITC.

Recall that weight sequence α^A was constructed to satisfy the sufficient conditions in Proposition 3. Thus, I already know that an EITC with negative marginal taxes at all

[34] In the phase-in range, the EITC exactly offsets the employee share of social security contributions.

income levels below the poverty line at \$11,770 is optimal. The simulation shows that optimal marginal taxes are even negative for incomes up to \$13,228, and participation taxes are negative for incomes up to \$32,144 (see blue line in Figure 4). Put differently, labor supply is upwards distorted at the intensive margin in the lowest 51 skill groups and upwards distorted at the extensive margin in the lowest 65 skill groups. This implies that 28.4% of all childless singles benefit from negative participation taxes as part of an EITC, and 21.9% of these EITC recipients face negative marginal taxes in the phase-in range. The share of non-working agents is reduced substantially to 12.6% (compared to 20% under the current US tax).

More specifically, the maximum tax credit is given by \$1,959 at income level \$13,288. For comparison, the optimal transfer to non-working agents is given by \$1,517. The ratio $T_A^P(y)/y$ of optimal participation taxes to gross labor incomes, which is sometimes referred to as the participation tax rate, is around -40% for very low incomes.[35] The ratio subsequently diminishes to levels around -15% at the end of the phase-in range. The average marginal tax in the phase-in range is given by $-13,6\%$, and the average marginal tax in the phase-out range is given by 9.4%.[36]

Hence, my numerical simulation shows that the effects of the mechanism studied in this paper - the tradeoff between intensive-margin efficiency and extensive-margin efficiency - are not only qualitatively, but also quantitatively important. When there are no (or only weak) concerns for redistribution among the poor, the optimal EITC may cover a much larger income range and feature a much larger maximum tax credit than the current EITC for childless workers in the US. Moreover, negative marginal taxes and participation taxes can be more than twice as large (in absolute terms) in the optimal scheme than in the current US scheme. Interestingly, the simulated tax schedules closely resemble those implied by recent proposals to strongly expand the EITC for childless workers in the US (for example, see Executive Office 2014 and House Budget Committee 2014).

6.3 Sensitivity analysis

To investigate the sensitivity of my results, I perform additional simulations in which I vary the key parameters of my calibrated model: intensive-margin elasticities, participation elasticities, and welfare weights. In this section, I focus on five alternative scenarios that might be expected to make an EITC less attractive and, potentially, suboptimal. Online Appendix D.3 provides simulations for additional specifications, a detailed discussion of all parameter choices and graphical illustrations of the corresponding optimal tax schedules.[37]

Figure 5 shows my simulation results for the first three scenarios. For comparison,

[35] The large negative participation taxes at the very bottom suggest that the optimal income tax may fall discontinuously at zero, in line with the results of Jacquet et al. (2013).

[36] The average marginal tax between y_k and y_j is computed as $[T(y_k) - T(y_j)] \,/\, (y_k - y_j)$. Alternatively, one can compute the implicit marginal tax at income y_j, which is given by $1 - h_y(y_j, \omega_j)$. The implicit marginal taxes are between 0 and 8% in the phase-in range and equal to zero in the phase-out range.

[37] In Online Appendix D.3, I also consider alternative discretizations of the skill set.

Figure 5: Sensitivity analysis, part 1.

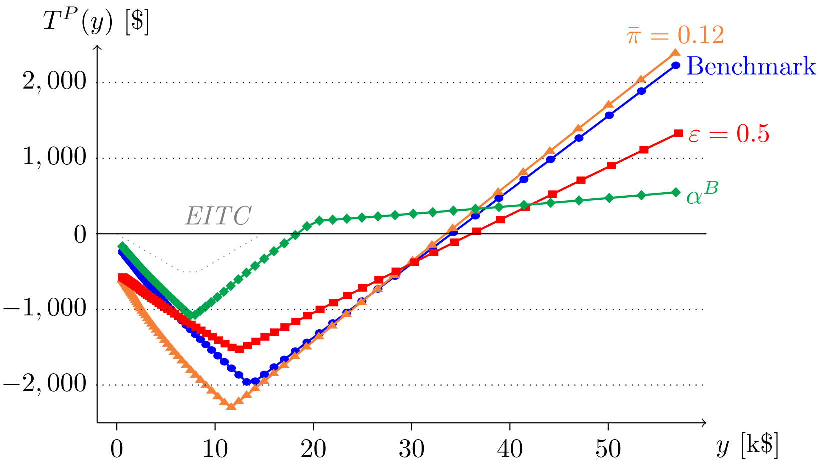

Notes: Figure 5 plots the optimal participation taxes for the benchmark case (blue line with dots) and alternative specifications with a larger intensive-margin elasticity $\varepsilon = 0.5$ (red line with squares), lower participation elasticities such that $\bar{\pi} = 0.12$ (orange line with triangles), or welfare weights α^B that drop sharply for workers with incomes above \$11,000 under the current US tax (green line with diamonds). For each specification, the marks represent the tuples $(y_j, T^P(y_j))$ for the workers in the different skill group. For comparison, the gray dotted line depicts the actual 2015 EITC.

the blue line with dots again shows the optimal tax in the benchmark simulation with an intensive-margin elasticity of 0.3, an average participation elasticity of 0.25, and welfare weights α^A that are flat for workers with incomes below \$11,000 and gradually decreasing for higher-skilled agents.

The first two specifications imply that the ratio of extensive-margin responses to intensive-margin responses is smaller than in the benchmark. Specifically, the red line with squares depicts the optimal participation tax for a specification with a larger intensive-margin elasticity of $\varepsilon = 0.5$. The orange line with triangles depicts the optimal tax for a case with a smaller average participation elasticity of $\bar{\pi} = 0.12$.[38] Compared to the benchmark, the optimal phase-in rate and the maximum tax credit are somewhat smaller in the former case and a bit larger in the latter case. In both cases, the phase-in range is slightly shorter than in the benchmark scenario. More importantly, however, my simulations for both scenarios confirm the qualitative insights from my benchmark simulation: The optimal income tax involves an EITC with negative marginal and participation taxes that is much larger than the 2015 EITC for childless singles.

The green line with diamonds shows the optimal tax for a third scenario with an alternative sequence α^B of welfare weights. As in the benchmark, sequence α^B assigns a welfare weight of 1.04 to all workers with annual incomes below the poverty line of \$11,770, i.e., to the bottom 5.6% of the skill distribution. In contrast to the benchmark, however, I assume that the welfare weights drop immediately to a constant value below the average weight of 1 for all workers with incomes above the poverty line (instead of decreasing gradually). As a result, the welfare weights also satisfy the necessary conditions

[38] The specification with $\bar{\pi} = 0.12$ corresponds to the recent estimates by Kleven (2020*a*) for single mothers, who finds much lower participation responses than the previous literature.

for an EITC with negative marginal taxes, but violate the sufficient conditions established in Proposition 3. Quantitatively, the simulation results for this case differ more strongly from the benchmark results: The marginal tax is negative for incomes below \$7,477 only, the participation tax is negative for incomes below \$18,162, and the maximum credit is reduced to \$1,078 (see green line in Figure 5). Qualitatively, however, my results remain unchanged: The optimal income tax involves an EITC that is larger than the actual 2015 EITC for childless singles.

Figure 6: Sensitivity analysis, part 2.

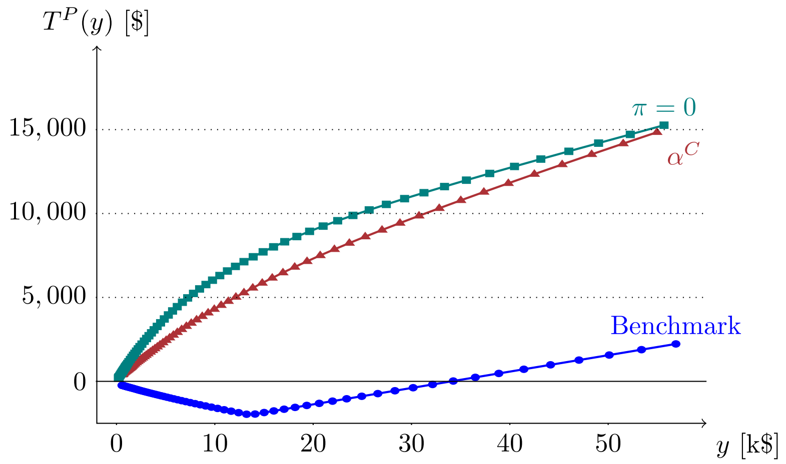

Notes: Figure 6 plots the optimal participation taxes for the benchmark case (blue line with dots) and alternative specifications with zero participation responses $\pi = 0$ (teal line with squares), or welfare weights α^C with a large skill gradient at the bottom (brown lines with triangles). For each specification, the marks represent the tuples $(y_j, T^P(y_j))$ for the workers in the different skill group.

In contrast, the simulation results change completely for the last two alternative specifications, which I illustrate in Figure 6. Specifically, the brown line with triangles depicts the optimal participation taxes for a specification with a sequence α^C of welfare weights that are convexly decreasing over the skill distribution, as in the calibrations of Saez (2002) and Jacquet et al. (2013). This case represents a society with a strong desire for redistribution among low-skill worker (i.e., from the poor to the very poor). The teal line with squares depicts the optimal tax schedule for a specification with the benchmark welfare weights α^A, but without labor supply responses at the extensive margin. In this case, labor supply responds only at the intensive margin as in Mirrlees (1971).[39] In both cases, the optimal marginal taxes and participation taxes are strictly positive everywhere below the top. Moreover, the optimal unemployment benefit is close to \$12,000 the two last cases, while is was below \$2,500 in the benchmark simulation and in each of the specifications depicted in Figure 5.

To sum up, my sensitivity analysis suggests that an EITC with negative marginal and participation taxes at the bottom remains optimal as long as, first, labor supply responds at both margins and, second, society considers the poor and the very poor almost equally

[39]To maintain consistency with the CPS data, I assume in this scenario that around 20% of the population are unproductive and do not work irrespective of the tax schedule.

deserving (i.e., if welfare weights are flat at the bottom). In contrast, an EITC cannot be optimal if labor supply responds only at the intensive margin, or if society considers the very poor much more deserving than the poor (i.e., if welfare weights are steeply falling at the bottom). Online Appendix D.3 provides further details and numerical results for additional specifications.

7 Conclusion

This paper has studied optimal income taxation in a model with labor supply responses at both the intensive margin and the extensive margin, as in Saez (2002) and Jacquet et al. (2013). Building on this earlier work, it is the first paper to provide both necessary and sufficient conditions for the optimality of an Earned Income Tax Credit with negative marginal taxes and negative participation taxes at the bottom. In particular, I show that an EITC is optimal if, first, participation elasticities are non-increasing along the skill distribution, second, society has only weak concerns for redistribution from the poor to the very poor and, third, there is a sufficiently large share of high-skill workers. The case for an EITC is particularly strong if the government's objective is to alleviate poverty, and if all agents with earned incomes below the poverty line are considered as equally deserving. As shown above, this result is driven by a tradeoff between intensive-margin and extensive-margin aspects of efficiency, which has not been discussed in the previous literature.

References

Bargain, O., Dolls, M., Neumann, D., Peichl, A. & Siegloch, S. (2014), 'Comparing inequality aversion across countries when labor supply responses differ', *International Tax and Public Finance* **21**(5), 845–873.

Bastani, S., Moberg, Y. & Selin, H. (2020), 'The Anatomy of the Extensive Margin Labor Supply Response', *Scandinavian Journal of Economics* (forthcoming).

Bastian, J. & Jones, M. R. (2020), Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Public Assistance Spending, Working paper. URL: *http://jacobbastian.squarespace.com/research*.

Beaudry, P., Blackorby, C. & Szalay, D. (2009), 'Taxes and employment subsidies in optimal redistribution programs', *American Economic Review* **99**(1), 216–42.

Besley, T. & Coate, S. (1992), 'Workfare versus welfare: Incentive arguments for work requirements in poverty-alleviation programs', *American Economic Review* **82**(1), 249–261.

Besley, T. & Coate, S. (1995), 'The design of income maintenance programmes', *Review of Economic Studies* **62**(2), 187–221.

Bierbrauer, F., Boyer, P. & Hansen, E. (2020), 'Pareto-improving tax reforms and the Earned Income Tax Credit'. CESifo Working Paper No. 8358, URL: *https://www.cesifo.org/DocDL/cesifo1_wp8358.pdf*.

Chetty, R., Friedman, J. N. & Saez, E. (2013), 'Using differences in knowledge across neighborhoods to uncover the impacts of the EITC on earnings', *American Economic Review* **103**(7), 2683–2721.

Chetty, R., Guren, A., Manoli, D. & Weber, A. (2013), 'Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities', *NBER Macroeconomics Annual* **27**(1), 1–56.

Choné, P. & Laroque, G. (2010), 'Negative marginal tax rates and heterogeneity', *American Economic Review* **100**(5), 2532–47.

Choné, P. & Laroque, G. (2011), 'Optimal taxation in the extensive model', *Journal of Economic Theory* **146**(2), 425–453.

Diamond, P. (1980), 'Income taxation with fixed hours of work', *Journal of Public Economics* **13**(1), 101–110.

Diamond, P. A. (1998), 'Optimal income taxation: An example with a u-shaped pattern of optimal marginal tax rates', *American Economic Review* **88**(1), 83–95.

Executive Office of the President and U.S. Treasury Department (2014), The President's proposal to expand the Earned Income Tax Credit, March 3, 2014. URL: *https://obamawhitehouse.archives.gov/sites/default/files/docs/eitc_report_0.pdf*.

Foster, J., Greer, J. & Thorbecke, E. (1984), 'A class of decomposable poverty measures', *Econometrica* **52**(3), 761–766.

Hansen, E. (2018), Optimal income taxation with labor supply responses at two margins: When is an Earned Income Tax Credit optimal?, Working paper. URL: *https://cmr.uni-koeln.de/sites/cmr/pdf/Hansen_Emanuel/Hansen_EITCOptimal_June2018.pdf*.

Holtz-Eakin, D., Joulfaian, D. & Rosen, H. S. (1993), 'The Carnegie Conjecture: Some empirical evidence', *Quarterly Journal of Economics* **108**(2), 413–435.

House Budget Committee Majority Staff, Chairman Paul Ryan (2014), Expanding opportunity in America. A discussion draft from the House Budget Committee, July 24, 2014. URL: *https://republicans-budget.house.gov/uploadedfiles/expanding_opportunity_in_america.pdf*.

Imbens, G. W., Rubin, D. B. & Sacerdote, B. I. (2001), 'Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players', *American Economic Review* **91**(4), 778–794.

Jacquet, L., Lehmann, E. & Van der Linden, B. (2013), 'Optimal redistributive taxation with both extensive and intensive responses', *Journal of Economic Theory* **148**(5), 1770–1805.

Juhn, C., Murphy, K. M. & Topel, R. H. (1991), 'Why has the natural rate of unemployment increased over time?', *Brookings Papers on Economic Activity* **1991**(2), 75–142.

Juhn, C., Murphy, K. M. & Topel, R. H. (2002), 'Current unemployment, historically contemplated', *Brookings Papers on Economic Activity* **2002**(1), 79–136.

Kanbur, R., Keen, M. & Tuomala, M. (1994), 'Optimal non-linear income taxation for the alleviation of income-poverty', *European Economic Review* **38**, 1613–1632.

Kleven, H. (2020*a*), The EITC and the Extensive Margin: A Reappraisal, NBER Working Paper 26405. URL: *http://www.nber.org/papers/w26405*.

Kleven, H. J. (2020*b*), 'Sufficient statistics revisited', *Annual Review of Economics* (forthcoming).

Lockwood, B. B. (2020), 'Optimal income taxation with present bias', *American Economic Journal: Economic Policy* **12**(4), 298–327.

Lorenz, N. & Sachs, D. (2016), 'Identifying Laffer bounds: A sufficient-statistics approach with an application to Germany', *Scandinavian Journal of Economics* **118**(4), 646–665.

Maag, E., Steuerle, C. E., Chakravarti, R. & Quakenbush, C. (2012), 'How marginal tax rates affect families at various levels of poverty', *National Tax Journal* **65**(4), 759–782.

Meghir, C. & Phillips, D. (2010), Labour supply and taxes, *in* J. Mirrlees, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles & J. Poterba, eds, 'Dimensions of Tax Design: The Mirrlees Review', Oxford University Press, chapter 3, pp. 202–274.

Meyer, B. D. & Rosenbaum, D. T. (2001), 'Welfare, the Earned Income Tax Credit and the labor supply of single mothers', *Quarterly Journal of Economics* **116**(3), 1063–1114.

Miller, C., Katz, L. F., Azurdia, G., Isen, A., Schultz, C. & Aloisi, K. (2018), 'Boosting the Earned Income Tax Credit for Singles: Final Impact Findings from the Paycheck Plus Demonstration in New York City'. MDRC. URL: *https://www.mdrc.org/publication/boosting-earned-income-tax-credit-singles*.

Mirrlees, J. A. (1971), 'An exploration in the theory of optimum income taxation', *Review of Economic Studies* **38**(114), 175–208.

Nichols, A. & Rothstein, J. (2015), The Earned Income Tax Credit (EITC), NBER Working Paper No. 21211. URL: *http://www.nber.org/papers/w21211*.

OECD (2017), *Taxing Wages 2017*, OECD Publishing, Paris.

Piketty, T. (1997), 'La redistribution fiscal face au chômage', *Revue française d'économie* **12**(1), 157–201.

Rochet, J.-C. & Stole, L. A. (2002), 'Nonlinear pricing with random participation', *Review of Economic Studies* **69**(1), 277–311.

Saez, E. (2001), 'Using elasticities to derive optimal income tax rates', *Review of Economic Studies* **68**, 205–229.

Saez, E. (2002), 'Optimal income transfer programs: Intensive versus extensive labor supply responses', *Quarterly Journal of Economics* **117**(3), 1039–1073.

Saez, E., Slemrod, J. & Giertz, S. H. (2012), 'The elasticity of taxable income with respect to marginal tax rates: A critical review', *Journal of Economic Literature* **50**(1), 3–50.

Saez, E. & Stantcheva, S. (2016), 'Generalized social welfare weights for optimal tax theory', *American Economic Review* **106**(1), 24–45.

Weinzierl, M. (2014), 'The promise of positive optimal taxation: Normative diversity and a role for equal sacrifice', *Journal of Public Economics* **118**, 128–142.

Appendix

A Proofs for Section 4

Section 4 studies the revenue and welfare effects of the tax perturbations depicted in Figures 2a and 2b. The following proofs build on the perturbation approach developed by Piketty (1997) and Saez (2002), and on its extension to multi-bracket reforms in Bierbrauer et al. (2020). I start with a preliminary result that prepares for the proofs of Lemmas 1 and 2.

Lemma A.1. *Let the skill set be given by $\Omega = [\underline{\omega}, \bar{\omega}]$. Consider a continuously differentiable, weakly convex tax schedule T such that $y_T(\underline{\omega}) > 0$, and a generic tax reform (τ, h) that replaces $T(y)$ by $T(y) + \tau h(y)$. For $\tau = 0$, the marginal revenue effect of this reform is given by*

$$\begin{aligned} R_\tau(0,h) \quad &= \quad F_y(0)\ h(0) + \int_{\underline{y}}^{\bar{y}} f_y(y)\ h(y)\ dy - \int_{\underline{y}}^{\bar{y}} f_y(y)\ h'(y)\ y\ \varepsilon(y)\ \frac{T'(y)}{1-T'(y)}\ dy \\ &\quad - \int_{\underline{y}}^{\bar{y}} f_y(y)\ [h(y) - h(0)]\ \eta(y)\ T^P(y)\ dy\ , \end{aligned} \tag{A.1}$$

where $\varepsilon(y) = \varepsilon_T(\omega_T(y))$, $\eta(y) = \eta_T(\omega_T(y))$ and $\omega_T = y_T^{-1}$ is the inverse function of y_T.

Proof. Consider a tax reform (τ, h), where the scalar τ gives the magnitude and the function $h : y \mapsto h(y)$ gives the direction of the reform. After this reform, net tax revenue can be written as $R(\tau, h) = E_\Omega [R(\tau, h \mid \omega)]$, where $R(\tau, h \mid \omega)$ is the net revenue collected from individuals with skill type ω. This skill-specific revenue is given by

$$R(\tau, h \mid \omega) = \hat{p}(\omega,\tau)\,[T(\hat{y}(\omega,\tau)) + \tau h(\hat{y}(\omega,\tau))] + (1-\hat{p}(\omega,\tau))\ \ [T(0) + \tau h(0)]\ ,$$

where $\hat{y}(\omega,\tau)$ denotes the optimal earnings of agents with skill type ω conditional on working, and $\hat{p}(\omega,\tau)$ denotes the participation rate among the agents with skill type ω. The derivative of the skill-specific revenue with respect to τ is given by

$$\begin{aligned} R_\tau(\tau, h \mid \omega) \quad &= \quad \hat{p}(\omega,\tau)\ h(\hat{y}(\omega,\tau)) + [1-\hat{p}(\omega,\tau)]\ \ h(0) \\ &\quad + \hat{p}(\omega,\tau)\ \hat{y}_\tau(\omega,\tau)\ \ \left[T'(\hat{y}(\omega,\tau)) + \tau h'(\hat{y}(\omega,\tau))\right] \\ &\quad + \hat{p}_\tau(\omega,\tau)\ \ [T(\hat{y}(\omega,\tau)) + \tau h(\hat{y}(\omega,\tau)) - T(0) - \tau h(0)]\ , \end{aligned}$$

where the first line captures the mechanical effects for working and non-working agents in this skill group, the second line captures revenue changes due to intensive-margin responses that are induced by changes in the marginal tax, and the third line captures revenue changes due to extensive-margin responses that are induced by changes in the participation tax. The derivatives of $\hat{y}(\omega,\tau)$ and $\hat{p}(\omega,\tau)$ with respect to the reform magnitude τ are given by

$$\begin{aligned} \hat{y}_\tau(\omega,\tau) \quad &= \quad -h'(\hat{y}(\omega,\tau))\ \frac{\partial \hat{y}(\omega,\tau)}{\partial[1-T'(\hat{y}(\omega,\tau))]}\ ,\ \text{and} \\ \hat{p}_\tau(\omega,\tau) \quad &= \quad -\left[h(\hat{y}(\omega,\tau)) - h(0)\right]\ \frac{\partial \hat{p}(\omega,\tau)}{\partial\left[\hat{y}(\omega,\tau) - T^P(\hat{y}(\omega,\tau))\right]}\ . \end{aligned}$$

For the next step, I evaluate the derivative $R_\tau(\tau, h \mid \omega)$ for $\tau = 0$. Using $\hat{y}(\omega, 0) = y_T(\omega)$ and

$\hat{p}(\omega,0) = p_T(\omega)$, and rewriting the intensive-margin and extensive-margin responses in terms of the (semi-)elasticities (8) and (9), this gives

$$\begin{aligned} R_\tau(0,h \mid \omega) \quad = \quad & p_T(\omega)\ h(y_T(\omega)) + [1-p_T(\omega)]\ \ h(0) \\ & -\ p_T(\omega)\ h'(y_T(\omega))\ \varepsilon_T(\omega)\ y_T(\omega)\ \frac{T'(y_T(\omega))}{1-T'(y_T(\omega))} \\ & -\ p_T(\omega)\ \ [h(y_T(\omega)) - h(0)]\ \ \eta_T(\omega)\ T^P(y_T(\omega))\ . \end{aligned}$$

Denote by g_ω the density of the unconditional skill distribution. Then, integrating over all skill groups, the marginal revenue effect R_τ results as

$$\begin{aligned} R_\tau(0,h) \quad = \quad & \int_{\underline{\omega}}^{\bar{\omega}} g_\omega(\omega)\, [1-p_T(\omega)]\, d\omega\ h(0) + \int_{\underline{\omega}}^{\bar{\omega}} g_\omega(\omega) p_T(\omega) h(y_T(\omega)) d\omega \\ & - \int_{\underline{\omega}}^{\bar{\omega}} g_\omega(\omega)\ p_T(\omega)\ h'(y_T(\omega))\ \varepsilon_T(\omega)\ y_T(\omega)\ T'(y_T(\omega)) d\omega \\ & - \int_{\underline{\omega}}^{\bar{\omega}} g_\omega(\omega)\ p_T(\omega)\ \ [h(y_T(\omega)) - h(0)]\ \ \eta_T(\omega)\ T^P(y_T(\omega)) d\omega\ . \end{aligned}$$

For any continuously differentiable and weakly convex tax schedule T, $y_T(\omega)$ is strictly increasing in ω due to the single-crossing property $h_{y\omega} < 0$. This implies that optimal income choices give rise to an income distribution such that $F_y(0) = \int_{\underline{\omega}}^{\bar{\omega}} g_\omega(\omega)\ [1-p_T(\omega)] d\omega$, $F_y(y) = F_y(0) + \int_{\underline{\omega}}^{\omega_T(y)} g_\omega(\omega)\ p_T(\omega)\ d\omega$ and $f_y(y) = g_\omega(\omega_T(y))\ p_T(\omega_T(y))\ \partial\omega_T(y)/\partial y$. Hence, a substitution of the integration variable gives expression (A.1) in Lemma A.1.

□

Proof for Lemma 1

Proof. Consider the tax reform (τ, h_a) depicted in Figure 2a. Formally, the direction h_a of this reform is given by

$$h_a(y) = \begin{cases} \ell & \text{for} \quad y = 0\ , \\ -\phi_a \ell & \text{for} \quad y \in (0, y_a - \ell]\ , \\ -\phi_a(y_a - y) & \text{for} \quad y \in (y_a - \ell, y_a)\ , \\ 0 & \text{for} \quad y \geq y_a \end{cases} \tag{A.2}$$

for any $\ell > 0$. The derivative $h_a'(y)$ equals ϕ_a for incomes in $(y_a - \ell, y_a)$ and zero otherwise. By Lemma A.1, the marginal revenue of this reform applied to a generic tax schedule T is

$$\begin{aligned} R_\tau(0,h_a) \quad = \quad & \ell\ F_y(0) - \phi_a \ell \int_{\underline{y}}^{y_a-\ell} f_y(y) dy - \phi_a \int_{y_a-\ell}^{y_a} f_y(y)(y_a - y) dy \\ & -\phi_a \int_{y_a-\ell}^{y_a} f_y(y)\ y\ \varepsilon(y) \frac{T'(y)}{1-T'(y)} dy \\ & +\phi_a \ell \int_{\underline{y}}^{y_a-\ell} f_y(y)\ \eta(y)\ T^P(y) dy + \phi_a \int_{y_a-\ell}^{y_a} f_y(y)\ \eta(y)\ (y_a - y)\ T^P(y) dy \\ & +\ell \int_{\underline{y}}^{\bar{y}} f_y(y)\ \eta(y)\ T^P(y) dy\ , \end{aligned}$$

where the first line captures mechanical revenue effects and the second line captures behavioral revenue effects due to intensive-margin responses. The third and fourth line captures behavioral revenue effects from extensive-margin responses to the tax cut $h(y_T(\omega))$ for low-income workers and the tax hike $h(0)$ for non-working agents, respectively. The second-order derivative of revenue R with respect to τ and ℓ, evaluated for $\tau = 0$ and $\ell = 0$, follows as

$$\begin{aligned} R_{\tau\ell}(0, h_a)|_{\ell=0} &= F_y(0) - \phi_a \left[F_y(y_a) - F_y(\underline{y})\right] \\ &\quad -\phi_a \ f_y(y_a) \ y_a \ \varepsilon(y_a) \ \frac{T'(y_a)}{1 - T'(y_a)} \\ &\quad +\phi_a \int_{\underline{y}}^{y_a} f_y(y) \ \eta(y) \ T^P(y) \ dy + \int_{\underline{y}}^{\bar{y}} f_y(y) \ \eta(y) \ T^P(y) \ dy \ . \end{aligned}$$

With $\phi_a = F_y(0)/[F_y(y_a) - F_y(\underline{y})]$, the sum of the mechanical effects in the first line is zero. Moreover, Assumption 1 with $t_p = 0$ implies that $T'(y) = 0$ and $T^P(y) = 0$ for all $y \in (0, y_p)$. Hence, for $y_a < y_p$, the second-order derivative $R_{\tau\ell}(0, h_a)|_{\ell=0}$ simplifies to

$$R_{\tau\ell}(0, h_a)|_{\ell=0} = \int_{y_p}^{\bar{y}} f_y(y) \ \eta(y) \ T^P(y) \ dy \ ,$$

which is the positive fiscal externality from a small reform.

I have assumed that this revenue gain is redistributed back using a uniform lump-sum transfer to all agents, with average welfare weight 1. Hence, the marginal welfare effect of a reform in direction h_a is given by

$$W_\tau(0, h_a) = R_\tau(0, h_a) + F_y(0) \ \hat{\alpha}_0 \ h(0) + \int_{\underline{y}}^{\bar{y}} f_y(y) \ \alpha(\omega_T(y)) \ h(y) \ dy \ ,$$

where $\hat{\alpha}_0 = E_{\Omega\times\Delta}\left[\alpha(\omega) \mid y^*(\omega, \delta) = 0\right]$ is the average welfare weight of non-working agents. Similarly, I denote by $\hat{\alpha}(\underline{y}, y_a) = E_{\Omega\times\Delta}\left[\alpha(\omega) \mid y^*(\omega, \delta) \in \left[\underline{y}, y_a\right)\right]$ the average welfare weight of workers with incomes between $\underline{y}$ and y_a. Then, the second-order derivative of welfare with respect to τ and ℓ is

$$\begin{aligned} W_{\tau\ell}(0, h_a)|_{\ell=0} &= -F_y(0) \ \hat{\alpha}_0 \ + \ \phi_a \left[F_y(y_a) - F_y(\underline{y})\right] \ \hat{\alpha}(\underline{y}, y_a) \ + \ R_{\tau\ell}(0, h_a)|_{\ell=0} \\ &= F_y(0) \ \left[\hat{\alpha}(\underline{y}, y_a) - \hat{\alpha}_0\right] \ + \ \int_{y_p}^{\bar{y}} f_y(y) \ \eta(y) \ T^P(y) \ dy \ , \end{aligned}$$

where I have again used that $\phi_a \left[F_y(y_a) - F_y(\underline{y})\right] = F_y(0)$ by construction. The welfare effect of a small reform in direction h_a with τ and ℓ positive but close to zero is approximately given by $W(\tau, h_a) - W(0, h_a) \approx \tau\ell \ W_{\tau\ell}(0, h_a)|_{\ell=0}$. This welfare effect is strictly positive if and only if condition (10) in Lemma 1 is satisfied. □

Proof of Lemma 2

Proof. The following proof proceeds along the same lines as the one for Lemma 1. For the tax reform (τ, h_b) depicted in Figure 2b, the direction h_b is given by

$$h_b(y) = \begin{cases} 0 & \text{for} \quad y = 0 \ , \\ \ell & \text{for} \quad y \in (0, y_a - \ell] \ , \\ y_a - y & \text{for} \quad y \in (y_a - \ell, y_a] \ , \\ -\phi_b (y - y_a) & \text{for} \quad y \in (y_a, y_a + \ell) \ , \\ -\phi_b \ell & \text{for} \quad y \in [y_a + \ell, y_b] \\ -\phi_b (y_b - y) & \text{for} \quad y \in (y_b, y_b + \ell) \\ 0 & \text{for} \quad y \geq y_b + \ell \ . \end{cases} \tag{A.3}$$

for any $\ell > 0$. Using equation (A.1), the marginal revenue of this reform is

$$\begin{aligned} R_\tau(0, h_b) \;=\; & \ell \int_{\underline{y}}^{y_a - \ell} f_y(y) \ dy + \int_{y_a - \ell}^{y_a} f_y(y) \ (y_a - y) \ dy \\ & -\phi_b \int_{y_a}^{y_a + \ell} f_y(y) \ (y - y_a) \ dy - \phi_b \ell \int_{y_a + \ell}^{y_b} f_y(y) \ dy - \phi_b \int_{y_b}^{y_b + \ell} f_y(y) \ (y_b - y) \ dy \\ & + \int_{y_a - \ell}^{y_a} f_y(y) \ y \ \varepsilon(y) \ \frac{T'(y)}{1 - T'(y)} \ dy + \phi_b \int_{y_a}^{y_a + \ell} f_y(y) \ y \ \varepsilon(y) \ \frac{T'(y)}{1 - T'(y)} \ dy \\ & -\phi_b \int_{y_b}^{y_b + \ell} f_y(y) \ y \ \varepsilon(y) \ \frac{T'(y)}{1 - T'(y)} \ dy \\ & -\ell \int_{\underline{y}}^{y_a - \ell} f_y(y) \ \eta(y) \ T^P(y) \ dy - \int_{y_a - \ell}^{y_a} f_y(y) \ \eta(y) \ (y_a - y) \ T^P(y) \ dy \\ & +\phi_b \int_{y_a}^{y_a + \ell} f_y(y) \ \eta(y) \ (y - y_a) \ T^P(y) \ dy + \phi_b \ell \int_{y_a + \ell}^{y_b} f_y(y) \ \eta(y) \ T^P(y) \ dy \\ & +\phi_b \int_{y_b}^{y_b + \ell} f_y(y) \ \eta(y) \ (y - y_a) \ T^P(y) \ dy \ , \end{aligned}$$

where lines 1 and 2 capture mechanical effects, and lines 3 and 4 capture behavioral effects due to intensive-margin responses for workers with incomes in $(y_a - \ell, y_a + \ell)$ and $(y_b, y_b + \ell)$, respectively. Line 5 captures behavioral effects due to extensive-margin responses for workers with incomes below y_a, and lines 6 and 7 capture behavioral effects due to extensive-margin responses for workers with incomes in $(y_a, y_b + \ell)$. The second-order derivative of R with respect to τ and ℓ, evaluated for $\tau = 0$ and $\ell = 0$, follows as

$$\begin{aligned} R_{\tau\ell}(0, h_b)|_{\ell=0} \;=\; & \int_{\underline{y}}^{y_a} f_y(y) \ dy - \phi_b \int_{y_a}^{y_b} f_y(y) \ dy \\ & +(1 + \phi_b) \ f_y(y_a) \ y_a \ \varepsilon(y_a) \ \frac{T'(y_a)}{1 - T'(y_a)} - \phi_b \ f_y(y_b) \ y_b \ \varepsilon(y_b) \ \frac{T'(y_b)}{1 - T'(y_b)} \\ & - \int_{\underline{y}}^{y_a} f_y(y) \ \eta(y) \ T^P(y) \ dy + \phi_b \int_{y_a}^{y_b} f_y(y) \ \eta(y) \ T^P(y) \ dy \ . \end{aligned}$$

With $\phi_b = \left[F_y(y_a) - F_y(\underline{y})\right] / \left[F_y(y_b) - F_y(y_a)\right]$, the sum of the mechanical effects in the first line is again zero. Moreover, Assumption 1 implies that $T'(y) = 0$ and $T^P(y) = t_p$ for all $y \in (0, y_p)$.

Hence, with $y_a < y_b < y_p$, the second-order derivative of R simplifies to

$$\begin{aligned} R_{\tau\ell}(0, h_b)|_{\ell=0} &= -t_p \left[\int_{\underline{y}}^{y_a} f_y(y)\ \eta(y)\ dy - \phi_b \int_{y_a}^{y_b} f_y(y)\ \eta(y)\ dy \right] \\ &= -t_p \left[F_y(y_a) - F_y(\underline{y}) \right]\ \left[\hat{\eta}(\underline{y}, y_a) - \hat{\eta}(y_a, y_b) \right]\ , \end{aligned}$$

where I have used that $\phi_b \left[F_y(y_b) - F_y(y_a) \right] = F_y(y_a) - F_y(\underline{y})$ and I have denoted by $\hat{\eta}(y_1, y_2) = \int_{y_1}^{y_2} f_y(y)\ \eta(y)\ dy / \left[F_y(y_2) - F_y(y_1) \right]$ the average semi-elasticity of participation among workers with incomes between y_1 and y_2. By the last line, the fiscal externality of the reform is positive if $t_p < 0$ and $\hat{\eta}(\underline{y}, y_a) > \hat{\eta}(y_a, y_b)$. Again, I have assumed that this revenue gain is redistributed back to the agents using a uniform lump-sum transfer. Hence, the second-order derivative of welfare with respect to τ and ℓ is given by

$$\begin{aligned} W_{\tau\ell}(0, h_b)|_{\ell=0} &= -\left[F_y(y_a) - F_y(\underline{y}) \right]\ \hat{\alpha}(\underline{y}, y_a)\ +\ \phi_b \left[F_y(y_b) - F_y(y_a) \right]\ \hat{\alpha}(y_a, y_b) \\ &\quad +\ R_{\tau\ell}(0, h_b)|_{\ell=0} \\ &= \left[F_y(y_a) - F_y(\underline{y}) \right]\ \left\{ \hat{\alpha}(y_a, y_b) - \hat{\alpha}(\underline{y}, y_a) - t_P \left[\hat{\eta}(\underline{y}, y_a) - \hat{\eta}(y_a, y_b) \right] \right\}\ . \end{aligned}$$

Again, the welfare effect of a small reform with $\tau > 0$ and $\ell > 0$ can be approximated by $W(\tau, h_b) - W(0, h_b) \approx \tau\ell\ W_{\tau\ell}(0, h_b)|_{\ell=0}$, which is positive if and only if condition (11) in Lemma 2 is satisfied. □

B Proofs for Subsections 5.1 and 5.2

In the following, I provide the formal proofs of Propositions 1 and 2. For this purpose, I proceed in several steps. First, I provide two preliminary results that allow to simplify the optimal tax problem for the model introduced in Section 3 with a discrete skill set Ω. Second, I provide two intermediate results in which I identify the first-and-half-best allocation and derive conditions under which this allocation satisfies or violates the incentive-compatibility constraints between workers in different skill groups. In the third and final step, I show that the labor supply distortions in the second-best allocation can be determined based on which IC constraints are violated and satisfied in the first-and-half-best allocation.

I start by showing that all implementable and second-best Pareto-efficient allocations have a distinct set of properties. Hence, I can restrict the following analysis to the set of allocation with these properties.

Lemma B.1. *Every implementable and second-best Pareto-efficient allocation is characterized by two vectors $(c_j)_{j=0}^n$ and $(y_j)_{j=1}^n$ such that*

- *all agents with skill type $\omega_j \in \Omega$ and fixed cost type δ below the threshold $\delta_j := c_j - h(y_j, \omega_j) - c_0$ receive bundle (c_j, y_j), and*
- *all other agents receive bundle $(c_0, 0)$.*

Proof. Consider an allocation $(c, y) : \Omega \times \Delta \to \mathbb{R}^2$ that is incentive-compatible and second-best Pareto-efficient. Incentive compatibility requires that, for any pair of types (ω^a, δ^a) and (ω^b, δ^b),

$$c^a - h(y^a, \omega^a) - \mathbb{1}_{y^a>0}\, \delta^a \;\geq\; c^b - h(y^b, \omega^a) - \mathbb{1}_{y^b>0}\, \delta^a \text{ and} \tag{B.1}$$

$$c^b - h(y^b, \omega^b) - \mathbb{1}_{y^b>0}\, \delta^b \;\geq\; c^a - h(y^a, \omega^b) - \mathbb{1}_{y^a>0}\, \delta^b \tag{B.2}$$

where $c^i = c(\omega^i, \delta^i)$ and $y^i = y(\omega^i, \delta^i)$ denote the consumption and output levels allocated to type (ω^i, δ^i), respectively. In the following, I show by contradiction that these IC constraints have three implications for allocation (c, y).

First, assume that, under (c, y), there are two types $(\omega^a, \delta^a) \neq (\omega^b, \delta^b)$ that provide zero output, $y^a = y^b = 0$, but receive different consumption levels $c^a \neq c^b$. Then, one of the IC constraints (B.1) or (B.2) is violated. Hence, incentive compatibility requires that all non-working agents receive the same consumption level c_0.

Second, assume that, under (c, y), there are two types with $\omega^a = \omega^b = \omega_j$ and $\delta^a \neq \delta^b$ that provide strictly positive output but receive different bundles such that $c^a - h(y^a, \omega_j) \neq c^b - h(y^b, \omega_j)$. Then, one of the IC constraints (B.1) or (B.2) is violated again. Hence, incentive compatibility requires that the gross utility $c^i - h(y^i, \omega_j)$ is equal to the same number z_j for each working agent with skill type ω_j. It does not require that the bundles (c^a, y^a) and (c^b, y^b) are identical, however.

Third, assume that, under (c, y), there are two types with $\omega^a = \omega^b = \omega_j$ and $\delta^a > \delta^b$ such that only the agent with the higher fixed cost provides positive output, $y^a > y^b = 0$. Then, the IC constraints (B.1) and (B.2) imply that $c^a - h(y^a, \omega_j) - \delta^a \geq c^b \geq c^a - h(y^a, \omega_j) - \delta^b$, respectively. Both constraints can only be satisfied simulataneously if $\delta^b \geq \delta^a$, which gives a contradiction. Hence, incentive compatibility requires that, in each skill group j, there is a unique participation threshold $\delta_j = z_j - c_0$ such that all agents with skill ω_j and fixed cost below δ_j provide strictly positive output, while all agents with skill ω_j and fixed cost above δ_j provide zero output.

Fourth, second-best Pareto efficiency requires that all working agents with skill type ω_j are allocated the same consumption-output bundle. For a proof by contradiction, assume that, in the incentive-compatible and Pareto-efficient allocation (c, y), two types with $\omega^a = \omega^b = \omega_j$ and $\delta^a \neq \delta^b$ receive different bundles that satisfy $y_a < y_b$ and $c^a - h(y^a, \omega_j) = c^b - h(y^b, \omega_j) = z_j$. I will now show that the government can achieve a Pareto improvement. For this purpose, define the set B of (c, y) bundles such that $c = k(y) = z_j + h(y, \omega_j)$ and $y \in [y^a, y^b]$. Note that, as long as the agents a and b are provided with (potentially different) bundles in B, their utility stays constant and the allocation remains incentive-compatible. For the government, the net resource cost of providing bundle $(k(y), y)$ to an agent is given by $k(y) - y = z_j + h(y, \omega_j) - y$. This net resource cost is strictly convex in y by $h_{yy} > 0$. Hence, there is a unique bundle (c_j, y_j) that minimizes the net resource cost over set B. Consequently, by changing the allocation so that both agents a and b are provided bundle (c_j, y_j), we can free up resources without making any agent worse off and use these resources to make all agents in the population better off. But this means that the allocation considered above was not Pareto-efficient.

Finally, if all working agents with skill type ω_j are allocated the same bundle (c_j, y_j), the skill-specific participation threshold is given by $\delta_j = z_j - c_0 = c_j - h(y,\omega_j) - c_0$. □

For the next step, note that, by Lemma B.1, a potentially optimal allocation is effectively given by two vectors $(c_j)_{j=1}^n$, $(y_j)_{j=1}^n$ that characterize the bundles of working types, the vector $(\delta_j)_{j=1}^n$ of participation thresholds and the consumption level c_0 of non-working agents. The next preliminary result shows that the optimal tax problem can be rewritten in a compact way that only involves the vectors $(y_j)_{j=1}^n$ and $(\delta_j)_{j=1}^n$, which characterize labor supply distortions at both margins. Hence, we can eliminate the consumption levels of working and non-working agents from the optimal tax problem, thereby simplifying the following analysis.

Lemma B.2. *The optimal allocation problem is equivalent to maximizing*

$$\sum_{j=1}^{n} f_j \left\{ G_j(\delta_j) \left[y_j - h\left(y_j, \omega_j\right) + \delta_j \left(\alpha_j - 1\right)\right] - \alpha_j \int_{\underline{\delta}}^{\delta_j} \delta dG_j(\delta) \right\} \tag{B.3}$$

over $(y_j)_{j=1}^n$ and $(\delta_j)_{j=1}^n$, subject to the incentive compatibility constraints along the skill dimension that, for any pair of skill groups j and k in $J = \{1, 2, \ldots, n\}$,

$$\delta_j - \delta_k \geq h\left(y_k, \omega_k\right) - h\left(y_k, \omega_j\right) \ . \tag{B.4}$$

Proof. I start by showing in two steps that (B.3) results from substituting the feasibility condition into the welfare function. For the first step, Lemma B.1 implies that welfare is given by

$$E_{\Omega \times \Delta} \left[\alpha(\omega) \left\{c(\omega, \delta) - h\left(y(\omega, \delta), \omega\right) - \mathbb{1}_{y(\omega,\delta)>0}\ \delta\right\}\right] =$$
$$\sum_{j=1}^{n} f_j \alpha_j \left\{ \int_{\underline{\delta}}^{\delta_j} \left[c_j - h(y_j, \omega_j) - \delta\right] dG_j(\delta) + \left[1 - G_j(\delta_j)\right] c_0 \right\} \ ,$$

with $\delta_j = c_j - h(y_j, \omega_j) - c_0$ denoting the participation threshold in skill group j. Using this threshold condition, the utility of a working type in skill group j can be rewritten as $c_j - h(y_j, \omega_j) - \delta = c_0 + \delta_j - \delta$. Hence, all agents receive at least utility c_0, and a working agent with type (ω_j, δ) receives an additional utility of $\delta_j - \delta$. We can thus rewrite welfare as

$$\sum_{j=1}^{n} f_j \alpha_j \left\{ c_0 + \int_{\underline{\delta}}^{\delta_j} [\delta_j - \delta] \, dG_j(\delta) \right\} = c_0 + \sum_{j=1}^{n} f_j \alpha_j \int_{\underline{\delta}}^{\delta_j} [\delta_j - \delta] \, dG_j(\delta) \; , \tag{B.5}$$

where I have used that the average welfare weight is given by $\sum_{j=1}^{n} f_j \alpha_j = 1$.

For the second step, Lemma B.1 implies that the feasibility condition is given by

$$E_{\Omega \times \Delta} \left[y(\omega, \delta) - c(\omega, \delta) \right] = \sum_{j=1}^{n} f_j \left\{ G_j(\delta_j)[y_j - c_j] - [1 - G_j(\delta_j)] \, c_0 \right\} \geq 0 \; .$$

Again, I can use the threshold condition to rewrite the net tax payment of a worker in skill group j as $y_j - c_j = y_j - h(y_j, \omega_j) - \delta_j - c_0$. Hence, we can interpret c_0 as the base transfer to all agents, and $y_j - h(y_j, \omega_j) - \delta_j$ as the gross tax payment of a worker in skill group j. Substituting this expression into the feasibility condition gives

$$\sum_{j=1}^{n} f_j G_j(\delta_j)[y_j - h(y_j, \omega_j) - \delta_j] \geq c_0 \; . \tag{B.6}$$

In the optimal allocation, this feasibility condition holds with equality by standard arguments. Hence, we can use (B.6) to eliminate the consumption level c_0 of non-working agents in the welfare function (B.5). After some rearrangements, this gives expression (B.3) in Lemma B.2.

Finally, consider the IC constraint between working agents with skill types ω_j and ω_k, which is usually written as $c_j - h\left(y_j, \omega_j\right) \geq c_k - h\left(y_k, \omega_j\right)$. Using the threshold conditions on δ_j and δ_k, I can replace $c_j = \delta_j + h(y_j, \omega_j) + c_0$ and $c_k = \delta_k + (y_k, \omega_k) + c_0$ in this IC constraint. Some rearrangements then give inequality (B.4) in Lemma B.2. Together, constraint (B.4) and the threshold condition also ensure that any non-working agent with type (ω_j, δ) prefers his bundle to the one of any working agent with skill type ω_k because $c_0 \geq c_j - h\left(y_j, \omega_j\right) - \delta \geq c_k - h\left(y_k, \omega_j\right) - \delta$. Finally, the threshold condition directly implies that any working agent with skill type ω_j prefers his bundle to the one of (all) non-working agents. □

I continue with two intermediate results that, first, identify the first-and-half-best allocation and, second, derive conditions under which this allocation satisfies or violates any of the incentive-compatibility (IC) constraints between workers in the skill groups 1, 2 and 3. For consistency, I continue to refer to the first-and-half-best allocation as $(\tilde{c}, \tilde{y})$, even though I characterize it in terms of the vectors $(\tilde{y}_j)_{j=1}^{n}$ and $(\tilde{\delta}_j)_{j=1}^{n}$. Based on the latter result, I can then proceed to proof Propositions 1 and 2.

Lemma B.3. *[Jacquet et al. (2013)] The first-and-half-best allocation* $(\tilde{c}, \tilde{y})$ *is defined by*

$$h_y(\tilde{y}_j, \omega_j) = 1 \; , \tag{B.7}$$

$$\tilde{\delta}_j = \tilde{y}_j - h(\tilde{y}_j, \omega_j) + \frac{\alpha_j - 1}{\tilde{\eta}_j} \; , \tag{B.8}$$

where $\tilde{\eta}_j = g_j(\tilde{\delta}_j)/G_j(\tilde{\delta}_j)$ *is the semi-elasticity of participation of skill group* j*. The first-and-half-best participation tax is given by* $T^P_{fhb}(\tilde{y}_j) = (1-\alpha_j)/\tilde{\eta}_j$*.*

Proof. Following Jacquet et al. (2013), the first-and-half-best allocation is defined as the allocation that maximizes welfare (12) subject to the feasibility condition (13) and the set of threshold conditions for all skill groups, but ignoring the IC constraints between workers with different skill types. By Lemma B.2, the feasibility condition and the threshold conditions can be substituted into the welfare function to get the objective function (B.3). Hence, allocation $(\tilde{c}, \tilde{y})$ can be derived by a point-wise maximization of (B.3) with respect to y_j and δ_j without side constraints. The first-order condition with respect to y_j requires that $\tilde{y}_j$ satisfies (B.7), i.e., maximizes $y - h(y, \omega_j)$. Hence, labor supply in the first-and-half-best allocation is undistorted at the intensive margin in each skill group j.

The first-order condition with respect to δ_j requires that

$$g_j(\tilde{\delta}_j)\left[\tilde{y}_j - h\left(\tilde{y}_j, \omega_j\right) + \tilde{\delta}_j(\alpha_j - 1) - \alpha_j \tilde{\delta}_j\right] + G_j(\tilde{\delta}_j)\left(\alpha_j - 1\right) = 0\ . \tag{B.9}$$

Rearranging this first-order condition and using $\tilde{\eta}_j = g_j(\tilde{\delta}_j)/G_j(\tilde{\delta}_j)$ gives (B.8). Finally, the first-and-half-best participation tax is given by $T^P_{fhb}(\tilde{y}_j) = T_{fhb}(\tilde{y}_j) - T_{fhb}(0) = \tilde{y}_j - \tilde{c}_j + \tilde{c}_0 = \tilde{y}_j - h(\tilde{y}_j, \omega_j) - \tilde{\delta}_j = (1-\alpha_j)/\tilde{\eta}_j$, where I have used the threshold condition on $\tilde{\delta}_j$ and (B.8). Hence, the first-and-half-best participation tax is negative (positive) if α_j is larger (smaller) than 1. □

Lemma B.4. *Let Assumption 2 hold, let* $\alpha_j < 2$ *for all* $j \in J$ *and define* $a_2 = \omega_2/\omega_1$ *and* $a_3 = \omega_3/\omega_2$*. The first-and-half-best allocation satisfies*

(i) all IC constraints if $\alpha_1 = \alpha_2 = \alpha_3 = 1$*;*

(ii) the upward IC constraint of ω_1 *workers if and only if*

$$\alpha_2 \le \beta_U(\alpha_1) := 2 - \frac{a_2^{1+\sigma}}{(2-\alpha_1)^{-1} + \sigma a_2^{1+\sigma}(a_2^{1+1/\sigma} - 1)}\ . \tag{B.10}$$

If α_1 *is strictly larger than the critical value* $\bar{\beta} := 2 - \frac{a_2^{1+\sigma} - 1}{\sigma a_2^{1+\sigma}\left(a_2^{1+1/\sigma} - 1\right)} \in (1, 2)$*, then* $\beta_U(\alpha_1)$ *is located in* $(\bar{\beta}, \alpha_1)$*. If* $\alpha_1 < \bar{\beta}$*, in contrast, then* $\beta_U(\alpha_1) < \alpha_1 < \bar{\beta}$*;*

(iii) the upward IC constraint of ω_2 *workers if* $\alpha_3 \le \alpha_2 \le \alpha_1$*;*

(iv) the downward IC constraint of ω_3 *workers if and only if*

$$\alpha_3 \ge \beta_D(\alpha_2) := 2 - \frac{a_3^{1+\sigma}}{(2-\alpha_2)^{-1} + \sigma(1 - a_3^{-1-1/\sigma})}\ . \tag{B.11}$$

Let $\phi_f(\alpha_2) = [\alpha_2 - 1]\,/\,[\alpha_2 - \beta_D(\alpha_2)]$*. There is a function* $\hat{\phi}_\omega(\alpha_1, \alpha_2) > \omega_2$ *such that* (B.11) *holds if and only if*

$$\omega_3 \;>\; \hat{\phi}_\omega(\alpha_1, \alpha_2) \text{ and} \tag{B.12}$$

$$f_3 \;>\; \phi_f(\alpha_2)\left(1 + f_1 \frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right)\ . \tag{B.13}$$

Proof. For part (i) of Lemma B.4, assume that $\alpha_j = \alpha_k = 1$ for two skill groups j and k. By Lemma B.3, in the first-and-half-best allocation, the participation threshold in skill group j is given by $\tilde{\delta}_j = \tilde{y}_j - h(\tilde{y}_j, \omega_j)$. Correspondingly, the participation threshold in skill group k is given by $\tilde{\delta}_k = \tilde{y}_k - h(\tilde{y}_k, \omega_k)$. Using these expressions, the first-and-half-best allocation satisfies the IC constraint (B.4) between workers in skill groups j and k if

$$\tilde{y}_j - h(\tilde{y}_j, \omega_j) \geq \tilde{y}_k - h(\tilde{y}_k, \omega_j) \; .$$

This condition is always satisfied because, by Lemma B.3, the first-and-half-best income $\tilde{y}_j$ is implicitly defined by $h_y(\tilde{y}_j, \omega_j) = 1$, i.e., is the unique maximizer of $y - h(y, \omega_j)$. Intuitively, incentive compatibility is ensured because the workers in both skill groups receive their laissez-faire bundles.

To prove the remaining parts of Lemma B.4, I derive closed-form expressions for the first-and-half-best allocation, which can then be used to verify the local IC constraints. For this purpose, I combine Assumption 2 with the definitions of $\tilde{y}_j$ and $\tilde{\delta}_j$ in Lemma B.3. First, with the iso-elastic effort cost function (15), income $\tilde{y}_j$ has to satisfy $h_y(\tilde{y}_j, \omega_j) = \tilde{y}_j^{1/\sigma} \omega_j^{-1-1/\sigma} = 1$. Hence, it is given by $\tilde{y}_j = \omega_j^{1+\sigma}$. This also implies that $\tilde{y}_j - h(\tilde{y}_j, \omega_j) = \omega_j^{1+\sigma}/(1+\sigma)$. Second, with a uniform distribution of fixed costs, the semi-elasticity in skill group j is given by $\tilde{\eta}_j = g_j(\tilde{\delta}_j)/G_j(\tilde{\delta}_j) = 1/\tilde{\delta}_j$ as long as $\tilde{\delta}_j < \bar{\delta}$. Plugging these expressions into the first-order condition (B.9) gives

$$\tilde{\delta}_j = \frac{\tilde{y}_2 - h(\tilde{y}_2, \omega_2)}{2 - \alpha_j} = \frac{\omega_j^{1+\sigma}}{(1+\sigma)(2-\alpha_j)} \; , \tag{B.14}$$

which is well-defined and increasing in ω_j and α_j for any $\alpha_j < 2$. Using the closed-form expressions for $\tilde{y}_j$ and $\tilde{\delta}_j$, the first-and-half-best allocation satisfies the IC constraint (B.4) between workers in skill groups j and k if and only if

$$\frac{\omega_j^{1+\sigma}}{(1+\sigma)(2-\alpha_j)} - \frac{\omega_k^{1+\sigma}}{(1+\sigma)(2-\alpha_k)} \quad \geq \quad \frac{\sigma}{1+\sigma} \omega_k^{1+\sigma} \left[1 - \left(\frac{\omega_k}{\omega_j} \right)^{1+1/\sigma} \right] \; . \tag{B.15}$$

For parts (ii) and (iii), consider the local upward IC constraint of ω_j workers, which is given by inequality (B.15) with $k = j+1$. Rearranging this inequality and substituting $a_{j+1} = \omega_{j+1}/\omega_j$ gives

$$\begin{aligned} \frac{1}{2-\alpha_j} - \frac{a_{j+1}^{1+\sigma}}{2-\alpha_{j+1}} \quad &\geq \quad \sigma a_{j+1}^{1+\sigma} \left(1 - a_{j+1}^{1+1/\sigma} \right) \\ \Leftrightarrow \alpha_{j+1} \quad &\leq \quad \tilde{\beta}_U(\alpha_j, a_{j+1}) \coloneqq 2 - \frac{a_{j+1}^{1+\sigma}}{(2-\alpha_j)^{-1} + \sigma a_{j+1}^{1+\sigma} (a_{j+1}^{1+1/\sigma} - 1)} \; . \end{aligned} \tag{B.16}$$

Note that function $\tilde{\beta}_U$ is continuously differentiable and strictly increasing in α_j and smaller than 2 for any $\alpha_j < 2$ and $a_{j+1} > 1$. Hence, we know that the first-and-half-best allocation satisfies the upward IC of ω_j workers if $\alpha_{j+1} \leq \tilde{\beta}_U(\alpha_j, a_{j+1})$, and violates it if $\alpha_{j+1} > \tilde{\beta}_U(\alpha_j, a_{j+1})$.

Next, I show that, fixing any $a_{j+1} > 1$, function $\tilde{\beta}_U(\; , a_{j+1})$ has a unique fixed point $\tilde{\beta}(a_{j+1})$, which is located in $(1, 2)$. For this purpose, I can rearrange the condition $\tilde{\beta}_U(\alpha_j, a_{j+1}) \lesseqgtr \alpha_j$ to read

$$\alpha_j \gtreqless \tilde{\beta}(a_{j+1}) \coloneqq 2 - \frac{a_{j+1}^{1+\sigma} - 1}{\sigma\, a_{j+1}^{1+\sigma}\left(a_{j+1}^{1+1/\sigma} - 1\right)} .$$

The fixed point $\tilde{\beta}(a_{j+1})$ is uniquely defined and below 2 for any $a_{j+1} > 1$ and $\sigma > 0$. Hence, I know that the upward IC is violated if $\alpha_j = \alpha_{j+1} \in (\tilde{\beta}(a_{j+1}), 2)$. By Lemma B.4 (i), however, the IC is satisfied if $\alpha_j = \alpha_{j+1} = 1$. Hence, $\tilde{\beta}(a_{j+1})$ must be located in $(1, 2)$ for any combination of $a_{j+1} > 1$ and $\sigma > 0$. Finally, as $\tilde{\beta}_U$ is strictly increasing in α_j, $\alpha_j \gtreqless \tilde{\beta}(a_{j+1})$ implies $\tilde{\beta}_U(\alpha_j, a_{j+1}) \gtreqless \tilde{\beta}(a_{j+1})$.

For part (ii), note that $\beta_U(\alpha_1) = \tilde{\beta}_U(\alpha_1, a_2)$ and $\bar{\beta} = \tilde{\beta}(a_2)$. Hence, the previous arguments imply that the upward IC of ω_1 workers is satisfied if $\alpha_2 \leq \beta_U(\alpha_1) = \tilde{\beta}_U(\alpha_1, a_2)$. Moreover, $\beta_U(\alpha_1) \in (\bar{\beta}, \alpha_1)$ for any $\alpha_1 \in (\bar{\beta}, 2)$, while $\beta_U(\alpha_1) > \alpha_1$ for any $\alpha_1 < \bar{\beta}$.

For part (iii), note that $\alpha_3 \leq \alpha_2 \leq \alpha_1$ implies that α_3 is smaller than the average welfare weight 1, and strictly smaller than $\tilde{\beta}(a_3) > 1$. Hence, we have $\alpha_3 \leq \alpha_2 < \tilde{\beta}_U(\alpha_2, a_3)$ if $\alpha_2 < \tilde{\beta}(a_3)$. We have $\alpha_3 < \tilde{\beta}(a_3) \leq \tilde{\beta}_U(\alpha_2, a_3)$ if $\alpha_2 \geq \tilde{\beta}(a_3)$. In both cases, we have $\alpha_3 < \tilde{\beta}_U(\alpha_2, a_3)$, which ensures that the local upward IC of ω_2 workers is satisfied.

For part (iv), finally, I consider the downward IC constraint of ω_3 workers. Rearranging inequality (B.15) with $j = 3$ and $k = 2$ and using $a_3 = \omega_3/\omega_2$ gives

$$\begin{aligned} \frac{a_3^{1+\sigma}}{2-\alpha_3} - \frac{1}{2-\alpha_2} &\geq \sigma\left(1 - a_3^{-1-1/\sigma}\right) \\ \Leftrightarrow \alpha_3 &\geq 2 - \frac{a_3^{1+\sigma}}{(2-\alpha_2)^{-1} + \sigma\left(1 - a_3^{-1-1/\sigma}\right)} = \beta_D(\alpha_2) . \end{aligned}$$

By the normalization of welfare weights, the welfare weight of ω_3 workers equals $\alpha_3 = \alpha_2 - [\alpha_2 - 1 + f_1(\alpha_1 - \alpha_2)]/f_3$. This implies that the condition $\alpha_3 \geq \beta_D(\alpha_2)$ is equivalent to

$$f_3 \geq \frac{\alpha_2 - 1}{\alpha_2 - \beta_D(\alpha_2)}\left(1 + f_1\frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right) \tag{B.17}$$

where the first bracket equals function $\phi_f(\alpha_2)$ in part (iv) of Lemma B.4. Because $f_1+f_2+f_3 = 1$, inequality (B.17) can only be satisfied if its right-hand side is smaller than $1 - f_1$. This is equivalent to the condition $\beta_D(\alpha_2) < (1 - f_1\alpha_1)/(1 - f_1)$ or, using (B.11),

$$\begin{aligned} 2 - \frac{a_3^{1+\sigma}}{(2-\alpha_p)^{-1} + \sigma\left(1 - a_3^{-1-1/\sigma}\right)} &< \frac{1 - f_1\alpha_1}{1 - f_1} \\ \Leftrightarrow Z(a_3) \coloneqq a_3^{1+\sigma} + \left(\sigma a_3^{-1-1/\sigma} - \sigma - (2-\alpha_2)^{-1}\right)\left(1 + (\alpha_1 - 1)\frac{f_1}{1 - f_1}\right) &> 0 . \end{aligned}$$

In the final step of the proof, I show that, for any combination of $\alpha_1 \in (1, 2)$, $\alpha_2 \in (1, \alpha_1]$, $\sigma > 0$ and $f_1 \in (0, 1)$, there is a unique threshold $\hat{\phi}_\omega(\alpha_1, \alpha_2) > \omega_2$ such that the condition $Z(a_3) > 0$ is satisfied if and only if $\omega_3 > \hat{\phi}_\omega(\alpha_1, \alpha_2)$. For this purpose, note first that $Z(1) = 1 - (2 - \alpha_2)^{-1}[1 + (\alpha_1 - 1)f_1/(1 - f_1)]$, which is strictly negative because both $(2 - \alpha_2)^{-1}$ and $[1 + (\alpha_1 - 1)f_1/(1 - f_1)]$ are larger 1. Second, $Z(a_3)$ converges to infinity for $a_3 \to +\infty$. Third, the derivative of Z in a_3 is given by

$$Z'(a_3) = (1+\sigma)a_3^{\sigma} - (1+\sigma)a_3^{-2-1/\sigma}\left[1 + (\alpha_1 - 1)\frac{f_1}{1-f_1}\right] ,$$

which is strictly positive if and only if $a_3 > \tilde{a} := [1 + (\alpha_1 - 1)f_1/(1 - f_1)]^{1/(2+\sigma+1/\sigma)} > 1$. Thus, I have shown that $Z(1) < 0$, $Z'(a_3) \leq 0$ for all $a_3 \leq \tilde{a}$, $Z'(a_3) > 0$ for all $a_3 > \tilde{a}$ and $\lim_{a_3 \to \infty} Z(a_3) = \infty$. This implies that there is a unique threshold $\phi_a(\alpha_1, \alpha_2) > 1$ such that $Z(a_3) > 0$ if and only if $a_3 = \omega_3/\omega_2 > \phi_a(\alpha_1, \alpha_2)$ or, equivalently, $\omega_3 > \phi_\omega(\alpha_1, \alpha_2) := \omega_2\ \phi_a(\alpha_1, \alpha_2) > \omega_2$. Note finally that Z is strictly decreasing in α_1 such that $\hat{\phi}(\alpha_1, \alpha_2) \geq \hat{\phi}(\alpha_2, \alpha_2)$ for $\alpha_1 \geq \alpha_2$. □

The following corollary provides the conditions under which the first-and-half-best allocation satisfies the upward IC of ω_1 workers and the downward IC of ω_3 workers in the case in which $\alpha_1 = \alpha_2 > 1$ (i.e., the social planner has no concerns for redistribution at the bottom). As they follow directly from the conditions in Lemma B.4, I state the corollary without further proof.

Corollary 1. *Let Assumption 2 hold and $\alpha_1 = \alpha_2 = \alpha_p > 1$. The first-and-half-best allocation satisfies*

(i) the upward IC constraint of ω_1 workers if and only if $\alpha_p > \bar{\beta}$;

(ii) the downward IC constraint of ω_3 workers if and only if $f_3 > \phi_f(\alpha_p)$ and $a_3 > \phi_\omega(\alpha_p) =: \hat{\phi}_\omega(\alpha_p, \alpha_p)$.

Proof of Proposition 1 (i) and 2 (i)

Proof. By Proposition 2 (i), the optimal allocation can only involve upwards distortions at the intensive margin if the condition $\alpha_2 > \beta_U(\alpha_1)$ is satisfied, i.e., if the first-and-half-best allocation violates the upward IC of ω_1 workers. Proposition 2 (i) focuses on the special case where $\alpha_1 = \alpha_2 = \alpha_p$, in which the condition $\alpha_2 > \beta_U(\alpha_1)$ is equivalent to $\alpha_p > \bar{\beta}$.

To proof that $\alpha_2 > \beta_U(\alpha_1)$ is a necessary condition for upwards distortions at the intensive margin, note first that the optimal allocation can be obtained by maximizing the objective (B.3) subject to the subset of the local IC constraints (B.4) that are binding. By standard arguments, optimal labor supply y_j^* in skill group j can only upwards distorted at the intensive margin if the upward IC of ω_{j-1} workers is binding. Hence, labor supply in skill group 1 can never be upwards distorted in the optimum. In the following, I study which IC constraints can be binding if $\alpha_2 \leq \beta_U(\alpha_1)$ holds, i.e., if the necessary condition is violated. By Lemma B.4, this condition ensures that the first-and-half-best allocation $(\tilde{c}, \tilde{y})$ satisfies the upward IC of ω_1 workers. By Lemma B.4 part (iii), moreover, $(\tilde{c}, \tilde{y})$ never violates the upward IC of ω_2 workers.

There remain three possible cases. First, if $(\tilde{c}, \tilde{y})$ violates the downward ICs of ω_2 and ω_3 workers, both downward ICs are unambiguously binding in the optimal allocation. Moreover, both y_2^* and y_3^* are downwards distorted. Second, if $(\tilde{c}, \tilde{y})$ only violates the downward IC of ω_2 workers, this IC is unambiguously binding in the optimal allocation. The downward IC of ω_3 workers may be binding or slack in the optimal allocation, but the upward IC of ω_2 workers is always slack. Hence, y_1^* is downwards distorted and y_2^* is either undistorted or downwards distorted. Formal proofs for these two cases are available on request.

Third, if $(\tilde{c}, \tilde{y})$ only violates the downward IC of ω_3 workers, this IC is unambiguously binding in the optimal allocation. If this were the only binding IC constraints, δ_2 would be smaller than $\tilde{\delta}_2$ and y_2 would be smaller than $\tilde{y}_2$. While $\delta_2 < \tilde{\delta}_2$ makes the bundle of ω_2 workers less attractive, $y_2 < \tilde{y}_2$ makes the bundle more attractive for ω_1 workers. If the latter effect dominates, then the downward IC of ω_3 workers may be binding in the optimal allocation as well, in addition to the binding upward IC of ω_1. Even in this case, however, y_2^* cannot be upwards distorted. Using Lemma B.2, the Lagrangian is then given by

$$
\begin{aligned}
\mathcal{L} \quad = \quad & \sum_{j=1}^{3} f_j \left[G_j(\delta_j) \left(y_j - h\left[y_j, \omega_j\right)\right] + \delta_j \left[\alpha_j - 1\right]\right) - \alpha_j \int_{\underline{\delta}}^{\delta_j} \delta dG_j(\delta) \right] \\
& + \mu_1^U \left[\delta_1 - \delta_2 - h(y_2, \omega_2) + h(y_2, \omega_1)\right] \\
& + \mu_3^D \left[\delta_3 - \delta_2 - h(y_2, \omega_2) + h(y_2, \omega_3)\right] \ , \qquad \text{(B.18)}
\end{aligned}
$$

where μ_1^U denotes the Lagrange multiplier of the upward IC of ω_1 workers and μ_3^D denotes the multiplier of the downward IC of ω_3 workers. For $j \in \{1, 3\}$, the first-order condition with respect to y_j implies that $1 - h_y(y_j^*, \omega_j) = 0$. Hence, y_j^* is equal to $\tilde{y}_j$ and undistorted at the intensive margin for $j \in \{1, 3\}$. As $G_j(\delta) = \delta/\bar{\delta}$ and $g_j(\delta)/G_j(\delta) = \delta$, the first-order conditions with respect to δ_2 and δ_1 require that

$$
\begin{aligned}
\delta_1^* \quad &= \quad \frac{y_1^* - h(y_1^*, \omega_1)}{2 - \alpha_1} + \mu_1^U \frac{\bar{\delta}}{f_1(2 - \alpha_1)} > \tilde{\delta}_1 \ , \text{ and} \\
\delta_2^* \quad &= \quad \frac{y_2^* - h(y_2^*, \omega_2)}{2 - \alpha_2} - \left(\mu_1^U + \mu_3^D\right) \frac{\bar{\delta}}{f_2(2 - \alpha_2)} < \tilde{\delta}_2 \, ,
\end{aligned}
$$

where the inequalities follow because $\mu_1^U > 0$, $\mu_3^D > 0$, $\tilde{y}_2 = \arg\max y - h(y, \omega_2)$ and $\tilde{\delta}_j = \left[\tilde{y}_j - h(\tilde{y}_j, \omega_j)\right]/(2 - \alpha_j)$ by (B.8). As the upward IC of ω_1 workers is binding by construction, we also know that

$$
\delta_1^* - \delta_2^* = h(y_2^*, \omega_2) - h(y_2^*, \omega_1) > \tilde{\delta}_1 - \tilde{\delta}_2 \ .
$$

If the sufficient condition $\alpha_2 > \beta^U(\alpha_1)$ is not met, allocation $(\tilde{c}, \tilde{y})$ satisfies the upward IC of ω_1 so that $\tilde{\delta}_1 - \tilde{\delta}_2 \geq h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_1)$. Hence, we have

$$
h(y_2^*, \omega_2) - h(y_2^*, \omega_1) > h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_1) \ ,
$$

which ensures that $y_2^* < \tilde{y}_2$ by $h_{y\omega} < 0$. Hence, labor supply in skill group y_2 is unambiguously downwards distorted. Summing up, I have shown that optimal labor supply cannot be upwards distorted in any skill group if the condition $\alpha_2 > \beta^U(\alpha_1)$ is violated. Put differently, $\alpha_2 > \beta^U(\alpha_1)$ is a necessary condition for upwards distortions at the intensive margin to be optimal. □

Proof of Proposition 1 (ii) and 2 (ii)

Proof. By Proposition 2 (ii), the optimal allocation involves upwards distortions at both margins if the jointly sufficient conditions $\alpha_2 > \beta_U(\alpha_1)$, $\omega_3 > \hat{\phi}_\omega(\alpha_1, \alpha_2)$ and $f_3 > \phi_f(\alpha_2)\left(1 + f_1 \frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right)$ are satisfied. Proposition 1 (ii) focuses on the special case with $\alpha_1 = \alpha_2 = \alpha_p$, where the previously mentioned conditions are equivalent to $\alpha_p > \bar{\beta}$, $\omega_3 = \phi_\omega(\alpha_p)$ and $f_3 > \phi_f(\alpha_p)$. To prove these statements, I invoke Lemma B.4, by which the condition $\alpha_2 > \beta_U(\alpha_1)$ ensures that the first-and-half-best allocation violates the upward IC of ω_1 workers. The conditions $\omega_3 > \hat{\phi}_\omega(\alpha_1, \alpha_2)$ and $f_3 > \phi_f(\alpha_2)\left(1 + f_1 \frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right)$ jointly ensure that the first-and-half-best allocation satisfies the downward IC of ω_3 workers.

For the second-best allocation, the three conditions jointly ensure that the upward IC of ω_1 workers is binding in the second-best allocation and that the downward IC of ω_2 worker is slack by standard arguments. With respect to the IC constraints between ω_2 and ω_3 workers, there are two possible cases: (a) no other IC is binding, (b) the downward IC of ω_3 workers is binding. The latter case is possible because, if only the upward IC of ω_1 workers were binding, y_2 would be larger than $\tilde{y}_2$. This would make the bundle of ω_2 workers more attractive to ω_3 workers. In case (a), y_2^* is upwards distorted by standard arguments. I now show that y_2^* is also upwards distorted in case (b), which has not been studied in the optimal tax literature so far. Using Lemma B.2, the Lagrangian for this case is again given by (B.18). As argued above, the first-order condition with respect to y_j implies that y_j^* is equal to $\tilde{y}_j$, i.e., undistorted at the intensive margin for skill group $j \in \{1, 3\}$. As $g_j(\delta)/G_j(\delta) = \delta$, the first-order conditions with respect to $\hat{\delta}_2$ and $\hat{\delta}_3$ require that

$$
\begin{aligned}
\delta_2^* &= \frac{y_2^* - h(y_2^*, \omega_2)}{2 - \alpha_p} - \left(\mu_1^U + \mu_3^D\right) \frac{\bar{\delta}}{f_2(2 - \alpha_p)} < \tilde{\delta}_2 \text{ , and} \\
\delta_3^* &= \frac{y_3^* - h(y_3^*, \omega_1)}{2 - \alpha_3} + \mu_3^D \frac{\bar{\delta}}{f_3(2 - \alpha_3)} > \tilde{\delta}_3 \, ,
\end{aligned}
$$

where the inequalities follow because $\mu_1^U > 0$, $\mu_3^D > 0$, $\tilde{y}_2 = \arg\max y - h(y, \omega_2)$ and $\tilde{\delta}_j = \left[\tilde{y}_j - h(\tilde{y}_j, \omega_j)\right]/(2 - \alpha_j)$ by (B.8). As the downward IC of ω_3 workers is binding by construction, I also know that

$$
\delta_3^* - \delta_2^* = h(y_2^*, \omega_2) - h(y_2^*, \omega_3) > \tilde{\delta}_3 - \tilde{\delta}_2 \, .
$$

Recall that allocation $(\tilde{c}, \tilde{y})$ satisfies the downward IC of ω_3, $\tilde{\delta}_3 - \tilde{\delta}_2 \geq h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_3)$. Hence, we must have

$$
h(y_2^*, \omega_2) - h(y_2^*, \omega_3) > h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_3) \, ,
$$

which ensures that $y_2^* > \tilde{y}_2$ by $h_{y\omega} < 0$. Hence, y_2 is unambiguously upwards distorted in case (b) as well. Summing up, I have shown that optimal labor supply is upwards distorted at the intensive margin in skill group 2 if the jointly sufficient conditions on α_1, α_2, ω_3 and f_3 are met.

It remains to show that optimal labor supply is upwards distorted at the extensive margin in skill groups 1 and 2. For this purpose, consider the participation threshold of ω_1 workers. The first-order condition with respect to δ_1 requires that $\delta_1^* = \left[y_1^* - h(y_1^*, \omega_1) + \mu_1^U \bar{\delta}/f_1\right]/(2 - \alpha_1)$. As $\mu_1^U > 0$ and $2 - \alpha_1 < 1$, this implies that $\delta_1^* > y_1^* - h(y_1^*, \omega_1)$. Hence, labor supply by ω_1 workers is upwards distorted at the intensive margin.

Second, I have already shown that y_2^* is upwards distorted and, hence, larger than $\tilde{y}_2 > \tilde{y}_1$. This implies that $y_1^* - h(y_1^*, \omega_1) > y_2^* - h(y_2^*, \omega_1)$ because $y_1^* = \tilde{y}_1$ is the unique maximizer of $y - h(y, \omega_1)$. Hence, we know that $\delta_1^* > y_2^* - h(y_2^*, \omega_1)$. As the upward IC of ω_1 workers is binding, the participation threshold of ω_2 workers satisfies $\delta_2^* = \delta_1^* + h(y_2^*, \omega_1) - h(y_2^*, \omega_2)$. Substituting in $\delta_1^* + h(y_2^*, \omega_1) > y_2 *$, the optimal allocation has to satisfy $\delta_2^* > y_2^* - h(y_2^*, \omega_2)$. Hence, labor supply by ω_2 workers is upwards distorted at the extensive margin as well. □

Online Appendix (not meant for publication)

C Proofs for Subsection 5.3

I start by providing a preliminary result upon which Propositions 3 and 4 build. In particular, the following lemma provides conditions under which the *first-and-half-best* allocation $(\tilde{c}, \tilde{y})$ satisfies or violates each of the local IC constraints between workers with adjacent skill types.

Lemma C.1. *Let Assumptions 4 and 5 be satisfied. If a is sufficiently close to 1, the following statements are true for any skill group $j \in J \setminus \{n\}$:*

(i) There is a strictly increasing function $\beta_{Uj} : \mathbb{R} \to \mathbb{R}$ such that $(\tilde{c}, \tilde{y})$ satisfies the upward IC of ω_j workers if and only if $\alpha_{j+1} \leq \beta_{Uj}(\alpha_j)$. There is a number $\bar{\beta}_j > 1$ such that $\beta_{Uj}(\alpha_j) < \alpha_j$ if and only if $\alpha_j > \bar{\beta}_j$.

(ii) There is a strictly increasing function $\beta_{Dj} : \mathbb{R} \to \mathbb{R}$ such that $(\tilde{c}, \tilde{y})$ satisfies the downward IC of ω_{j+1} workers if and only if $\alpha_{j+1} \geq \beta_{Dj}(\alpha_j)$. There is a number $\underline{\beta}_j < 1$ such that $\beta_{Dj}(\alpha_j) < \alpha_j$ if and only if $\alpha_j > \underline{\beta}_j$.

Proof. Lemma C.1 establishes the existence of two functions β^U_j and β^D_j that determine for each couple (α_j, α_{j+1}) which of the IC constraints between skill groups j and $j+1$ are satisfied. The properties of both functions are illustrated in Figure 3. In the following, I first provide a complete and detailed derivation of part (i) on the existence and the properties of function β^U_j. I then sketch the corresponding proof for part (ii) on function β^D_j.

Part (i), first step: *If a is small enough, there is a number $\bar{\beta}_j > 1$ such that the upward IC of ω_j workers is violated if $\alpha_j = \alpha_{j+1} > \bar{\beta}_j$, and satisfied if $\alpha_j = \alpha_{j+1} < \bar{\beta}_j$.*

The claim generalizes the first step of the proof of Proposition 1 for arbitrary skill sets, effort cost functions and type distributions. By (B.4), allocation $(\tilde{y}, \tilde{\delta})$ satisfies the upward IC of ω_j workers if and only if

$$\tilde{\delta}_j - \tilde{\delta}_{j+1} \geq h(\tilde{y}_{j+1}, \omega_{j+1}) - h(\tilde{y}_{j+1}, \omega_j) \ . \tag{C.1}$$

First, I show that (C.1) is violated if α' is equal to some fixed number above 1 and a is close enough to 1. For this purpose, fix two numbers $\hat{a} > 1$ and $\gamma \in (1, 2)$ such that $(\tilde{y}, \tilde{\delta})$ is well-behaved and defined by first-order conditions for $\alpha' = \gamma$ and any $a \in [1, \hat{a}]$. In the limit case $a = 1$, we have $\omega_{j+1} = \omega_j$, $\tilde{y}_{j+1} = \tilde{y}_j$, $\eta_j(\delta) = \eta_{j+1}(\delta)$ and $\tilde{\delta}_{j+1} = \tilde{\delta}_j$. Hence, (C.1) is satisfied with a strict equality.

Now consider a small increase in a. For $a = 1$, the derivative of the left-hand side of (C.1) with respect to a is given by

$$\left. \frac{d\left[\tilde{\delta}_j - \tilde{\delta}_{j+1}\right]}{da} \right|_{a=1} = - \left. \frac{d\tilde{\delta}_{j+1}}{da} \right|_{a=1} = \frac{\omega_j h_\omega(\tilde{y}_j, \omega_j)\eta_j^2 + \omega_j \partial\eta_j/\partial\omega_j(\gamma - 1)}{\eta_j^2 + \partial\eta_j/\partial\delta_j(\gamma - 1)} \ ,$$

while the derivative of the right-hand side is given by

$$\left. \frac{d\left[h(\tilde{y}_{j+1}, \omega_{j+1}) - h(\tilde{y}_{j+1}, \omega_j)\right]}{da} \right|_{a=1} = \omega_j h_\omega(\tilde{y}_j, \omega_j) \ .$$

Hence, the derivative of the difference between the left-hand side and the right-hand side follows as

$$\frac{\omega_j(\gamma-1)}{\eta_j^2+\partial\eta_j/\partial\delta_j(\gamma-1)}\left[\partial\eta_j/\partial\omega_j-h_\omega(\tilde{y}_j,\omega_j)\partial\eta_j/\partial\delta_j\right] .$$

While the fraction before the bracket is strictly positive, the term in the bracket is strictly negative because $h_\omega < 0$ and because η_j is decreasing in both ω_j and δ_j by Assumption 4. Hence, there is a number $z > 1$ such that the left-hand side of (C.1) is smaller than the right-hand side for a for all $a \in (1, z)$. This implies that $(\tilde{y}, \tilde{\delta})$ violates the upward IC of ω_1 workers whenever $\alpha' > 1$ and a is close enough to 1.

Second, I show that there is a unique threshold $\bar{\beta}_j \in (1, \gamma)$ such that (C.1) is satisfied if and only if $\alpha' < \bar{\beta}_j$. The derivative of $\tilde{\delta}_k$ with respect to α_k can be calculated by implicit differentiation of (B.9). In particular, $\tilde{\delta}_k$ is strictly increasing in α_k and given by

$$\frac{d\tilde{\delta}_k}{d\alpha_k}=\left[\eta_k-\frac{\partial\eta_k}{\partial\delta_k}\left(\tilde{y}_k-h(\tilde{y}_k,\omega_k)-\tilde{\delta}_k\right)\right]^{-1}=\left[\eta_k-\frac{\partial\eta_k/\partial\delta_k}{\eta_k}\left(1-\alpha_k\right)\right]^{-1}>0 ,$$

for any $k \in J$ whenever $(\tilde{y}, \tilde{\delta})$ satisfies the second-order condition for a welfare optimum. Consequently, the difference $\tilde{\delta}_j - \tilde{\delta}_{j+1}$ is strictly decreasing in α',

$$\left.\frac{d[\tilde{\delta}_j-\tilde{\delta}_{j+1}]}{d\alpha'}\right|_{\alpha_j=\alpha_{j+1}=\alpha'}=\frac{1}{\eta_j-\frac{\partial\eta_j/\partial\delta_j}{\eta_j}(1-\alpha')}-\frac{1}{\eta_{j+1}-\frac{\partial\eta_{j+1}/\partial\delta_{j+1}}{\eta_{j+1}}(1-\alpha')} \quad < \quad 0$$

$$\Leftrightarrow \left(\frac{\partial\eta_j/\partial\delta_j}{\eta_j}-\frac{\partial\eta_{j+1}/\partial\delta_{j+1}}{\eta_{j+1}}\right)(1-\alpha') \quad < \quad \eta_j-\eta_{j+1} .$$

The first bracket on the left-hand side is equal to the semi-elasticity of the relative participation responses η_j/η_{j+1}. By Assumption 5, the absolute value of this semi-elasticity is smaller than $\eta_j - \eta_{j+1}$. For any $\alpha' \in (0, 2)$, the absolute value of $1 - \alpha'$ is smaller than 1. Consequently, the left-hand side of the previous inequality is smaller than the right-hand side of the last inequality. Hence, the left-hand side of (C.1) is strictly decreasing in α'. In contrast, the right-hand side of (C.1) does not depend on α'.

Recall from the first step 1 that, as long as a is small enough, the upward IC (C.1) is violated for $\alpha' = \gamma$. It is satisfied for $\alpha' = 1$ as shown in the proof of Proposition 1. Hence, the monotonicity in α' ensures that there is a unique threshold $\bar{\beta}_j \in (1, \gamma)$ such that the upward IC is satisfied if $\alpha' \leq \bar{\beta}_j$ and violated if $\alpha' \in (\bar{\beta}_j, 2)$.

Part (i), second step: Existence and properties of function β_{Uj}. The claim generalizes the statement proven in the first part of Proposition 2. The proof is simple and mainly exploits that $\tilde{\delta}_k$ is continuously differentiable and strictly increasing in α_k for any $k \in J$, as shown in the first step. This directly implies that the left-hand side of (C.1) is strictly increasing in α_j and strictly decreasing in α_{j+1} (and continuously differentiable in both parameters). Hence, if function β_{Uj} exists, (C.1) has to be satisfied with equality for any α_j and $\alpha_{j+1} = \beta_{Uj}(\alpha_j)$ such that there is a well-behaved welfare maximum. Moreover, β_{Uj} must be continuously differentiable and strictly increasing in α_j.

To prove the existence and the properties of function β_{Uj}, I again exploit the monotonicity of (C.1) in α_j and α_{j+1}. In the first part of this proof, I have shown that (C.1) is satisfied with equality for $\alpha_j = \alpha_{j+1} = \bar{\beta}_j$. First, fix α_j to be equal to some $\gamma > \bar{\beta}_j$. As the left-hand side of (C.1) is strictly increasing in α_j, the upward IC is also satisfied if $\alpha_j = \gamma$ and $\alpha_{j+1} = \bar{\beta}_j$. However, I have shown above that it is violated if $\alpha_{j+1} = \alpha_j = \gamma > \bar{\beta}_j$. Due to the monotonicity in α_{j+1}, there is a unique number $z \in (\bar{\beta}_j, \gamma)$ such (C.1) is satisfied with equality for $\alpha_{j+1} = z$ and $\alpha_j = \gamma$.

Second, fix α_j so be equal to some number $\gamma < \bar{\beta}_j$. As the left-hand side of (C.1) is strictly increasing in α_j, the upward IC is violated if $\alpha_j = \gamma$ and $\alpha_{j+1} = \bar{\beta}_j$. However, I have shown above that it is satisfied if $\alpha_{j+1} = \alpha_j = \gamma < \bar{\beta}_j$. Due to the monotonicity in α_{j+1}, there is a unique number $z \in (\bar{\beta}_j, \gamma)$ such (C.1) is satisfied with equality for $\alpha_{j+1} = z$ and $\alpha_j = \gamma$.

Hence, for any α_j such that $(\tilde{y}, \tilde{\delta})$ is well-defined, the function $\beta_{Uj}(\alpha_j)$ is well-defined as the unique number such that (C.1) is satisfied with equality for $\alpha_{j+1} = \beta_{Uj}(\alpha_j)$. This completes the proof of part (i) of Lemma C.1.

Part (ii): Existence and properties of function β_{Dj}. The proof of part (ii) applies similar arguments as the proof of part (i) to determine conditions on the welfare weights α_j and α_{j+1} such that allocation $(\tilde{y}, \tilde{\delta})$ satisfies the downward IC of ω_{j+1} workers, i.e.,

$$\tilde{\delta}_{j+1} - \tilde{\delta}_j \geq h(\tilde{y}_j, \omega_j) - h(\tilde{y}_j, \omega_{j+1}) \; . \tag{C.2}$$

First, it can be shown that (C.2) is violated if $\alpha_j = \alpha_{j+1}$ is equal to some number $\gamma < 1$ and a is slightly larger than 1. Again, this is done by comparing the derivatives of the left-hand side and the right-hand side of (C.2) in a, evaluated for $a = 1$. Second, for $\alpha_j = \alpha_{j+1}$ equal to some number α', the left-hand side is strictly increasing in α' (this directly follows from the observation that the left-hand side of (C.2) is equal to minus one times the left-hand side of (C.1)). For $\alpha' = 1$, (C.2) is satisfied. Hence, if a is small enough, there is a unique number $\underline{\beta}_j \in (\gamma, 1)$ such that the downward IC of ω_{j+1} workers is satisfied if $\alpha' \geq \underline{\beta}_j$ and violated if $\alpha' < \underline{\beta}_j$.

Finally, I can again exploit the monotonicity of $\tilde{\delta}_j$ and $\tilde{\delta}_{j+1}$ in the welfare weights α_j and α_{j+1}, respectively, to show that there is a uniquely defined function β_{Dj} such that (C.2) is satisfied with equality for any combination of α_j and $\alpha_{j+1} = \beta_{Dj}(\alpha_j)$. For $\alpha_j < \underline{\beta}_j$, we have $\beta_{Dj}(\alpha_j) \in (\alpha_j, \underline{\beta}_j)$. For $\alpha_j > \underline{\beta}_j$, we have $\beta_{Dj}(\alpha_j) \in (\underline{\beta}_j, \alpha_j)$. Moreover, the differentiability and monotonicity of $\tilde{\delta}_k$ in α_k ensures that β_{Dj} is continuously differentiable and strictly increasing in α_j. □

Proof of Proposition 3

Proof. The conditions in Proposition 3 imply that the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ (i) violates the upward IC of ω_j workers for at least one $j \in \{1, \ldots, k-1\}$ and either violates or satisfies with strict equality the upward IC of ω_j workers for each $j \in \{1, \ldots, k-1\}$, and (ii) satisfies the downward IC of ω_j workers for each $j \in \{k+1, \ldots, n\}$, i.e.,

$$\tilde{\delta}_j - \tilde{\delta}_{j+1} < h(\tilde{y}_{j+1}, \omega_{j+1}) - h(\tilde{y}_{j+1}, \omega_j) \quad \text{for at least one } j \in \{1, \ldots, k-1\} \; , \tag{C.3}$$

$$\tilde{\delta}_j - \tilde{\delta}_{j+1} \leq h(\tilde{y}_{j+1}, \omega_{j+1}) - h(\tilde{y}_{j+1}, \omega_j) \qquad \text{for all } j \in \{1, \ldots, k-1\} \; , \tag{C.4}$$

$$\tilde{\delta}_j - \tilde{\delta}_{j-1} \geq h(\tilde{y}_{j-1}, \omega_{j-1}) - h(\tilde{y}_{j-1}, \omega_j) \qquad \text{for all } j \in \{k, \ldots, k-1\} \; . \tag{C.5}$$

I now show that this ensures an upward distortion in the second-best output of ω_j workers, $y_j^* > \tilde{y}_j$ for all $j \in \{2, \ldots, k\}$. This requires to determine the set of binding IC constraints in the second-best allocation (y^*, δ^*). Assume that the upward IC of ω_j workers is binding for j in the set $J_U \subset J$ only and that the downward IC of ω_j workers is binding for j in the set $J_D \subset J$ only. Then, the Lagrangian of the optimal tax problem is given by

$$\begin{aligned}\mathcal{L} \;=\; & \sum_{j=1}^{n} f_j \left[G_j(\hat{\delta}_j) \Big(y_j - h\left[y_j, \omega_j\right)] + \hat{\delta}_j \left[\alpha_j - 1\right] \Big) - \alpha_j \int_{\underline{\delta}}^{\hat{\delta}_j} \delta dG_j(\delta) \right] \\ & + \sum_{j \in J_U} \mu_j^U \left[\hat{\delta}_j - \hat{\delta}_{j+1} - h(y_{j+1}, \omega_{j+1}) + h(y_{j+1}, \omega_j) \right] \\ & + \sum_{j \in J_D} \mu_j^D \left[\hat{\delta}_j - \hat{\delta}_{j-1} - h(y_{j-1}, \omega_{j-1}) + h(y_{j-1}, \omega_j) \right] \;,\end{aligned}$$

where μ_j^U and μ_j^D are the Lagrange multiplier of the upward IC and the downward IC, respectively, of ω_j workers.

In the following, I show that the upward IC of each skill group $j \in \{1, \ldots, k-1\}$ is binding and that $y_j^* > \tilde{y}_j$ for each $j \in \{2, \ldots, k\}$ in (y^*, δ^*). To provide the relevant arguments in a simple way, I go through three alternative cases for an economy with 4 skill groups.

Case 1: $k = 3$, upward IC of ω_2 workers violated. Assume that $n = 4$, $k = 3$ and that equation (C.3) holds for $j = 2$, (C.4) holds for $j = 1$ and (C.5) holds for $j = 4$. It is immediately clear that (y^*, δ^*) cannot be given by an allocation in which no IC is binding, nor by an allocation in which only some downward ICs are binding. Moreover, I can rule out the possibility that only the upward IC of ω_2 workers is binding, i.e., that $J_U = \{2\}$ and $J_D = \emptyset$. In this case, the FOCs would imply that $\delta_1^* = \tilde{\delta}_1$, $y_2^* = \tilde{y}_2$ and δ_2^* was given by

$$\delta_2^* = y_2^* - h(y_2^*, \omega_2) + \frac{\alpha_2 - 1}{\eta_2^*} + \frac{\mu_2^U}{f_2 g_2(\delta_2^*)} > \tilde{\delta}_2 = \tilde{y}_2 - h(\tilde{y}_2, \omega_2) + \frac{\alpha_2 - 1}{\tilde{\eta}_2} \;, \qquad \text{(C.6)}$$

where η_2^* and $\tilde{\eta}_2$ denote the semi-elasticities of participation of ω_2 workers in (y^*, δ^*) and in $(\tilde{y}, \tilde{\delta})$, respectively. We find that $\delta_2^* > \tilde{\delta}_2$ because $\mu_2^U > 0$ and $\eta_2^* \leq \tilde{\eta}_2$ for $\delta_2^* > \tilde{\delta}_2$ by Assumption 4 (moreover, the second-order condition must be satisfied in a welfare maximum). Combining this inequality with $\delta_1^* = \tilde{\delta}_1$, $y_2^* = \tilde{y}_2$ and (C.4) for $j = 1$ implies that $\delta_1^* - \delta_2^* < h(y_2^*, \omega_2) - h(y_2^*, \omega_1)$, i.e., the upward IC of ω_1 workers is violated. Hence, this allocation cannot be second-best. Rather, both upward ICs of ω_1 and ω_2 workers must be binding as long as the upward IC of ω_3 workers and the downward IC of ω_4 workers are not taken into account.

It remains to determine the set of binding IC constraint and the distortions in y_2^* and y_3^* when the full set of IC constraints is taken into account. For this purpose, we can easily generalize the arguments given in the proof of Propositions 1 (ii) and 2 (ii). As a result, the upward ICs of ω_1 and ω_2 workers are binding in (y^*, δ^*), while the downward IC of ω_4 workers might be binding or not. In any case, however, both y_2^* and y_3^* are upwards distorted at the intensive margin.

Case 2: $k = 3$, upward IC of ω_1 workers violated. Assume that $n = 4$, $k = 3$ and that equation (C.3) holds for $j = 1$, (C.4) holds for $j = 2$ and (C.5) holds for $j = 4$. Again, (y^*, δ^*) cannot be given by an allocation in which no IC is binding nor an allocation in which only some downward ICs are binding. Moreover, I can rule out the possibility that only the upward IC of ω_1 workers is binding, i.e., $J_U = \{1\}$ and $J_D = \emptyset$. In this case, the FOCs would imply that $\delta_3^* = \tilde{\delta}_3$, $y_3^* = \tilde{y}_3$, $y_2^* > \tilde{y}_2$ and that δ_2^* was given by

$$\delta_2^* = y_2^* - h(y_2^*, \omega_2) + \frac{\alpha_2 - 1}{\eta_2^*} - \frac{\mu_1^U}{f_2 g_2(\delta_2^*)} < \tilde{\delta}_2 = \tilde{y}_2 - h(\tilde{y}_2, \omega_2) + \frac{\alpha_2 - 1}{\tilde{\eta}_2}\,. \tag{C.7}$$

We find that $\delta_2^* < \tilde{\delta}_2$ because $\mu_1^U > 0$, $y_2^* - h(y_2^*, \omega_2) < \tilde{y}_2 - h(\tilde{y}_2, \omega_2)$ and $\eta_2^* \geq \tilde{\eta}_2$ for $\delta_2^* < \tilde{\delta}_2$ by Assumption 4. Combining this inequality with $\delta_3^* = \tilde{\delta}_3$, $y_3^* = \tilde{y}_3$ and (C.4) for $j = 2$ implies that $\delta_2^* - \delta_3^* < h(y_3^*, \omega_3) - h(y_3^*, \omega_2)$, i.e., the upward IC of ω_2 workers is violated. Again, this means that both upward ICs of ω_1 and ω_2 workers must be binding as long as the upward IC of ω_3 workers and the downward IC of ω_4 workers are not taken into account. As in case 1, we can apply the arguments given in the proof of Propositions 1 (ii) and 2 (ii) to show that both upwards ICs are binding in (y^*, δ^*) and that both y_2^* and y_3^* are upwards distorted at the intensive margin.

Case 3: $k = 2$. Assume that $n = 4$ and $k = 2$. This implies that (C.3) holds for $j = 1$ and that (C.5) holds for $j \in \{3, 4\}$. For concreteness, assume also that (C.4) is not satisfied for $j \in \{2, 3\}$. As before, (y^*, δ^*) cannot be given by an allocation in which no ICs are binding or only some downwards ICs are binding. Consider a relaxed problem that takes into account only the upward IC constraints of ω_1 and ω_2 workers, and only the downward IC constraints of ω_2 and ω_3 workers. The proof of Propositions 1 (ii) and 2 (ii) directly implies that, in the solution (y^R, δ^R) to this relaxed problem, the upward IC of ω_1 workers is binding and y_2^R is upwards distorted. Besides, there are three possible constellations: (1) If no other IC is binding in (y^R, δ^R), then the same is true in the second-best allocation (y^*, δ^*). (2) If the upward IC of ω_2 workers is binding in (y^R, δ^R), then we can apply the arguments in the proof of Propositions 1 (ii) and 2 (ii) to show the same is true in (y^*, δ^*) and that both y_2^* and y_3^* are upwards distorted. (3) If the downward IC of ω_3 workers is binding in (y^R, δ^R), then it is also binding in (y^*, δ^*) but y_2^* is upwards distorted by the arguments given the proof of Propositions 1 (ii) and 2 (ii). Moreover, the downward IC of ω_4 workers may be binding or not and, correspondingly, y_3^* may be undistorted or downwards distorted.

For any economy with any $n > 4$ and $k \in \{2, \ldots, n-1\}$, the complete proof follows from a repeated application of the arguments given in cases 1 to 3. □

Proof of Proposition 4

Proof. In the following, I derive conditions under which some decreasing sequence of welfare weights $(\alpha_1, \ldots, \alpha_n)$ satisfies the conditions in Proposition 3, i.e., upward distortions at both margins are optimal. While this result generalizes Proposition 1, the formal arguments in the following proof differ substantially.

Specifically, the proof shows that a sequence α of welfare weights can at the same time (a) satisfy conditions (i) and (ii) in Proposition 3, (b) be weakly decreasing such that $\alpha_{j+1} \leq \alpha_j$, and (c) imply an average weight of $\sum_{j=1}^{n} f_j \alpha_j = 1$. The strategy of the proof is to construct a candidate sequence $q = (q_1, \ldots, q_{n-1})$ and to derive conditions under which q satisfies (a) to (c). For this purpose, fix $\bar{q} > 1$ to be given by a number such that $\tilde{\delta}_j$ is defined by the first-order condition (B.9) for any $j \in J$ and $\alpha_j \in [0, \bar{q}]$. Now, consider the sequence q such that $q_j = \bar{q} > 1$ for all $j \in \{1, \dots, k\}$ and $q_{j+1} = \max\left\{\beta_{Dj}(q_j), \underline{q}\right\}$ for all $j \in \{k, \dots, n-1\}$, where $\underline{q} := \max_{j \geq k} \underline{\beta}_j < 1$.

Property (a): Conditions in Proposition 3. First, by Lemma C.1 (i), there is for each $j \in J$ a number $\bar{\beta}_j > 1$ such that $\beta_{Uj}(q_j) < q_j$ if $q_j > \bar{\beta}_j$. As shown in the proof of Lemma C.1 (i), $\bar{\beta}_j < \bar{q}$ if $a = \omega_{j+1}/\omega_j$ is close enough to 1. Then, $q_{j+1} > \beta_{Uj}(q_j)$ for all $j \leq k-1$. Second, by construction, we have $q_{j+1} \geq \beta_{Dj}(q_j)$ for all $j \geq k$ such that $q_{j+1} > \underline{q}$. Finally, by C.1 (ii), $\underline{q} = \max_{j \geq k} \underline{\beta}_j$ ensures that $q_{j+1} \geq \beta_{Dj}(q_j)$ for all $j \geq k$ such that $q_{j+1} = \underline{q}$.

Property (b): Decreasing weights. First, by construction, $q_{j+1} \leq q$ for all $j \leq k-1$. Second, by Lemma C.1, $\underline{q} = \max_{j \geq k} \underline{\beta}_j$ implies that $q_{j+1} = \max\left\{\underline{q}, \beta_{Dj}(q_j)\right\} \leq q_j$ for all $j \in J$ (with a strict inequality if $q_j > \underline{\beta}_j$).

Property (c): Average weight of 1. First, with decreasing weights, $\sum_{j=1}^n f_j q_j = 1$ and $q_1 > 1$ requires that $q_n < 1$. As $q_j = \bar{q} > 1$ for all $j \leq k$, this obviously requires that $n \geq k+1$. In the following, I show that there is a finite number $m \geq k+1$ such that $q_n < 1$ whenever $n \geq m$. To see why, assume that there is a lower bound $\varepsilon > 0$ such that $q_j - q_{j+1} > \varepsilon$ for each $j \geq k$ such that $\alpha_j \geq 1$. If this is true, then $\alpha_j < \min\{1, \bar{q} - (j-k)\varepsilon\}$ for each $j \geq k+1$. Consequently, there is a unique natural number $m \leq k + (\bar{q}-1)/\varepsilon$ such that $\alpha_j < 1$ is ensured for all $j \in [m, n]$.

I now show that $q_j - q_{j+1}$ is indeed smaller than some number $\varepsilon > 0$ if $q_j \geq 1$, $\omega_{j+1}/\omega_j = a > 1$ and Assumption 3 holds. For this purpose, recall that $\tilde{\delta}_{j+1}$ is an increasing function of α_{j+1} and that, with $q_{j+1} = q_j \geq 1$, $\tilde{\delta}_{j+1} - \tilde{\delta}_j$ would be weakly larger than $\tilde{y}_{j+1} - h(\tilde{y}_{j+1}, \omega_{j+1}) - [\tilde{y}_j - h(\tilde{y}_j, \omega_j)]$. With $q_{j+1} = \beta_{Dj}(q_j)$, instead, the downward IC of ω_{j+1} workers is binding, $\tilde{\delta}_{j+1} - \tilde{\delta}_j = h(\tilde{y}_j, \omega_j) - h(\tilde{y}_j, \omega_{j+1})$. Hence, the welfare weights q_j and q_{j+1} satisfy

$$\int_{q_{j+1}}^{q_j} \frac{d\tilde{\delta}_{j+1}(\alpha)}{d\alpha_{j+1}} d\alpha \geq \tilde{y}_{j+1} - h(\tilde{y}_{j+1}, \omega_{j+1}) - [\tilde{y}_j - h(\tilde{y}_j, \omega_{j+1})] \ . \tag{C.8}$$

Define $\zeta_4 := \max_{j \geq k,\ \alpha_j \in [1, \bar{q}]} d\tilde{\delta}_j(\alpha_j)/d\alpha_j > 0$. Then, the left-hand side of (C.8) is weakly smaller than $(q_j - q_{j+1})\zeta_4$. Due to $h_{y\omega} < 0$ and $h_{yy} > 0$, in contrast, the right-hand side is larger than

$$\begin{aligned} \tilde{y}_{j+1} - h(\tilde{y}_{j+1}, \omega_{j+1}) - [\tilde{y}(z) - h(\tilde{y}(z), z)] &= \int_{\tilde{y}_j}^{\tilde{y}(z)} [1 - h_y(\tilde{y}(\omega), \omega_{j+1})]\, d\omega \\ &> (\tilde{y}(z) - \tilde{y}_j)\ [1 - h_y(\tilde{y}(z), \omega_{j+1})] \end{aligned}$$

with $z \in (\omega_j, \omega_{j+1})$ and $\tilde{y}(z) = \arg\max y - h(y, z)$. Assumption 3 (i) implies that

$$\tilde{y}(z) - \tilde{y}_j > \int_{\omega_j}^{z} \nu_1 \frac{\tilde{y}(\omega)}{d\omega} d\omega > \nu_1 \tilde{y}_j \ln\left(\frac{z}{\omega_j}\right) \ ,$$

while Assumption 3 (ii) implies that

$$1 - h_y(\tilde{y}(z), \omega_{j+1}) = \int_{z}^{\omega_{j+1}} h_{yy}(\tilde{y}(\omega), \omega_{j+1}) \frac{\tilde{y}(\omega)}{d\omega} d\omega > \int_{z}^{\omega_{j+1}} \frac{\nu_1}{\nu_2} \frac{h_y(\tilde{y}(\omega), \omega_{j+1})}{\omega} d\omega$$

$$> \frac{\nu_1}{\nu_2} h_y(\tilde{y}(z), \omega_{j+1}) \ln\left(\frac{\omega_{j+1}}{z}\right) .$$

For $z = \omega_j a^{1/2}$, the right-hand side of (C.8) is hence larger than

$$\zeta_5 := \frac{\nu_1^2}{\nu_2} (1/2 \ln a)^2 \min_{j \geq k} \tilde{y}_j h_y(\tilde{y}_j, \omega_{j+1}) > 0 .$$

Substituting all terms into (C.8), it follows that $q_j - q_{j+1} > \zeta_5/\zeta_4$ and, hence, bound away from 0 for all $j \geq 0$. As argued above, this ensures that there is a number $m \leq k + (\bar{q} - 1)\zeta_4/\zeta_5$ such that $q_j < 1$ for all $j \in [m, n]$.

Finally, the average welfare weight is exactly equal to 1 if

$$\begin{aligned}
\sum_{j=1}^{n} f_j q_j &= \sum_{j=1}^{m-1} f_j E\left[q_j \mid j < m\right] + \sum_{j=m}^{n} f_j E\left[q_j \mid j \geq m\right] = 1 \\
\Leftrightarrow \sum_{j=m}^{n} f_j &= \bar{z} := \frac{E\left[q_j \mid j < m\right] - 1}{E\left[q_j \mid j < m\right] - E\left[q_j \mid j \geq m\right]}
\end{aligned}$$

For $\sum_{j=m}^{n} f_j > \bar{z}$, there exists another (similar) weight sequence q' with average value 1 that satisfies the conditions in Proposition 3 and that is strictly decreasing with $q'_j > q'_{j+1}$ for all $j \in \{1, \ldots, n-1\}$. □

D Supplementary material

Appendix D provides additional results and numerical analyses. Additionally, it reviews previous results in the optimal tax literature, summarizes empirical estimates of labor supply elasticities and discusses the arguments for and formalization of poverty alleviation objectives.

D.1 Limit result

In the previous subsections, I have provided necessary and sufficient conditions for the optimality of an EITC in models with a discrete skill set. A natural question is whether the results also extend to a model with a continuum of skill types. The tax reform analysis in Section 4 shows that, in the continuous model, the introduction of a small EITC with negative marginal and participation taxes at the bottom increases welfare if the desire for redistribution among the poor is sufficiently small. However, this analysis did not clarify whether the optimal income tax involves negative marginal and participation taxes as well; there might be alternative tax reforms that allow to increase welfare even more. Section 5 provides sufficient conditions for the optimality of an EITC in the discrete model, but the formal proofs exploit the discreteness of the skill set and, hence, cannot be transferred to the continuous model.

I can tackle this question, however, by studying the results of the model for the limit case where a discrete skill $\{\omega_1, \omega_2, \ldots, \omega_n\}$ converges to the set of all rational numbers in the interval $[\omega_1, \omega_n]$. For this purpose, I first impose the same functional-form assumptions on the utility function and the type distribution as in the stylized model studied in Subsection 5.2. Hence, I assume that the effort cost function h is isoelastic according to (15) and that fixed costs are uniformly distributed on an interval $[0, \bar{\delta}]$ in all skill groups. Second, I consider an economy with a large finite number n of skill groups and constant relative distances between adjacent skill types, $\omega_{j+1}/\omega_j = a > 1$ for all $j \in J_{-n}$, as in Subsection 5.3. The highest skill type follows as $\omega_n = \omega_1 a^{n-1}$. I denote by $F(\omega')$ the share of agents with skill types in the interval $[\omega_1, \omega']$. An economy in this class is characterized by a collection $(n, a, F, \sigma, \omega_1, \bar{\delta})$. In the following, I let the relative distance a shrink to 1 and the number n of skill groups go to infinity, while keeping the smallest and the largest skill type fixed. Hence, I consider the limit case where the skill set converges to an infinite number of evenly spaced skill types in the interval $[\omega_1, \omega_n]$. The following proposition demonstrates that the optimal marginal tax can be negative on a substantial income range even in this limit case.

Proposition D.1. *Fix $\sigma > 0$, $\bar{\delta}$, ω_1 and ω_n. Consider three skill types ω_a, ω_b and ω_c in (ω_1, ω_n) such that $\tilde{y}(\omega_b) = 2\tilde{y}(\omega_a)$ and $\tilde{y}(\omega_c) = 4\tilde{y}(\omega_a)$. Let a and n converge to 1 and to ∞, respectively, in such a way that $\omega_n = \omega_1 a^{n-1}$ stays constant. In the limit, there are strictly decreasing welfare weights $\alpha : \omega \mapsto \alpha(\omega)$ such that optimal labor supply is upwards distorted at both margins for each skill group $\omega \in (\omega_1, \omega_a)$ if*

$$\frac{1 - F(\omega_c)}{F(\omega_b)} \geq 2 \; . \tag{D.1}$$

The formal proof is available on request. Proposition D.1 clarifies that an EITC with negative marginal taxes can be optimal even if the limit case where the discrete skill set converges to a continuum. More specifically, it shows that a tax with strictly negative marginal taxes for all incomes below some threshold $y(\omega_a)$ is optimal for some strictly decreasing welfare weights α if condition (D.1) is satisfied. As in Proposition 1 for the stylized model with three skill types and in Proposition 4 for the general discrete model, the sufficient condition is expressed in terms of the skill distribution; it requires that there exists a sufficiently large share of agents that are sufficiently more productive than those workers with incomes in the phase-in range. In particular, this condition compares, first, the shares of high-skill agents with conditionally optimal incomes above $\tilde{y}(\omega_c) = 4\tilde{y}(\omega_a)$ and, second, the share of low-skill agents with conditionally optimal incomes below $\tilde{y}(\omega_b) = 2\tilde{y}(\omega_a)$. Intuitively, the latter subset includes both the agents in the phase-in range and in the phase-out range of the EITC. An EITC with negative marginal taxes up to threshold $\tilde{y}(\omega_a)$ can be rationalized whenever the share $1 - F(\omega_c)$ of high-skill workers is at least twice as large as the share $F(\omega_b)$ of low-skill workers. Remarkably, this condition does not depend on the level of the intensive-margin elasticity $\sigma > 0$.

Implicitly, Proposition D.1 also provides a lower bound on the size of the phase-in range that can be rationalized with decreasing welfare weights. For a given skill distribution F, I can determine the highest skill level ω_a that still satisfies condition (D.1).[40] As an illustration, consider the subgroup of childless singles in the US in 2015, which is also considered in the numerical simulations in Section 6. I can estimate the skill distribution in this subgroup using data from the March 2016 CPS. Based on this data, Proposition D.1 implies that an EITC with negative marginal taxes up to an annual income of \$12,600 and negative participation taxes up to an annual income of \$25,200 can be rationalized. By contrast, the 2015 US EITC only implied negative marginal taxes for incomes below \$6,580 and negative participation taxes for incomes below \$14,820.

D.2 Construction of weight sequences for simulations

In my benchmark simulation in Section 6, I use the sequence α^A of welfare weights, which assigns the welfare weight of $\alpha_j^A = 1.04$ to all skill types with a conditionally optimal income below \$11,770 (the US 2015 poverty line for childless singles) under laissez-faire. This condition is satisfied for the 49 lowest skill groups, constituting the bottom 5.6% of the skill distribution according to CPS data. For each skill group $j \in \{49, \ldots, 74\}$, I set the weight α_j^A to the lowest value such that the first-and-half-best allocation still satisfies the downward IC of workers with skill type ω_j. Formally, this implies that $\alpha_j^A = \beta_{j-1}^D(\alpha_{j-1}^A)$. The welfare weight of each skill group $j \geq 75$ is set to a value $\underline{\alpha}^A$ such that the average welfare weight equals 1. The implied weight sequence α_j is monotonically decreasing over the skill distribution. In Figure D.1a, sequence α^A is depicted by the blue line.

In the sensitivity analysis, I additionally consider two sequences of welfare weights that – just as sequence α^A – assign a welfare weight of 1.04 to the lowest-skilled workers. In contrast to the benchmark weights, however, these alternative welfare weights are not gradually decreasing, but drop sharply for higher skill typs from above-average values to below-average values. As a result, these welfare weights satisfy the necessary condition for an EITC with negative marginal taxes at the bottom, but violate the sufficient conditions established in Proposition 3. Specifically, sequence α^B of welfare weights again assigns the welfare weight 1.04 to the agents in the lowest

[40]Note that both ω_b and ω_c are increasing with ω_a by construction. Hence, the left-hand side of condition (D.1) is strictly decreasing in ω_a. This implies that the maximum of skill levels ω_a that satisfy condition (D.1) is well-defined.

49 skill groups, i.e., with conditionally optimal incomes below the poverty line. Then, the welfare weights drop to attain a value of $\underline{\alpha}^B \approx 0.9976$ for all higher-skilled agents. This ensures an average welfare weight of 1 in the population. In Figure D.1a, sequence α^B is depicted by the brown line.

Similarly, sequence α^D drops sharply from above-average welfare weights to below-average weights at some skill threshold. Hence, these welfare weights also violate the sufficient conditions for an EITC from Proposition 3. In contrast to the previous sequences, α^D assigns the welfare weight of 1.04 to a larger set of agents: to all those with conditionally optimal incomes below \$28, 500 under laissez-faire. This includes approximately the lowest-skilled quartile of the population according the CPS data (26.3%). To ensure an average welfare weight of 1, the welfare weight of all higher-skilled agents is set to a value of $\underline{\alpha}^D \approx 0.986$. In Figure D.1a, sequence α^D is depicted by the green line.

Finally, the sensitivity analysis also considers a sequence α^C of welfare weights that are convexly decreasing over the income distribution. Hence, the skill gradient of these welfare weights is particularly large (negative) at the bottom, representing a social planner with a particularly strong desire for redistribution from the rich to the poor. In particular, I assume that the weight of the workers in skill group j is given by $\alpha_j^C = 3/4 + 0.3\ (\omega_1/\omega_j)^{1/3}$, while the average welfare weight of non-working agents is given by approximately 1.46. A graphical illustration of sequence α^C is given in Figure D.5a.

D.3 Sensitivity analysis

In this section, I provide an extensive sensitivity analysis for the numerical simulations reported in Section 6. I start by investigating how the optimal income tax is affected by variations in the skill gradient of welfare weights for medium- and high-skilled workers. Figure D.1a depicts three alternative weight sequences, all of which are flat at the bottom: the benchmark welfare weights α^A and the alternative welfare weights α^B and α^C defined in Subsection D.2 above. Figure D.1b depict the optimal participation taxes for all three welfare weight sequences, given the benchmark levels of the intensive-margin elasticity ($\varepsilon = 0.3$) and the participation elasticities (average value $\bar{\pi} = 0, 25$).

Figure D.1b shows that the optimal income tax for each of the welfare weights α^A, α^B, and α^D involves an EITC with negative marginal taxes and negative participation taxes for low-income workers. This was not clear a priori, because both the welfare weights α^B and α^D violate the sufficient conditions for an EITC established in Proposition 3. Quantitatively, the optimal EITC for sequence α^B is smaller than for the benchmark sequence α^A, regarding both the length of the phase-in and phase-out ranges and the maximum tax credit. In contrast, the optimal EITC for sequence α^D involves longer phase-in and phase-out ranges and a larger maximum credit.

The following figures illustrate the effects of variations in the labor supply elasticities at both margins. To get a broader picture, each figure shows the effects on the simulated participation taxes $T^P(y) = T(y) - T(0)$ that are optimal given welfare weights α^A in the left panel, and the sequence of welfare weights α^D in the right panel.

Figure D.2 shows how variations in the intensive-margin elasticity affect the optimal tax schedule for both weight sequences. Intuitively, a higher intensive-margin elasticity implies

Figure D.1: Optimal taxes for different welfare weights.

(a) Welfare weight sequences. (b) Optimal tax schedules.

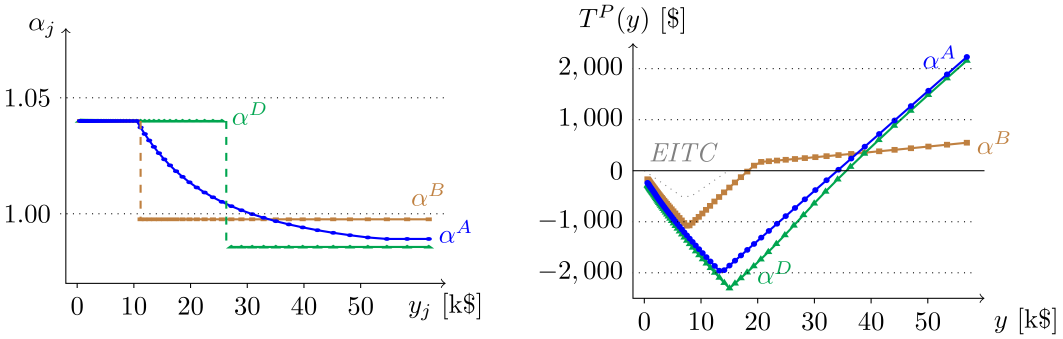

Notes: Figure D.1a in the left panel plots three alternative sequences of welfare weights: the benchmark welfare weights α^A (blue line), sequence α^B that assigns above-average weights to the bottom 5.6% of the skill distribution (brown), and sequence α^C that assigns above-average weights to the bottom 26.3% (green). Figure D.1b in the right panel plots the optimal participation taxes for the welfare weight sequences α^A (blue line), α^B (brown) and α^C (green).

that the first-and-half-best tax schedule would induce larger distortions at the intensive margin. Hence, the planner is willing to deviate further from this "target function" in order to implement the optimal compromise between distortions at both margins. My simulation results confirm this intuition. In both panels, the benchmark simulations with an elasticity of $\varepsilon = 0.3$ are depicted by the blue lines. The green lines with triangles depict the participation taxes that are optimal for a smaller elasticity of $\varepsilon = 0.1$. The red lines with squares depict the optimal participation taxes for a larger elasticity $\varepsilon = 0.5$. The larger the intensive-margin elasticity, however, the flatter are the tax functions, both in the phase-in and in the phase-out ranges. For a higher elasticity, this implies both smaller marginal subsidies and a smaller maximum credit. Besides, a higher elasticity at the intensive margin seems to be related to a somewhat shorter phase-in range. The qualitative properties of the optimal tax remain constant, however: In all cases, it is given by an EITC with negative marginal taxes for incomes up to at least $12,500 and negative participation taxes for incomes up to at least $30,000 (see Figure D.2).

Figure D.2: Optimal taxes for different intensive-margin elasticities.

(a) Optimal tax for weight sequence α_A (b) Optimal tax for weight sequence α_D.

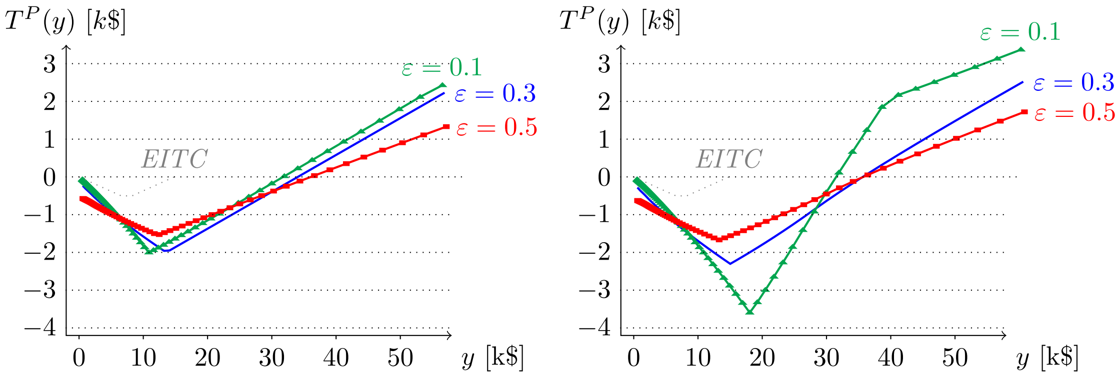

Notes: Figure D.2 depicts the optimal participation tax schedules for welfare weights α_A (left panel) and for welfare weights α_D (right panel), given three different levels of the intensive-margin elasticity: the benchmark level $\varepsilon = 0.3$ (blue lines), a lower elasticity of $\varepsilon = 0.1$ (green lines with triangles), and a higher elasticity of $\varepsilon = 0.5$ (red lines with squares).

Figure D.3 shows how variations in the participation elasticities affect the optimal tax schedule for both weight sequences. Intuitively, a decrease in the participation elasticity implies that a given participation tax/subsidy creates less distortions at the extensive margin. On the one hand, this implies that the social planner focuses more on reducing intensive-margin distortions, which makes an EITC less attractive. On the other hand, the first-and-half-best participation taxes become more negative at the bottom and more positive at the top. Based on my simulations, however, the effects on the second-best tax schedules are quite limited. In both panels, the blue lines depict the benchmark simulations for participation elasticities that are falling from 0.4 to 0.18 along the skill distribution, giving an average elasticity of $\bar{\pi} = 0.25$. The green lines with triangles depict the optimal participation taxes for smaller participation elasticities that are falling from 0.16 to 0.1, yielding an average elasticity of $\bar{\pi} = 0.12$ (in line with Kleven 2020*a*). The red lines with squares depict the optimal taxes for larger elasticities that are falling from 0.5 to 0.4 as in the benchmark calibration by Jacquet et al. (2013), giving an average elasticity of $\bar{\pi} = 0.43$. In the last case, I also set the intensive-margin elasticity to 0.25 to be consistent with Jacquet et al. (2013). Both figures show that a variations in the average level and in the skill gradient of participation elasticities have only minor effects on the properties of the optimal income tax. Unreported simulations show, moreover, that the results are almost unchanged for a specification with a participation elasticity equal to 0.25 in all skill groups.

Figure D.3: Optimal taxes for different participation elasticities.

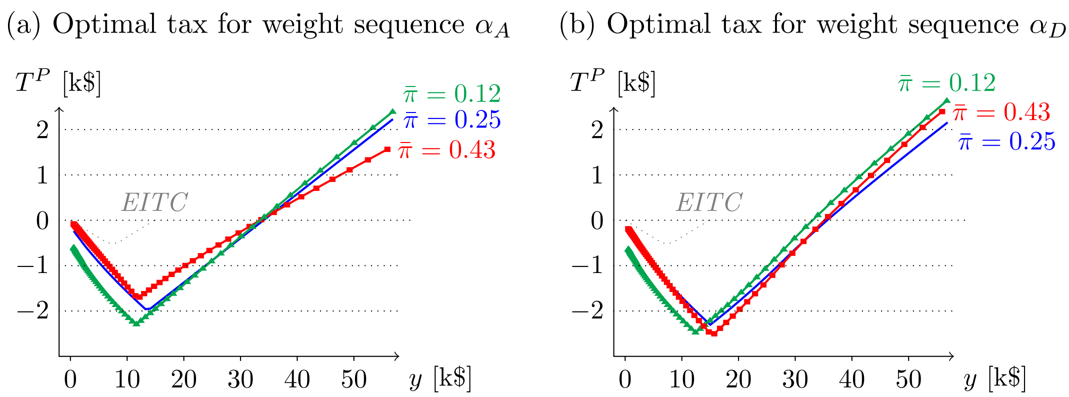

Notes: Figure D.3 depicts the optimal participation tax schedules for welfare weights α_A (left panel) and for welfare weights α_D (right panel), given the benchmark case with an average participation elasticity of $\bar{\pi} = 0.25$ (blue lines), a lower participation elasticity of $\bar{\pi} = 0.12$ (green lines with triangles), and a higher elasticity of 0.43 (red lines with squares).

Figure D.4 shows how the discretization of the skill set affects the optimal tax schedule for both weight sequences. In both panels, the benchmark simulations for a skill set with $n = 96$ skill groups are depicted by blue lines. The relative distance between each pair of adjacent skill types is equal to $\omega_{j+1}/\omega_j = 1.05$. To study the effects of the discretization, I consider two alternative calibrations with skill sets such that the lowest and the highest skill types are identical to the benchmark calibration, but the number of skill types and the relative distance between adjacent types are varied. For each discretization, I have re-estimated the skill distribution based on March 2016 CPS data. In particular, the green lines with triangles depict the participation taxes that are optimal for a coarser skill set with 48 skill types. The red lines with squares depict the optimal taxes for a finer skill set with 192 skill types. Both figures show that a finer discretization of the skill set leads to flatter tax schedules both in the phase-in and the phase-out income ranges. At the same time, there are only small effects on the length of the phase-in range.

Figure D.4: Optimal taxes for different discretizations of the skill set.

(a) Optimal taxes for weight sequence α_A (b) Optimal taxes for weight sequence α_D

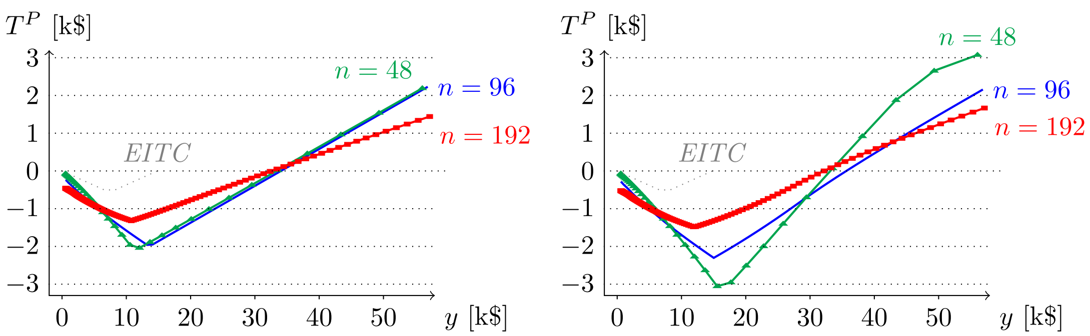

Notes: Figure D.4 depicts the optimal participation tax schedules for welfare weights α_A (left panel) and welfare weights α_D (right panel), given three alternative discretizations of the skill set: a case with 48 skill types (green lines with triangles), the benchmark case with 96 skill types (blue lines), and a case with 192 skill types (red lines with squares).

The reason for these effects is that, with a larger number of less distant skill types, there is a larger set of IC constraints along the skill dimension, which also become more restrictive.[41] In the cases considered in Figure D.4, the inverse elasticity rule (18) maximizes extensive-margin efficiency, but conflicts with incentive compatibility along the skill dimension. To implement the inverse elasticity rule, the upward ICs of low-skill workers must be slackened, which requires to distort labor supply upwards at the intensive margin. The finer the skill set, the more the upward IC constraints must be slackened by means of larger upwards distortions at the intensive margin. Thus, the tradeoff between intensive-margin efficiency and extensive-margin efficiency becomes more severe. As a result, the optimal tax schedule becomes flatter because it deviates further from the (unchanged) first-and-half-best tax. A further increase in the number of skill groups would lead to a further flattening of the optimal tax schedule, but would continue to involve strictly negative marginal and participation taxes at the bottom.

Finally, I investigate how the optimal tax changes if the welfare weights are convexly decreasing along the skill distribution as in the calibrations of Saez (2002) and Jacquet et al. (2013). In this case, the skill gradient is particularly large at the bottom, representing a social planner with a strong desire for redistribution between the poor and the very poor. Figure D.5a depicts the sequence α^C of welfare weights, defined in Subsection D.2 above. Figure D.5b plots the optimal income tax for welfare weights α^C that results from my numerical simulations. In contrast to all previous figures, both participation taxes and marginal taxes are strictly positive at all income levels below the very top.

[41] To see this, fix the bundles (c_j, y_j) and (c_{j+1}, y_{j+1}) of the workers with skill types ω_j and ω_{j+1}, respectively. Assume that, given these bundles, the upward IC constraint of ω_j workers with respect to the bundle of ω_{j+1} workers is binding such that $c_j - h(y_j, \omega_j) = c_{j+1} - h(y_{j+1}, \omega_j)$. Now, I introduce another group of workers with an intermediate skill $\omega_{j+1/2} \in (\omega_j, \omega_{j+1})$ and some bundle of income $y_{j+1/2} \in (y_j, y_{j+1})$ and consumption $c_{j+1/2} \in (c_j, c_{j+1})$. It can be shown that, for any such bundle $(y_{j+1/2}, c_{j+1/2})$, either the upward IC of ω_j workers with respect to the bundle of $\omega_{j+1/2}$ workers, or the upward IC of $\omega_{j+1/2}$ workers with respect to the bundle of ω_{j+1} workers is violated.

Figure D.5: An example with strictly positive marginal taxes.

(a) Welfare weights α^C (b) Optimal tax schedule for sequence α^C

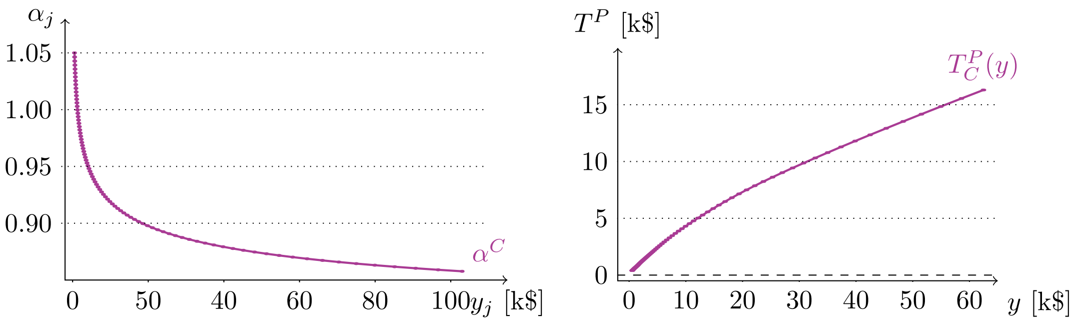

Notes: Figure D.5a in the left panel plots for each skill group j the welfare weight α_j^C against the skill-specific gross income level y_j under the current US tax system. Figure D.5b in the right panel depicts the optimal participation tax $T_C^P(y)$ for welfare weights α_C and benchmark values of the intensive-margin elasticity ($\varepsilon = 0.3$) and the participation elasticity (average value 0.25). Each dot represents the tuple $(y_j, T_C^P(y_j))$ for the workers in one skill group.

D.4 Optimality of positive marginal taxes

Above, I have shown numerically that the results of Saez (2002) and Jacquet et al. (2013) are recovered in my model if the skill gradient of welfare weights at the bottom is large. To complement this numerical result, I provide an analytical condition for the optimality of strictly positive marginal taxes.

Proposition D.2. *Let Assumptions 4, 5 and 3 be satisfied. Optimal output is downwards distorted at the intensive margin in all skill groups $j \in J_{-n}$ if $\alpha_{j+1} \leq \beta_{Dj}(\alpha_j)$ for all $j \in J_{-n}$ with at least one strict inequality, where function β_{Dj} is defined in Lemma C.1.*

Proposition D.2 extends the main result of Jacquet et al. (2013) to a model with a discrete instead of a continuous skill set. A formal proof is available on request. Basically, the proposition says that optimal labor supply is downwards distorted at the intensive margin if the first-and-half-best allocation violates the downward IC constraints of the workers in all skill groups $\{2, \ldots, n\}$. The proof shows that, in this case, every single downward IC constraint is unambiguously binding in the optimal allocation. By Lemma C.1, the condition $\alpha_{j+1} \leq \beta_{Dj}(\alpha_j)$ is satisfied if either the difference between the welfare weights α_j and α_{j+1} is large enough, or both welfare weights are sufficiently much below the average weight of 1 (see also Figure 3).

D.5 Optimal income taxes for single parents

The following section provides the results of numerical simulations for a second subgroup of the US population, single parents. To calibrate the model for this population group, I consider the same functional forms as for childless singles: (a) utility (1) is quasilinear, (b) the effort cost function (15) gives rise to a constant elasticity at the intensive margin, (c) the distributions of fixed costs (19) are logistic in each skill group, (d) the participation elasticity π_j and the participation share π_j vary along the skill distribution according to equations (20) and (21), respectively. The last point implies that parameters μ, $\underline{\pi}$ and $\bar{\pi}$ have to be set in order to match empirically plausible labor supply elasticities for single parents.

For this purpose, I again follow the survey by Saez et al. (2012), the meta-study by Chetty, Guren, Manoli & Weber (2013) and the study by Chetty, Friedman & Saez (2013) on EITC recipients. With respect to the intensive margin, there is no convincing evidence for differences between single parents and other population groups. Hence, I again set the intensive-margin parameter to $\sigma = 0.3$.

With respect to the extensive margin, in contrast, the literature suggests substantial differences across population groups. In particular, single mothers are commonly perceived as the group that is most responsive at the participation margin (along with married women). Early studies estimate extremely large elasticities for single mothers receiving the EITC, sometimes even above 1 (see, e.g., Meyer & Rosenbaum 2001). Chetty, Guren, Manoli & Weber (2013) review a number of studies based on quasi-experimental variation and re-calculate the elasticities to be consistent with each other. The resulting elasticities in those studies that focus on single mothers are in the range of 0.43, which is still larger than for other population groups. The two most recent studies exploiting EITC reforms point towards smaller levels: Bastian & Jones (2020) estimate a participation elasticity of 0.33 for single mothers, while Kleven (2020*a*) finds an elasticity closer to, and not significantly different from 0.

Based on this literature, I set the parameters $\underline{\pi}$ equal to 0.6 and $\bar{\pi}$ equal to 0.3. This gives an average participation elasticity of 0.42, close to the suggestion of Chetty, Guren, Manoli & Weber (2013). In particular, this also implies that elasticities are falling along the skill distribution, in line with the existing evidence for all subgroups of the population. As the participation rate of single parents is 82.1% in the CPS data (similar to the one of childless singles), I set $\underline{L} = 0.7$ and $\bar{L} = 0.85$.

To calibrate the skill distribution of single parents, I again use data from the March 2016 CPS. For this purpose, I consider only single parents at ages 25 to 60 that do neither live with an unmarried spouse nor with any family member except for own children below the age of 18. This ensures that there is no joint labor supply decision problem. The restricted sample contains $7,141$ observations. In the CPS data, single parents are on average less productive than childless singles (average annual income \$32,958 compared to \$42,223 for childless singles). To back out the skill distribution, I again use the same linear approximation of the US tax system.[42] I can then use the first-order condition for an individual optimum to assign skill levels to all single parents with positive incomes. In particular, I use the same skill set with $n = 96$ skill types and estimate the share of single parents in each skill group based on a kernel density estimation.

To calibrate the planner's redistributive preferences, I cannot simply use the same weight sequences as for childless singles. As the skill distribution varies across subgroups, the average weight among single parents would differ from 1. Hence, I again construct two exogenous sequences of welfare weights that are monotonically decreasing and have an average value of 1 (within the group of single parents). First, I consider a sequence $\tilde{\alpha}^A$ that assigns a welfare weight

[42]For single parents, this approximation is less convincing than for childless singles, especially at low income levels where the effective marginal tax depends on the statutory tax schedule, the EITC phase-in and phase-out rates and the marginal reduction rates of welfare transfers. As shown by Maag et al. (2012), the effective marginal tax varies strongly across states and income levels. For example, if a single parents moves from zero income to the federal poverty line, he faces an average marginal tax between a minimum of -13.3% and a maximum of 25.5%. Between the poverty level and twice the poverty level, the average marginal tax is above 40% in most states. Given the lack of superior alternatives, I stick to the same linear approximation as for childless singles.

of 1.04 to the agents in the lowest 45 skill groups, i.e., to all working agents with incomes below \$9,050. The welfare weights of higher-skilled agents are assumed to be gradually decreasing (as with sequence α^A for childless singles). Second, I consider a sequence $\tilde{\alpha}_B$ that assigns a welfare weight of 1.03 to the lowest 66 skill groups. These agents constitute the lower-skilled half of the single parent subpopulation (more precisely: 52.3%) and include all working agents with incomes below \$31,000. The average weight of the higher-skilled agents is set to be constant as well and equal to 0.967. As in the case of childless singles, welfare weights $\tilde{\alpha}_A$ satisfy the sufficient conditions for an EITC with negative marginal taxes, while welfare weights $\tilde{\alpha}_B$ only satisfy the necessary conditions.

Figure D.6: Optimal income taxes for single parents.

(a) Optimal tax for weight sequence $\tilde{\alpha}_A$ (b) Optimal tax for weight sequence $\tilde{\alpha}_B$

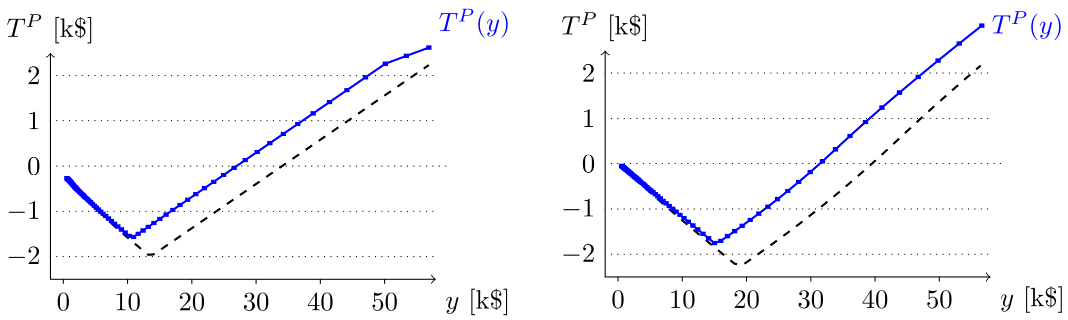

Notes: Figure D.6 depicts the optimal participation taxes for single parents given welfare weights $\tilde{\alpha}^A$ (left panel) and welfare weights $\tilde{\alpha}^B$ (right panel). In each panel, the optimal participation tax for single parents is depicted by the solid blue line. Each square represents the tuple $(y_j, T^P(y_j))$ for one of the 96 skill groups. For comparison, the optimal participation tax for childless singles is depicted by the dashed black line in each panel.

The simulation results are depicted in Figures D.6a for sequence $\tilde{\alpha}^A$ and D.6b for sequence $\tilde{\alpha}^B$. In particular, the solid blue lines in both figures show the optimal participation taxes for single parents. For welfare weights $\tilde{\alpha}^A$, the optimal tax involves an EITC with negative marginal taxes for incomes up to \$10,935, a maximum tax credit of \$1,566 and negative participation taxes for incomes up to \$26,574. For welfare weights $\tilde{\alpha}^B$, the optimal tax involves an EITC with negative marginal taxes for incomes up to \$15,017, a maximum tax credit of \$1,753 and negative participation taxes for incomes up to \$29,912. The transfers to non-working agents only reach levels of \$1,231 and \$1,573, respectively. However, these transfers have to interpreted as transfers paid by higher-skilled single parents only. If the social planner has a desire for redistribution from other population groups to the group of single parents, then there is an additional lump-sum transfer towards all single parents. With quasi-linear utility, specifically, these between-subgroup transfers only imply parallel shifts in the budget sets; they do not affect the optimal levels of marginal taxes and participation taxes.

These simulated tax schedules have to be interpreted as the optimal integrated schemes that result from the combination of income taxes and welfare transfers (and, potentially, child care costs). As argued above, the integrated tax-transfer systems effectively imply a larger "net EITC" in some US states and a smaller one in other states (see Maag et al. 2012). Hence, my simulation results can be broadly interpreted as providing support for the current levels of the EITC for single parents. They do not provide support for another substantial extension of the

EITC, however.

For comparison, Figures D.6a and D.6b also depict by the black dashed lines the optimal tax schedules that would result for childless singles given similar welfare sequences. According to these simulation results, the optimal EITC should be larger for childless singles than for single parents. This is probably in contrast to the current US policies, although single parents in the US do not only benefit from a more generous EITC but also face larger positive reduction rates from a number of welfare transfers such as TANF, SNAP, Public Housing etc. The main reason for the different simulation results is that the share of high-skilled agents is much larger in the group of childless singles than in the group of single parents. As explained in Subsection 5.2, this implies that a generous EITC can more easily be financed in the group of childless singles without imposing strong distortions at the intensive margin. In contrast, the larger participation elasticities of single parents seem to have only a negligible effect on the optimal size of the subgroup EITC.

D.6 Generalized welfare weights and poverty alleviation

In this section, I discuss whether an EITC is the optimal policy for alleviating poverty. Following Mirrlees (1971), optimal income taxation is commonly studied under the assumption that the social objective is given by a utilitarian welfare function (see Weinzierl 2014 for a recent critique). As an alternative, Kanbur et al. (1994) as well as Besley & Coate (1992, 1995) suggest to use the goal of alleviating poverty as measured by the available income. They advocate this objective as being more consonant with public debates and, consequently, as providing better insights into real-world policy choices. In particular, they argue that neither policy-makers nor taxpayers seem to value the leisure enjoyed by the poor (which is an argument of standard utility functions), but rather seem to focus on income as a more visible sign of poverty (see, e.g., Besley & Coate 1995, p. 189, and Kanbur et al. 1994, pp. 1615-1616).

Support for this view comes from the recent public debate surrounding a potential expansion of the EITC for childless workers. Most prominently, President Barrack Obama and Paul Ryan, then Republican Chairman of the House of Representative Budget Committee, independently proposed to expand the EITC by doubling the phase-in and phase-out rates, raising the phase-out start and the eligibility threshold, and relaxing age restrictions for childless workers (Executive Office 2014, House Budget Committee 2014). The Obama proposal emphasizes the goal to reduce poverty for childless low-income workers. In particular, the proposal estimates that "the increase in the credit would lift about half a million people above the poverty line and reduce the depth of poverty for 10 million more" (Executive Office 2014: 2). It also criticizes that the current US tax code pushes childless workers with low incomes "into or deeper into poverty" (Executive Office 2014: 3), both directly and indirectly through discouraging work.[43] The Ryan proposal suggests a number of reforms to reduce poverty and increase economic self-sufficiency. It argues that the EITC is the most successful program in fighting poverty among families, and that its expansion would significantly reduce poverty among childless workers. The Ryan proposals also emphasizes that an EITC expansion would provide larger incentives for people to

[43] Besides, the proposal argues that an EITC expansion would increase employment rates, and that this might benefit society through positive external effects such as increasing marriage rates, supporting child outcomes and reducing incarceration rates (Executive Office 2014: 9).

work and "earn enough money to place them above the poverty line" (House Budget Committee 2014: 7). A number of further proposals provide similar arguments for an even more generous EITC expansion, emphasizing the goals of lifting people above the poverty line, reducing the depth of poverty for others and, in particular, eliminating the possibility that low-income workers are taxed into poverty. For example, such poverty-related arguments were made to support two recent proposals for EITC expansions, introduced in the House of Representatives in February 2017, and in the Senate in June 2017.[44]

I continue by sketching the most common poverty measures and their formalization in optimal tax problems. Let $\bar{c}$ denote the poverty line, expressed in terms of consumption or available income (after tax and transfers).[45] Foster et al. (1984) introduce a class of poverty measures given by

$$P_a(c, \bar{c}) = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{\bar{c} - c_i}{\bar{c}} \right)^a \mathbb{1}_{c_i \leq \bar{c}} \; .$$

where n is the number of agents in the population and c_i is agent i's available income. For $a = 0$, this measure is equal to the poverty rate or head count ratio, the share of a population with available incomes below $\bar{c}$. For $a = 1$, it is equal to the poverty gap, the average (percentage) shortfall of available income from the poverty line. In contrast to the poverty rate, the poverty gap also accounts for the intensity (depth) of poverty. For $a > 1$, the measure assigns higher weights to larger shortfalls from the poverty line. Institutions such as the World Bank and the United Nations commonly use the first two measures, the poverty rate and the poverty gap.

Kanbur et al. (1994) were the first to study optimal income taxation under the goal of poverty alleviation instead of welfare maximization. They formalize the objective of poverty alleviation by using what they call a generalized poverty gap measure (with $a > 1$). For an intensive-margin model, they find that the optimal marginal tax might be negative at incomes below the poverty line. Numerical simulations show large positive marginal taxes to be optimal even at the very bottom, however. They also point out that their formalization of poverty alleviation is inconsistent with the Pareto principle, because it fails to account for the agents' disutility of providing output instead of enjoying leisure. To reconcile the poverty gap criterion with the Pareto principle, Saez & Stantcheva (2016) suggest to apply their approach of generalized welfare weights by setting welfare weights equal to $\alpha_p > 1$ for all agents with consumption below the poverty line $\bar{c}$, and equal to $\alpha_{np} \in [0, 1)$ for all agents with consumption above $\bar{c}$. In my numerical simulations, I have considered the sequence α_B of welfare weights, which has exactly this shape. As shown in Subsection 6.2, the optimal tax for these welfare weights is given by an EITC with negative marginal taxes and negative participation taxes in my model with labor supply responses at two margins. With intensive-margin responses alone, in contrast, the optimal marginal taxes for this criterion are strictly positive both below and above the poverty line $\bar{c}$ (Saez & Stantcheva 2016).

[44]For the details on these proposals, see *https://www.congress.gov/bill/115th-congress/house-bill/822* and *https://www.congress.gov/bill/115th-congress/senate-bill/1371*.

[45]Sometimes, the poverty line is also defined in terms of pre-tax income.

D.7 Perturbation approach and optimal tax formula

One common method to solve for the optimal income tax is the perturbation approach, introduced by Piketty (1997) and Saez (2001). Jacquet et al. (2013) and Lorenz & Sachs (2016) use this approach in a model with two margins and a continuum of skill (or income) levels. The analysis starts by fixing an initial tax schedule T and the income distribution F_y it implements. It then considers a perturbation of the tax system that increases the marginal tax T' by a small amount τ on a small interval $(y', y' + \ell)$. For the model introduced above, the welfare effect of a small reform with $\tau \to 0$ and $\ell \to 0$ is zero if and only if the marginal tax at income level y' satisfies

$$\frac{T'(y')}{1 - T'(y')} = \frac{1 - F_y(y')}{\varepsilon(y')\ y'\ f_y(y')} \left[1 - \hat{\alpha}(y', \bar{y}) - \frac{1}{1 - F_y(y')} \int_{y'}^{\bar{y}} f_y(y)\ \eta(y)\ T^P(y) dy\right] , \qquad \text{(D.2)}$$

where $\varepsilon(y)$ denotes the intensive-margin elasticity of workers with income y, $\eta(y)$ denotes the semi-elasticity of participation of workers with income y, and $\bar{\alpha}(y')$ is the average welfare weights associated to agents with incomes above y' under tax schedule T. If tax schedule T is optimal, it satisfies equation (D.2) at any income level $y' \in [\underline{y}, \bar{y}]$. Note that the formula does not account for income effects, which are ruled out in my model by the assumption of quasi-linear preferences. Jacquet et al. (2013) and Lorenz & Sachs (2016) provide generalized versions of this optimal tax formula that allow for income effects. Saez (2002) provides a discretized version of (D.2), where the right-hand side contains the sum $\sum_{y'=y}^{\bar{y}} f_y(y)\eta(y)T^P(y')$ over a finite number of income levels interpreted as occupations (instead of the integral term in D.2).

Equation (D.2) has a similar structure as the well-known ABC formula for the intensive model (Diamond 1998). The only difference is given by the integral term on the right-hand side, which accounts for the extensive-margin responses at income levels above y'. Due to this integral term, however, the optimal tax formula for the two-margin model is substantially less informative about the qualitative properties of the optimal tax schedule T than the ABC formula. To see this, note that the optimal marginal tax at income y' also depends on the values of the semi-elasticity η, the income density f_y and the optimal participation tax T^P at all income levels above y'. Most importantly, the optimal participation tax T^P is an endogenous object – more precisely, the object of interest in the optimal tax problem. Hence, we cannot determine T' at any income level without precise knowledge of the entire tax function. Formally, (D.2) represents a differential equation with an integral term, which in general cannot be solved analytically. This complication makes it even hard to determine the optimal sign of $T'(y)$. On the right-hand side of (D.2), all terms in front of the bracket and the difference $1 - \hat{\alpha}(y', \bar{y})$ are positive for any $y \in [\underline{y}, \bar{y}]$. However, we have to subtract the extensive-margin terms under the integral, which is also positive in the plausible case that $T^P(y) > 0$ for most high-skill workers. But then, the optimal sign of $T'(y)$ depends on whether the term $1 - \hat{\alpha}(y', \bar{y})$ or the integral term is larger, given the optimal participation tax $T^P(y)$ for all $y > y'$.[46]

[46] Without the integral term, by contrast, the optimal marginal tax can be determined point-wise at each income level: It only depends on the values of ε, $\hat{\alpha}$ and the hazard rate at income y'. In particular, the optimal marginal tax is unambiguously positive if welfare weights are monotonically decreasing with income (and skill) so that $\hat{\alpha}(y', \bar{y})$ is below 1 for any $y' \in [\underline{y}, \bar{y}]$.

Based on Saez (2002), some papers argue that a negative marginal tax is more likely to be optimal if labor supply responds *more strongly* at the extensive margin than at the intensive margin. Indeed, an increase in the semi-elasticity $\eta(y)$ decreases the right-hand side of (D.2) if $T^P(y)$ is positive and held constant. However, $T^P(y)$ itself varies with η in a non-trivial way: In the extensive model, $T^P(y)$ is decreasing (increasing) in $\eta(y)$ if $\alpha(y)$ is above (below) 1 (see Diamond (1980), Saez (2002), Choné & Laroque (2011)). Moreover, the intensive-margin elasticity $\varepsilon(y')$ has only a direct effect on the absolute value of the right-hand side, not on its sign; the indirect effects on T^P are not obvious. To conclude, it is hard to sign the optimal marginal tax T' based on the optimal tax formula (D.2). This problem cannot even be alleviated by imposing structural assumptions on the labor supply elasticities ε and η.

D.8 The mechanism design approach with a continuous skill set

Lemma B.2 provides a simplified expression of optimal allocation problem that is valid both with a discrete and a continuous skill set. In the case of a continuous skill set, the dimensionality of the problem can be reduced further by rewriting the incentive compatibility constraints in terms of an envelope condition. Then, the participation threshold in each skill group ω is given by

$$\hat{\delta}(\omega, \delta_0, y) := \delta_0 - \int_{\underline{\omega}}^{\omega} h_\omega(y(\omega'), \omega') d\omega' \ , \tag{D.3}$$

where δ_0 is the participation threshold of the lowest skill type $\underline{\omega}$. Plugging (D.3) into equation (B.3) allows to write the optimal tax problem as the unconstrained problem to maximize

$$\begin{aligned}\tilde{W}(\delta_0, y) \quad = \quad & \int_{\omega \in \Omega} f(\omega) G_\delta(\hat{\delta}(\omega, \hat{\delta}_0, y) \mid \omega) \left[y(\omega) - h(y(\omega), \omega) + (\alpha(\omega) - 1)\hat{\delta}(\omega, \hat{\delta}_0, y) \right] d\omega \\ & - \int_{\omega \in \Omega} f(\omega)\alpha(\omega) \int_{\underline{\delta}}^{\hat{\delta}(\omega, \hat{\delta}_0, y)} \delta g_\delta(\delta \mid \omega) d\delta d\omega \end{aligned} \tag{D.4}$$

over δ_0 and the output function $y : \Omega \to \mathbb{R}$.

Dealing with an unconstrained maximization problem seems straightforward. Indeed, we can easily differentiate (D.4) point-wise with respect to $y(\omega)$ to get the first-order condition

$$f(\omega') \ G(\hat{\delta}(\omega', \ ) \mid \omega') \cdot \frac{1 - h_y(y(\omega'), \omega')}{h_{y\omega}(y(\omega'), \omega')} =$$

$$\int_{\omega'}^{\bar{\omega}} f(\omega) G(\hat{\delta}(\omega, \ ) \mid \omega) \left\{ \alpha(\omega) - 1 + \frac{g(\hat{\delta}(\omega, \ ) \mid \omega)}{G(\hat{\delta}(\omega, \ ) \mid \omega)} \left[y(\omega) - h(y(\omega), \omega) - \hat{\delta}(\omega, \ ) \right] \right\} d\omega \ . \tag{D.5}$$

The term $1 - h_y(y(\omega'), \omega')$ on the left-hand side captures the intensive-margin distortion and the *implicit marginal tax* at income level $y(\omega')$ in terms of the model's primitives, while the term $y(\omega) - h(y(\omega), \omega) - \hat{\delta}(\omega)$ at the right-hand side captures the extensive-margin distortions and the *implicit participation tax* at income $y(\omega)$.

As in the case of the optimal tax formula (D.2) derived from the perturbation approach, however, the sign of either side cannot be determined for skill type ω' in isolation. By equation (D.5), the intensive-margin distortions in skill group ω' rather depends on the extensive-margin distortions in all skill groups above ω' (and vice versa). As a result, the first-order condition

(D.5) is not directly informative about the optimal signs of marginal taxes and participation taxes. Actually, the first-order condition (D.5) only differs from the optimal tax formula (D.2) derived in Subsection D.7 in that the former is expressed in terms of the model's primitives, while the latter is expressed in terms of empirically estimable objects.

D.9 Empirical evidence on labor supply elasticities

In Section 6, I calibrate my model to match empirical moments for the subgroup of childless singles in the US. This requires to choose calibration targets for the labor supply elasticities at both margins. The current state of the empirical literature is summarized in the survey by Saez et al. (2012) and the meta-study by Chetty, Guren, Manoli & Weber (2013). Unfortunately, there is no clear consensus on the empirical levels of elasticities at both margins. In particular, participation elasticities seem to differ substantially across subgroups of the population. For my calibration, I mainly use the preferred estimates suggested by Chetty, Guren, Manoli & Weber (2013) and Saez et al. (2012). Besides, I consider the studies by Chetty, Friedman & Saez (2013) on the elasticities on EITC recipients and by Bargain et al. (2014) on the elasticities of childless singles as particularly informative for my purposes.

With respect to the intensive-margin elasticity, Saez et al. (2012) consider the best available estimates to be in the range between 0.12 and 0.4. Chetty, Guren, Manoli & Weber (2013) suggest 0.33 as their preferred estimate, and Chetty, Friedman & Saez (2013) estimate an average elasticity for EITC recipients of 0.21 (wage earnings) and 0.36 (total earnings). Bargain et al. (2014) estimate an elasticity of 0.13 for childless singles, with somewhat higher levels among low-skill workers.

With respect to the extensive-margin elasticity, Chetty, Guren, Manoli & Weber (2013) suggest 0.25 as their preferred estimate for the population average based on a meta-analysis. They also argue that elasticities are probably higher around 0.43 in groups such as single mothers who have a lower attachment to the labor force. More recently, Bastian & Jones (2020) estimated a participation elasticity of 0.33 for single mothers, while Kleven (2020*a*) finds an even smaller elasticity. For childless singles in the US, Bargain et al. (2014) estimate an average elasticity around 0.24. Besides, most studies find that participation elasticities are strictly decreasing along the income distribution (see Juhn et al. 1991, 2002 for the US and Meghir & Phillips 2010 for the UK). Miller et al. (2018) and Bastani et al. (2020) provide evidence that participation responses are even decreasing within the group of low-income workers. In particular, Miller et al. (2018) document the results from a recent randomized control trial that studied the effects of a substantial increase in the EITC for childless workers, finding much larger responses among the most vulnerable subgroups. Bastani et al. (2020) exploit a tax reform in Sweden to estimate the extensive-margin responses of married females from low-income families. In contrast, Bargain et al. (2014) find participation elasticities to be by and large constant along the income distribution.

Summing up, there remains a considerable uncertainty about the levels of labor supply responses at both margins. In my benchmark simulations for childless singles, I mainly follow the suggestions of Chetty, Guren, Manoli & Weber (2013). In particular, I consider an intensive-margin elasticity of 0.3 in all skill groups. With respect to the extensive margin, I assume that participation elasticities are decreasing from 0.4 in the lowest skill group to 0.18 in the highest skill group. This gives an average elasticity of 0.25, the preferred value of Chetty, Guren, Manoli & Weber (2013). In my simulations for single parents, I consider an intensive-margin elasticity of 0.3 as well. With respect to the extensive margin, I consider somewhat higher participation elasticities for single parents than for childless singles. Specifically, I assume that participation elasticities are falling from 0.6 in the lowest skill group to 0.3 in the lowest skill group. This gives an average participation elasticity of 0.42, close to the value of 0.43 found in Chetty, Guren, Manoli & Weber (2013).

D.10 Definition and illustration of labor supply distortions

In Subsection 5.1, I provide definitions for labor supply distortions at both margins. These definitions are based on the following thought experiment, which I illustrate in Figures D.7 and D.8 below. Consider an initial allocation in which agent i's bundle is given by $\left(c^i, y^i\right) \geq 0$. Now consider providing agent i with a different bundle $(\tilde{c}, \tilde{y}) \geq 0$ such that $\tilde{y} - y^i = \tilde{c} - c^i \neq 0$. The set of these potential deviations is given by a straight line through $\left(c^i, y^i\right)$ with slope equal to 1, the economy's *marginal rate of transformation* between consumption and output. Agent i's labor supply is said to be distorted if there is a bundle $(\tilde{c}, \tilde{y})$ on this line that i strictly prefers to $\left(c^i, y^i\right)$.

First, it might be possible to increase i's utility through a marginal deviation from $\left(c^i, y^i\right)$. This will be the case if and only if i's marginal rate of substitution, $h_y\left(y^i, \omega^i\right)$, differs from 1. If $h_y\left(y^i, \omega^i\right) < 1$, i would strictly prefer an output-increasing deviation. Then, i's labor supply is said to be *downwards distorted at the intensive margin*. Correspondingly, if $h_y\left(y^i, \omega^i\right) > 1$, i would strictly prefer an output-decreasing deviation, and i's labor supply is said to be *upwards distorted at the intensive margin*.

Second, it might be possible to increase i's utility through a large deviation from $\left(c^i, y^i\right)$ that changes his participation status from non-working (zero output) to working (positive output). Consider an initial allocation with $y^i = 0$ and the deviation to bundle $\left(c^i + \tilde{y}, \tilde{y}\right)$ for some $\tilde{y} > 0$. Agent i would be strictly better off with the new bundle than with his initial bundle if and only if i's total costs of providing output $\tilde{y}$ are below the additional utility from consuming $\tilde{y}$, $h\left(\tilde{y}, \omega^i\right) + \delta^i < \tilde{y}$. Hence, i's labor supply is said to be *downwards distorted at the extensive margin* if both $y^i = 0$ and $\delta^i < \max_{y>0} y - h(y, \omega^i)$.

Correspondingly, it might be possible to increase i's utility through a large deviation from $\left(c^i, y^i\right)$ that changes his participation status from working (positive output) to non-working (zero output). Consider an initial allocation with $y^i > 0$. Agent i would be strictly better off with bundle $\left(c^i - y^i, 0\right)$ than with his initial bundle if and only if i's total costs of providing output y^i exceed the utility from consuming y^i, $h\left(y^i, \omega^i\right) + \delta^i > y^i$. Hence, i's labor supply is said to be *upwards distorted at the extensive margin* if both $y^i > 0$ and $\delta^i > y^i - h\left(y^i, \omega^i\right)$.

Figures D.7 and D.8 illustrate these definitions of labor supply distortions graphically. In Figure D.7, point A marks the initial bundle (c^i, y^i) allocated to agent i, with a strictly positive output $y^i > 0$. The set of hypothetical deviations is given by the solid lines through the points A and B. The indifference curves of Agent i are given by the union of the dashed line and point Z, corresponding to the discontinuity in i's utility due to the fixed cost δ^i. Figure D.7 shows a case where, in point A, the slope of the indifference curve – i.e., the marginal rate of substitution – is below 1. Hence, i's utility could be increased by moving slightly upwards the

Figure D.7: Labor supply distortions, example 1.

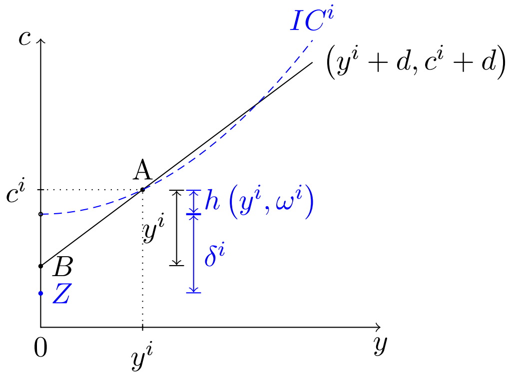

Notes: Figure D.7 depicts an example with labor supply distortions at both margins. Point A represents the current bundle (c^i, y^i) allocated to agent i. The solid black line through points A and B depicts the set of bundles that correspond to the same net resource cost as bundle (c^i, y^i). The union of the dashed blue line and point Z represent an indifference curve of agent i. Labor supply is downwards distorted at the intensive margin because the utility of agent i could be increased by giving him a bundle at the solid black line with a slightly higher output level y than in point A. Labor supply is also upwards distorted at the extensive margin because the utility of agent i could be increased by giving bundle B with zero output instead of bundle A.

solid line. Alternatively, i's utility could also be increased by a large deviation downwards to point B, where output provision is zero. Hence, i's labor supply is at the same time downwards distorted at the intensive margin and upwards distorted at the extensive margin.

In Figure D.8, the initial bundle (c^i, y^i) is again marked by point A. In this case, the initial output of agent i is zero, $y^i = 0$. Again, the set of hypothetical deviations is given by the solid lines through the points A and B. The indifference curves of Agent i are given by the union of the dashed line and point A, corresponding to the discontinuity in i's utility due to the fixed cost δ^i. Figure D.8 shows a case in which point B is located above the indifference curve that corresponds to i's initial utility level. Hence, agent i's utility can be increased by jumping upwards to point B with positive output $y_B = \arg\max_{y>0} \{y - h(y, \omega^i)\}$. Hence, i's labor supply is downwards distorted at the extensive margin.

Figure D.8: Labor supply distortions, example 2.

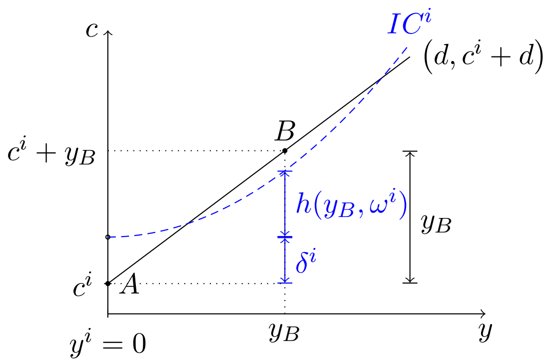

Notes: Figure D.7 depicts an example with a downwards distortion at the extensive margin. Point A represents the current bundle (c^i, y^i) allocated to agent i. The solid black line through points A and B depicts the set of bundles that correspond to the same net resource cost as bundle (c^i, y^i). The union of the dashed blue line and point A represent an indifference curve of agent i. Labor supply is downwards distorted at the extensive margin because the utility of agent i could be increased by giving him bundle B with output level $y_B = \arg\max_y y - h(y, \omega^i)$ instead of point A.

By Lemma B.1, it is possible to characterize the labor supply distortions for the agents in each skill group $j \in J$ simultaneously. In particular, labor supply in skill group j is said to be distorted at the intensive margin if the marginal rate of substitution $h(y_j, \omega_j)$ differs from one for all working agents with skill type ω_j. Similarly, labor supply in skill group j is said to be downwards distorted at the extensive margin if the skill-specific participation threshold δ_j is located below $\max_y y - h(y, \omega_j)$. It is said to be upwards distorted at the extensive margin if the participation threshold δ_j is above $y_j - h(y_j, \omega_j)$.