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

Emanuel Hansen*

December 31, 2019

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* with negative marginal taxes and negative participation taxes at the bottom is optimal if, first, participation elasticities are non-increasing 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 workers 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. The author thanks Felix Bierbrauer, Ulrich Glogowsky, Christian Hellwig, Martin Hellwig, Sebastian Koehne, Wojciech Kopczuk, Etienne Lehmann, Thorsten Louis, Johannes Münster, Nico Pestel, Dominik Sachs, Paul Schempp, Joel Slemrod, Dezsö Szalay and John Weymark as well as participants at the ASSA Meetings 2019 in Atlanta, the NTA Annual Meeting 2018 in New Orleans, the ZEW Public Finance Conference 2017 in Mannheim, the Taxation Theory Conference 2016 in Toulouse, the IIPF Doctoral School 2015 in Warsaw, the IIPF 2014 in Lugano, the VfS 2014 in Hamburg and at seminars in Bonn, Cologne and Umea for helpful comments and discussions. The usual disclaimer applies.

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 that transfer resources towards the poor in the US (Nichols & Rothstein 2015). The EITC differs from traditional welfare programs in that it entails negative marginal taxes and negative participation 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. 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] Saez (2002) has already shown 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.

In the theoretical part of this paper, I provide analytical results on the optimality of an EITC. For this purpose, I study optimal taxation in a 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 the intensive margin and the extensive margin. I focus on cases in which participation elasticities are non-increasing over the skill distribution, in line with the existing empirical evidence (e.g., Juhn et al. 1991, 2002, Meghir & Phillips 2010, Bargain et al. 2014, Bastani et al. 2019). In this setting, I derive the following necessary and sufficient conditions. First, an EITC with negative marginal taxes for incomes below

[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 other, Chetty, Guren, Manoli & Weber (2013).

some phase-in endpoint $\bar{y}$ *can only* be optimal if the desire for redistribution among the EITC recipients is weak, i.e., if the skill/income gradient of welfare weights is close enough to zero at the bottom of the skill distribution.[3] Second, the optimality of an EITC is *ensured* if, additionally, there is a sufficiently large share of agents that are more productive than the EITC recipients. As will become clear, this sufficient condition (only) imposes a restriction on the level of the phase-in endpoint $\bar{y}$.

To simplify the exposition, I first derive a version of my results in a stylized model with three skill groups and particularly simple functional forms. This stylized model allows to derive necessary and sufficient conditions for an EITC in a straightforward way, and to explain the economic mechanism at work in more detail. In the stylized model, the sufficient condition is satisfied if both the population share and the productivity of the highest-skilled group are large enough. Second, I demonstrate that my results extend to a more general 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 this general model, the sufficient condition is always satisfied for an EITC with a small phase-in endpoint $\bar{y}$, but not for an EITC with a large phase-in endpoint. 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 C.1).

In the numerical part of this paper, I simulate 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 condition discussed above is 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 (CPS) 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 a limited effect on the quantitative properties of the optimal tax schedule. By contrast, the assumed preferences for redistribution have a crucial impact: An EITC is optimal whenever 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

[3] I restrict my attention to cases in which welfare weights are non-increasing along the skill distribution (i.e., the social planner has a standard desire for redistribution from high-skill to low-skill workers).

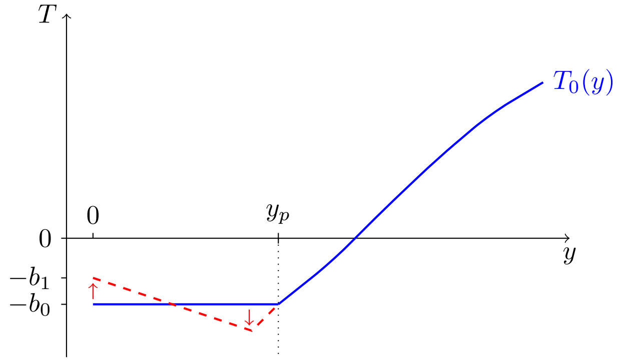

Figure 1: The efficiency effects of a small EITC.

and Jacquet et al. 2013). Additionally, I provide the results of numerical simulations for the subgroup of single parents in the US.[4]

The intuition behind my results comes from a trade-off between intensive-margin distortions and extensive-margin distortions. Figure 1 depicts a tax perturbation that helps to understand this trade-off. Assume that the social planner uses an income tax to transfer resources from rich agents with incomes above some threshold y_p to poor agents with incomes below y_p. Specifically, let the status quo tax schedule T_0 be given by the solid blue line in Figure 1: the poor receive the same transfer $b_0 > 0$ for any income below y_p, financed by strictly increasing taxes for incomes above y_p. Under T_0, marginal taxes and participation taxes are hence equal to zero for the poor and strictly positive for the rich. This implies that labor supply of the poor is undistorted at both margins, while labor supply of (potentially) rich agents is downwards distorted at both margins.

Now consider a tax perturbation that slightly reduces the transfers to non-working agents and working agents with very low incomes and slightly increases the transfers to agents with incomes close to y_p (see dashed red line in Figure 1). This introduces a (small) EITC with negative marginal taxes and negative participation taxes for incomes below y_p. At the *extensive margin*, the reduction of the transfer for non-workers from b_0 to b_1 leads to a first-order reduction in the existing downward distortions among the rich, while the introduction of negative participation taxes induces small (second-order) upward distortions among the poor. Letting the negative participation tax grow with income implies that the upwards distortions at the extensive margin are increased for workers with incomes close to y_p, while upwards distortions are decreased for workers with incomes close to 0. From an extensive-margin perspective, this is optimal by the inverse elasticity rule if participation responses are decreasing in income (in line with the empirical evidence).[5] At the *intensive margin*, however, the negative marginal tax introduces small (second-order) upwards distortions among the poor. Hence, there is a tradeoff between *intensive efficiency* and *extensive efficiency*. Initially, however, there is

[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 C.5).

[5] If the participation *elasticity* is constant in income, the *semi-elasticity* of participation is still decreasing. As I show in Section 4, this is enough to ensure that an EITC can be optimal.

a first-order reduction in extensive-margin distortions but only a negligible (second-order) increase in intensive-margin distortions. As a result, the introduction of a small EITC as depicted by the dashed red line in Figure 1 unambiguously increases overall efficiency.

To determine the welfare effect, we need to take into account that the tax perturbation redistributes resources from the very poor with incomes close to 0 to the poor with incomes close to y_p. If the planner assigns much higher welfare weights to the very poor than to the poor, these equity losses outweigh the efficiency gains of an EITC introduction: welfare goes down. If the planner assigns (almost) equal welfare weights to the very poor and the poor, the efficiency gains dominate: The introduction of an EITC increases welfare.

The results of this paper give rise to the question what a convincing assumption on the skill (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 assumption of a utilitarian welfare function and standard utility functions 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 face on average higher fixed costs, 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.

The paper proceeds as follows. Section 2 briefly reviews the related literature. Section 3 introduces the basic framework and the optimal tax problem. Section 4 subsequently provides my results for a stylized model with three skill types and a general model with a finite, but arbitrarily large number of skill types. Section 5 provides numerical simulations for a calibrated version of the model. Section 6 concludes. Appendix A contains the formal proofs for the basic framework and the stylized model. Online Appendices B and C 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). They use different models and different approaches to solve for the optimal income tax, but come to similar results. In particular, 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. Jacquet et al. (2013) use a mechanism design approach to derive first-order conditions that characterize the optimal allocation. In Online Appendices C.7 and C.8, I present versions of their optimal tax formula and first-order conditions, respectively, and explain why neither of them pins down the optimal sign of the marginal tax. Both papers note that the optimal marginal tax *might* be negative, but neither of them provides examples in which this is indeed the case. Both papers also perform numerical simulations in which the optimal tax is always monotonically increasing over the range of positive income levels.[6]

An additional contribution of Jacquet et al. (2013) is to provide a sufficient condition for optimal marginal taxes to be positive everywhere below the top, based on the analysis of a relaxed problem (the *first-and-half-best*). Based on numerical simulations, they argue that marginal taxes should even be positive if this sufficient condition is not met. Lorenz & Sachs (2012) complement these results by providing a condition under which optimal participation taxes are positive everywhere as well.

To conclude, none of the previous papers on optimal income taxation with two margins provides general conditions or even a numerical example in which the optimal marginal tax at the bottom is negative. I fill this gap by providing both necessary and sufficient conditions, based on a refinement of the *first-and-half-best* approach.

Second, there are a few previous papers that provide alternative rationalizations for 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 (2017), a negative marginal tax helps to correct for the optimization errors that workers make because of a present bias. In my model, there is no informal labor

[6]Besides, they 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-employed 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.

market, welfare weights are monotonically decreasing along the skill distribution, and agents behave individually rational. By contrast, I show 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.

3 The basic framework

This section introduces a basic framework to study optimal taxation with labor supply responses at both margins. I start by defining the economic environment. Then, I derive individually optimal earnings choices and define the optimal tax problem.

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 component that the agents bears if and only if she provides strictly positive output. Formally, the agent's preferences can be represented by the utility function

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

The variable effort cost is captured by function h. It depends on the output level y^i and the parameter $\omega^i \in \Omega$, which is commonly referred to as agent i's *skill type*. I impose the standard assumptions that h satisfies $h(0, \omega) = 0$, $h_y(y, \omega) > 0$, $h_{yy}(y, \omega) > 0$ and $h_{y\omega}(y, \omega) < 0$ for all $y > 0$ and $\omega \in \Omega$. The fixed cost of providing output is captured by the parameter $\delta^i \in \Delta$, which I refer to as i's *fixed cost type*.

Each agent is privately informed about his skill type ω^i and his fixed cost type δ^i. The set of skill types Ω may be given by a discrete or continuous subset of $\mathbb{R}$. In principle, the set of fixed cost types Δ could also be given by a discrete or a continuous subset of $\mathbb{R}$. For reasons that become clear below, however, I focus on the case where Δ is an interval $\left[\underline{\delta}, \bar{\delta}\right]$. 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 impose an assumption on the sets Ω and Δ. In particular, I denote by $\underline{\omega}$ and $\bar{\omega}$ the lowest and the highest elements of Ω, respectively. Correspondingly, I denote by $\underline{\delta}$ and $\bar{\delta}$ denote the lowest and the highest elements of Δ, respectively. I assume that

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

This assumption implies that, under laissez-faire, some agents with any skill type $\omega \in \Omega$ provide strictly positive output, while some other agents 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 (c, y) 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 an 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

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

where R is an exogenous revenue requirement, which I set to $R = 0$ for the remainder of this paper. The optimal tax problem is to maximize welfare (3) over the set of tax schedules that satisfy (4), taking into account the labor supply responses captured by $y^*(\omega, \delta \mid T)$ for all types in $\Omega \times \Delta$.

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 affect 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. 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).[7]

3.2 Individual labor supply choices

I start by studying the individually optimal behavior of agents with different types. Fix a tax schedule $T : \mathbb{R}_0^+ \to \mathbb{R}$. For expositional reasons, focus on a tax schedule that is continuously differentiable and weakly convex in y.[8] An agent with type (ω, δ) maximizes

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

[8] If T is not continuously differentiable, $y_T(\omega)$ may be located at a kink and fail to satisfy (6). If T is strictly concave, (6) may be satisfied by multiple output levels. In the remainder of the paper, I do not impose either of these assumptions.

her utility (1) by choosing the output level

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

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

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

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

$$\delta_T(\omega) := y_T(\omega) - h(y_T(\omega), \omega) - [T(y_T(\omega)) - T(0)] \ . \tag{7}$$

For any T, hence, labor supply is characterized completely by the skill-specific incomes $y_T : \Omega \to \mathbb{R}_+$ and the skill-specific participation thresholds $\delta_T : \Omega \to \mathbb{R}$. Conditional on providing positive output, an agent's optimal intensive-margin choice depends only on her skill type ω. 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 $\delta_T(\omega)$. Hence, the resulting allocation involves pooling by all working agents with skill type ω and fixed-cost type below $\delta_T(\omega)$. There is also pooling by all agents with fixed cost type above the skill-specific threshold $\delta_T(\omega)$. As a result, T gives rise to an income distribution F_y with a mass point at zero and a distribution over the image of $y_T : \Omega \to \mathbb{R}$. For the case of a continuous skill set, I denote by f_y the density of the implemented income distribution.

As in the classical Mirrlees (1971) framework, the taxation of income in general induces labor supply distortions. In particular, labor supply is said to be *downwards distorted at the intensive margin* for workers with skill type ω if the marginal rate of substitution between income and consumption is strictly smaller than the economy's marginal rate of transformation, $h_y(y_T(\omega), \omega) < 1$. Labor supply is said to be *upwards distorted at the intensive margin* if $h_y(y_T(\omega), \omega)$ is strictly larger than 1. By the first-order condition (6), labor supply in skill group ω is downwards (upwards) distorted at the intensive margin, if the marginal tax at income level $y_T(\omega)$ is positive (negative).

In the model studied here, however, there is a second type of distortions: labor supply is said to be *downwards distorted at the extensive margin* for workers with skill ω if the participation threshold $\delta_T(\omega)$ is strictly smaller than $\max_{y>0} y - h(y, \omega)$, the maximal surplus that a worker with skill ω can generate. Labor supply is said to be *upwards distorted at the extensive margin* if participation threshold $\delta_T(\omega)$ is strictly larger than $y_T(\omega) - h(y_T(\omega), \omega)$, the surplus a worker with skill ω generates under tax T. By the cutoff condition (7), labor supply in skill group ω is downwards (upwards) distorted at the extensive margin, if the participation tax at income level $y_T(\omega)$ is positive (negative).

Correspondingly, tax changes can induce two types of labor supply responses. As usual, I measure intensive-margin responses by the elasticity of income with respect to the retention rate,

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

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

$$\eta_T(\omega) := \frac{\partial G_\delta(\delta_T(\omega) \mid \omega)}{\partial\left[y_T(\omega) - T^P(y_T(\omega))\right]} \frac{1}{G_\delta(\delta_T(\omega) \mid \omega)} = \frac{g_\delta(\delta_T(\omega) \mid \omega)}{G_\delta(\delta_T(\omega) \mid \omega)} ,$$

where $g_\delta(\cdot \mid \omega)$ and $G_\delta(\cdot \mid \omega)$ denote the pdf and the cdf of the fixed cost distribution in skill group ω, respectively. Alternatively, extensive-margin responses could be measured by the (standard) elasticity of participation. Both the semi-elasticity and the elasticity are only well-defined if the set of fixed cost types Δ is continuous.[9] In general, both $\varepsilon_T(\omega)$ and $\eta_T(\omega)$ vary with the skill type ω and the tax schedule T.

3.3 The optimal tax problem

I use a mechanism design approach to study the optimal income tax. This approach requires, in the first step, to solve for the optimal allocation and, in the second step, to identify the tax schedule that allows to decentralize this optimal allocation. Specifically, the optimal signs of optimal marginal and participation taxes can be identified based on the labor supply distortions at both margins as shown above.

The optimal allocation (c, y) maximizes welfare over the set of allocations that are implementable, i.e., feasible,

$$E_{\Omega\times\Delta}\left[y(\omega, \delta) - c(\omega, \delta)\right] \geq 0 , \tag{8}$$

and incentive compatible,

$$c(\omega', \delta') - h\left[y(\omega', \delta'), \omega'\right] - \mathbb{1}_{y(\omega',\delta')>0}\ \delta' \geq c(\omega'', \delta'') - h\left[y(\omega'', \delta''), \omega'\right] - \mathbb{1}_{y(\omega'',\delta'')>0}\ \delta' \tag{9}$$

for each pair (ω', δ') and (ω'', δ'') in $\Omega \times \Delta$. In words, the last constraint requires that each type (ω', δ') is assigned a bundle of consumption and income that makes him weakly better off than the bundle assigned to any other type (ω'', δ''). The two-dimensional type set makes this maximization problem more challenging.

The separability assumptions imposed on utility (1) allow to simplify the problem considerably, however. First, I can restrict the maximization to allocations in which all working agents with skill type ω receive the same consumption-income bundle, and all non-working agents receive the same consumption c_0 (see Lemma 4 in Appendix A). Recall that any well-behaved tax schedule T implements an allocation with the same kind of pooling. Second, incentive compatibility along the fixed cost dimension requires that an agent with skill type ω works if and only if his fixed cost type is below the participation

[9] By contrast, if Δ is a discrete set, the semi-elasticity is either zero or not defined for any tax T.

threshold $\hat{\delta}(\omega) = c(\omega) - h(y(\omega), \omega) - c_0$ for each skill level ω. Finally, if we rewrite the allocation in terms of income levels and participation thresholds (instead of consumption levels), we can substituting the feasibility condition into the objective and eliminate the consumption level c_0 of non-working agents from the maximization problem.

Lemma 1. *The optimal allocation problem is equivalent to maximizing*

$$E_\Omega \left[G_\delta(\hat{\delta}(\omega) \mid \omega) \Big(y(\omega) - h\left[y(\omega), \omega\right] + \hat{\delta}(\omega)\left[\alpha(\omega) - 1\right] \Big) - \alpha(\omega) \int_{\underline{\delta}}^{\hat{\delta}(\omega)} \delta dG_\delta(\delta \mid \omega) \right] \tag{10}$$

over $y : \Omega \to \mathbb{R}_0^+$ *and* $\hat{\delta} : \Omega \to \mathbb{R}$, *subject to the incentive compatibility constraint along the skill dimension that, for any* ω' *and* ω'' *in* Ω,

$$\hat{\delta}(\omega') - \hat{\delta}(\omega'') \geq h\left[y(\omega''), \omega''\right] - h\left[y(\omega''), \omega'\right] \ . \tag{11}$$

Lemma 1 expresses the optimal tax problem as the problem to maximize welfare over two skill-specific functions – incomes and participation thresholds – subject to a set of incentive-compatibility constraints along the skill dimension only. This is particular useful because the labor supply distortions at both margins are defined in terms of the skill-specific incomes and participation thresholds, as shown in Subsection 3.2 above.

Lemma 1 holds whether the skill set Ω is continuous or discrete. Below, I will use it to solve for the optimal allocation given a discrete set Ω. As I will show, the efficiency properties of the optimal allocation are easier to determine in this case than in the case of a continuous skill set.[10]

3.4 The first-and-half best approach

Jacquet et al. (2013) develop a method that sometimes allows for signing the optimal marginal tax despite the complications outlined above. For this purpose, they consider a setting with a continuous skill set Ω. In particular, they suggest to solve the relaxed problem of maximizing (10) over the functions $y : \Omega \to \mathbb{R}$ and $\delta : \Omega \to \mathbb{R}$ while ignoring the incentive compatibility constraint (11). This problem takes into account incentive compatibility along the fixed cost dimension, but not along the skill dimension. I follow Jacquet et al. (2013) in referring to the solution of this relaxed problem as the *first-and-half-best* allocation and denote it by $(\tilde{y}, \tilde{\delta})$. Obviously, solving the relaxed problem is much easier than solving the original problem. In particular, point-wise maximization over $y(\omega)$ and $\hat{\delta}(\omega)$ gives rise to simple skill-specific first-order conditions.

[10]In Online Appendix C.8, I show that the optimal allocation problem can be written even more compactly if the skill set is continuous. Then, it is straightforward to derive first-order conditions that jointly define the intensive-margin and extensive-margin distortions in the optimal allocation. However, these first-order conditions do not allow to determine the optimal direction of distortions easily.

Lemma 2. *[Jacquet et al. (2013)] The first-and-half-best allocation is defined by*

$$h_y(\tilde{y}(\omega),\omega) - 1 \;=\; 0\;, \tag{12}$$

$$\tilde{y}(\omega) - h(\tilde{y}(\omega),\omega) - \tilde{\delta}(\omega) \;=\; \frac{1-\alpha(\omega)}{\tilde{\eta}(\omega)}\;, \tag{13}$$

where $\tilde{\eta}(\omega) = g_\delta(\tilde{\delta}(\omega) \mid \omega)/G_\delta(\tilde{\delta}(\omega) \mid \omega)$ *is the semi-elasticity of participation at skill* ω*.*

Lemma 2 recovers *Lemma 1* in Jacquet et al. (2013) and extends it to models with a discrete skill set Ω. Condition (12) implies that labor supply in the first-and-half-best is undistorted at the intensive margin in all skill groups. Condition (13) implies that labor supply is downwards at the extensive margin in skill group ω if $\alpha(\omega) < 1$ and upwards distorted at the extensive margin if $\alpha(\omega) > 1$.

Using $\tilde{\delta}(\omega) = \tilde{c}(\omega) - h(\tilde{y}(\omega),\omega) - \tilde{c}_0$, equation (13) also implicitly defines the *first-and-half-best tax* as

$$T_{FH}(\tilde{y}(\omega)) = \tilde{y}(\omega) - \tilde{c}(\omega) = \frac{1-\alpha(\omega)}{\tilde{\eta}(\omega)} - \tilde{c}_0\;. \tag{14}$$

The corresponding participation tax is given by $T_{FH}(\tilde{y}(\omega)) - T_{FH}(0) = T_{FH}(\tilde{y}(\omega)) + \tilde{c}_0$. Its sign only depends on the difference $1-\alpha(\omega)$, while its level also depends on the semi-elasticity of participation $\tilde{\eta}(\omega)$. Generically, the implicit tax T_{FH} varies with income y, implying a marginal tax that differs from 0.[11] But then, individual maximization given T_{FH} would lead to labor supply distortions at the intensive margin, violating (12). Hence, the first-and-half-best allocation is not implementable.

Jacquet et al. (2013) show that the first-and-half-best is nevertheless helpful: If T_{FH} is monotonically increasing in y, then the optimal second-best marginal tax is strictly positive everywhere below the very top, just as in the intensive model. If T_{FH} is non-monotonic, in contrast, they are unable to sign the optimal marginal tax.[12] In the following, I refine their approach to derive conditions under which the first-and-half-best tax is non-monotonic *and* the second-best marginal tax is negative at the bottom.

4 Optimal taxes in the discrete model

From now on, I focus on discrete versions of the basic model introduced in Section 3.1. In particular, I assume that the skill set is given by a finite set $\Omega = \{\omega_1, \omega_2, \ldots, \omega_n\}$, where $\omega_{j+1} > \omega_j$ for all $j \in \{1, 2, \ldots, n-1\}$. To simplify the notation, I denote the share of agents with skill type ω_j by $f_j = f(\omega_j)$ and the distribution of fixed cost types among agents with skill type ω_j by $G_j(\cdot) = G_\delta(\cdot \mid \omega_j)$. Correspondingly, I denote the consumption, output, participation threshold and welfare weight of workers with skill type ω_j by c_j, y_j, δ_j and α_j for each j in $J := \{1, 2, \ldots, n\}$, respectively. I start by

[11]Specifically, T_{FH} is constant in y if and only if $\alpha(\omega) = 1$ for all $\omega \in \Omega$; it is then given by $T_{FH}(y) = 0$ and implements the laissez-faire allocation.

[12]Based on numerical simulations, however, they argue that the optimal marginal tax should even be positive if the monotonicity condition is violated.

providing necessary and sufficient conditions for an EITC with negative marginal taxes in a stylized model with three skill types. I then use the stylized model to provide the intuition for why an EITC is sometimes optimal. Afterwards, I extend my results to a more general model.

4.1 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. I normalize the lowest skill level to $\omega_1 = 1$. Hence, the relative distance between the first two skill groups is given by $\omega_2 = \omega_2/\omega_1$. The relative distance between the two highest skill types is given by ω_3/ω_2.

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

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

This implies that the (intensive-margin) elasticity $\varepsilon_T(\omega)$ of individually optimal income with respect to the retention rate $1 - T'(y)$ is equal to parameter $\sigma > 0$ for all agents.

Third, fixed costs are uniformly distributed on the same interval $\left[0, \bar{\delta}\right]$ in each skill group. Hence, the pdf of the fixed cost distribution satisfies

$$g_j(\delta) = g_\delta(\delta \mid \omega_j) = \begin{cases} 1/\bar{\delta} & \text{for} \quad \delta \in \left[0, \bar{\delta}\right] \\ 0 & \text{for} \quad \delta \notin \left[0, \bar{\delta}\right] \end{cases} \tag{16}$$

for each $j \in \{1, 2, 3\}$.[13] Given this distribution, the participation elasticity is equal to $1+\sigma$ for each skill group in the laissez-faire allocation, while the semi-elasticity of participation is decreasing along the skill distribution.

The following assumption summarizes these properties.

Assumption 1. *The skill set is* $\Omega = \{1, \omega_2, \omega_3\}$*, the effort cost function is given by* (15) *with* $\sigma > 0$ *and the fixed cost distribution is given by* (16)*.*

The sequence of welfare weights is decreasing with $\alpha_1 \geq \alpha_2 \geq \alpha_3$ and normalized to have a population average of 1. For any given pair (α_1, α_2), the welfare weight α_3 of high-skill workers is hence given by

$$\alpha_3 = \frac{1 - f_1\alpha_1 - f_2\alpha_2}{f_3} . \tag{17}$$

Under Assumption 1, the set of implementable allocations depends only on the collection of parameters $(\sigma, \bar{\delta}, \omega_2, \omega_3, f_1, f_2)$. In the following, I fix all these parameters and

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

vary only the welfare weights α_1 and α_2. This allows me to solve for the optimal allocation given alternative redistributional concerns. In particular, my goal is to determine the optimal direction of labor supply distortions at both margins for different combinations (α_1, α_2). For this purpose, I solve the optimal tax problem in the formulation provided in Lemma 1. I focus on parameter constellations such that there is a unique welfare maximum that can be identified using the first-order approach. In the stylized model, this is ensured if the welfare weights of all three skill groups are below 2.[14]

To identify conditions for the optimality of upwards distortions, I first investigate the first-and-half-best allocation (see Lemma 2). Under Assumption 1, I can derive the following closed-form expressions for the first-and-half-best:

$$\tilde{y}_j = \omega_j^{1+\sigma} , \tag{18}$$

$$\tilde{\delta}_j = \frac{\omega_j^{1+\sigma}}{(1+\sigma)(2-\alpha_j)} \tag{19}$$

for any $j \in \{1, 2, 3\}$. The first-and-half-best tax then follows as

$$T_j^{FH} = \tilde{y}_j - h(\tilde{y}_j, \omega_j) - \tilde{\delta}_j - \tilde{c}_0 = \frac{\omega_j^{1+\sigma}}{1+\sigma} \frac{1-\alpha_j}{2-\alpha_j} - \tilde{c}_0 . \tag{20}$$

These closed-form expressions simplify the analysis in two important ways. First, they make it easy to check whether the first-and-half best tax is monotonic or non-monotonic. Second, they allow to verify easily whether the first-and-half-best allocation satisfies the incentive compatibility (IC) constraint between workers with skill types ω_j and ω_k,

$$\tilde{\delta}_j - \tilde{\delta}_k \geq h\left(\tilde{y}_k, \omega_k\right) - h\left(\tilde{y}_k, \omega_j\right) .$$

4.1.1 No desire for redistribution at the bottom

I start with the limit case where α_1 and α_2 are equal to the same number $\alpha_p > 1$. 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. In this case, the first-and-half-best participation tax $T_j^{fhb} + \tilde{c}_0$ as defined by (20) is negative for ω_1 and ω_2 workers, and positive for ω_3 workers. Moreover, the tax is non-monotonic: it is decreasing between y_1^{fhb} and y_2^{fhb}, and increasing between y_2^{fhb} and y_3^{fhb},

$$T_2^{fhb} < T_1^{fhb} < -\tilde{c}_0 < T_3^{fhb} . \tag{21}$$

[14] 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. In particular, there may be local welfare minima or multiple maxima and mimina. In an earlier version of the paper, I derive a sufficient condition for the existence of a unique and well-defined optimum (Hansen 2018). In all cases I consider, an EITC turns out to be optimal for a subset of the parameter constellations such that a well-behaved welfare optimum exists.

Hence, the sufficient condition of Jacquet et al. (2013) for optimal marginal taxes to be positive is violated. But this does not yet ensure that the optimal marginal tax is negative, i.e., that the optimal allocation involves upward distortions at the intensive margin. The following proposition provides necessary and sufficient conditions for the optimal allocation to involve upwards distortions at both margins.

Proposition 1. *Under Assumption 1, there are three numbers* $\bar{\beta}(\omega_2) \in (1, 2)$, $\phi_\omega(\alpha_p) > 1$ *and* $\phi_f(\alpha_p, \omega_3/\omega_2) \in (0, 1)$ *such that, at the intensive margin, optimal labor supply in skill group* 2

(i) can only be upwards distorted if $\alpha_p > \bar{\beta}(\omega_2)$*;*

(ii) is upwards distorted if $\alpha_p > \bar{\beta}(\omega_2)$, $\omega_3/\omega_2 > \phi_\omega(\alpha_p)$ *and* $f_3 > \phi_f(\alpha_p, \omega_3/\omega_2)$*.*

In case (ii)*, optimal labor supply is also upwards distorted at the extensive margin in skill groups* 1 *and* 2*.*

I provide closed-form expressions for $\bar{\beta}$ and ϕ_ω, and the definition of ϕ_f in the appendix. By part (i), optimal marginal taxes can only be negative if the desire for redistribution towards the poor and the very poor is sufficiently strong (as measured by α_p). Figure 2a below illustrates this condition for a numerical example where the intensive-margin elasticity is given by $\sigma = 0.3$. The necessary condition is satisfied for all parameter constellations in the shaded area. As can be seen in Figure 2a, the threshold $\bar{\beta}$ is strictly increasing in ω_2, the relative distance between the lower two skill levels. If this distance vanishes, $\omega_2 \to 1$, the threshold $\bar{\beta}$ converges to 1. In the limit, hence, the necessary condition is satisfied whenever α_1 and α_2 are equal and above 1, the average welfare weight.

By part (ii), the optimality of negative marginal taxes is ensured 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. Figure 2b illustrates the interdependence between the two latter conditions for a numerical example with $\sigma =$

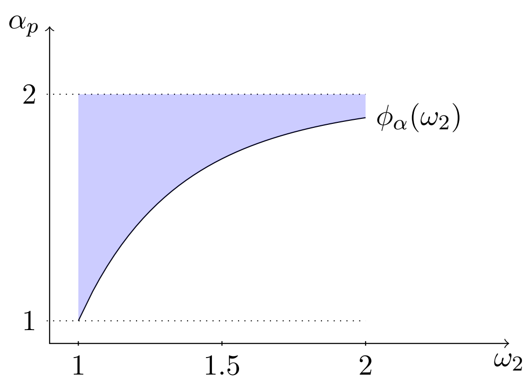

(a) Necessary condition on α_p.

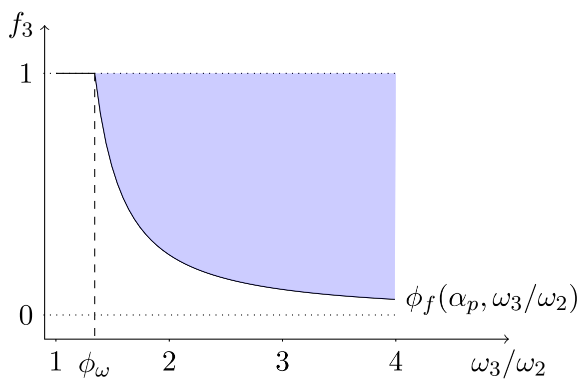

(b) Sufficient conditions on ω_3 and f_3.

Figure 2: Conditions for an optimal EITC. *Parameter values:* $\sigma = 0.3$, $\alpha_p = 1.2$ *(right panel). Shaded areas: Optimal allocation involves upwards distortions at both margins.*

0.3 and $\alpha_p = 1.2$. Again, the shaded area represents the set of parameter constellations for which the jointly sufficient conditions are met, i.e., the second-best allocation involves upwards distortions at both margins.

Finally, note that Proposition 1 only provides conditions for labor supply in skill group 2 to be upwards distorted at the intensive margin. In the formal appendix, I show that labor supply in skill groups 1 and 3 is never upwards distorted at this margin if Assumption 1 holds and welfare weights are decreasing with $\alpha_1 \geq \alpha_2 \geq \alpha_3$.

4.1.2 Sketch of proof

The formal proofs of Proposition 1 and the following Proposition 2 are provided in Appendix A. They proceed in two steps. In step (a), Lemma 5 in the appendix investigates incentive compatibility in the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$. First, I show that the upward IC of ω_1 workers is satisfied if and only if α_p is below a uniquely defined threshold $\bar{\beta}(\omega_2) > 1$. Second, the downward IC of ω_3 workers is satisfied if the welfare weight α_3 of high-skill workers is above some critical value $\beta_D(\alpha_p)$. Finally, I show that this condition on α_3 is equivalent to the condition that both ω_3 and f_3 are large enough to exceed the uniquely thresholds $\phi_\omega(\alpha_p)$ and $\phi_f(\alpha_p, \omega_3/\omega_2)$, respectively. The proof provides closed-form expressions of the functions $\bar{\beta}$, ϕ_ω and the critical value β_D as well as the precise definition of ϕ_f.

In step (b), the proof of part (i) is completed by showing that labor supply in the second-best allocation cannot be upwards distorted if the first-and-half-best allocation satisfies the upward IC constraints of ω_1 and ω_2 workers. The proof of part (ii) is completed by showing that labor supply in skill group 2 is always upwards distorted at the intensive margin if $(\tilde{y}, \tilde{\delta})$ violates the upward IC of ω_1 workers and satisfies the downward IC of ω_3 workers. Basically, distorting y_2 upwards allows to slacken the upward IC of ω_1 workers and to implement an allocation that is closer to the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$.

4.1.3 Intuition

For the intuition behind Proposition 1, recall first that the welfare weights α_1 and α_2 are both equal to $\alpha_P > 1$, while α_3 is below one. Hence, the planner has a desire to transfer resources from the rich to the poor and the very poor, but does not care for how these transfers are shared between the poor and the very poor. Put differently, the planner's goal is to transfer resources towards ω_1 and ω_2 in the way that induces as little labor supply distortions as possible. In the following, I explain why this goal is achieved by an EITC with negative marginal and participation taxes.

I start by explaining the properties of the first-and-half-best tax, which maximizes welfare if only the labor supply distortions at the extensive margin are taken into account. First, by (21), the first-and-half-best participation taxes to ω_1 and ω_2 workers are negative for any $\alpha_p > 1$. For the intuition, consider an initial tax schedule T_0 that provides equal transfers to non-working agents and working agents in skill groups 1 and 2, financed by taxes paid by ω_3 workers. Given T_0, labor supply by ω_1 and ω_2 workers is undistorted, while labor supply by ω_3 workers is downwards distorted at the extensive margin. While a small increase in the transfer to non-working agents leads to a first-order increase in the downwards distortions in skill group 3, a small increase in the transfers to ω_1 and ω_2 workers has no effect on labor supply by ω_3 workers and only leads to second-order distortions in the two low-skill groups. Hence, negative participation taxes allow the planner to reduce the amount of extensive-margin distortions.

Second, by equation (14), the planners sets the transfers to satisfy

$$T_j^{fhb} + c_0 \quad = \quad -\frac{\alpha_p - 1}{\tilde{\eta}_j} < 0 \text{ for each } j \in \{1, 2\} \ . \tag{22}$$

This is an *inverse elasticity rule*: As negative participation taxes induces upwards distortions, they should be more negative for the skill groups that responds less elastically as measured by the semi-elasticity of participation $\tilde{\eta}_j$. Under Assumption 1, the semi-elasticity of skill group 1 is always larger than the semi-elasticity of skill group 2. Hence, the first-and-half-best tax is strictly decreasing from $\tilde{y}_1$ to $\tilde{y}_2$ whenever the welfare weights α_1 and α_2 are equal to the same number $\alpha_p > 1$. The same is true if α_1 and α_2 are not equal, but close to each other and above 1. If α_1 is much larger than α_2, by contrast, the first-and-half-best tax is increasing.

Third, the higher α_p is, the larger are the transfers to both groups and the more decreasing is the first-and-half-best tax. Again, the intuition comes from the *inverse elasticity rule*: A larger share of each additional transfer should be provided to the less responsive group in order to limit the additional distortions. Specifically, Assumption 1 implies that the ratio of the semi-elasticities $\tilde{\eta}_1/\tilde{\eta}_2 > 1$ remains constant if α_p changes. This means that, if α_p increases, the transfers to both groups of workers should grow proportionally. As a result, the difference between both transfers is strictly increasing in α_p as well. In the model with three skill groups, the transfer to ω_2 workers are so much larger than the transfers to ω_1 workers that the upward IC of ω_1 workers is violated if and only if α_p is above a unique cutoff $\bar{\beta}$. In this case, the first-and-half-best allocation is not implementable.

To design the optimal second-best tax, the social planner also has to respect the IC constraints along the skill dimension. For $\alpha_p > \bar{\beta}$, however, the upward IC of ω_1 workers makes it impossible to follow the inverse elasticity rule. To slacken this constraint, the planner has to distort 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. Note, however, that introducing a slightly negative marginal tax comes with first-order reductions of extensive-margin distortions, but introduces only second-order distortions at the intensive margin. This suggests that the optimal second-best tax is decreasing at the bottom as well (as the dashed red line in Figure 1).

Additionally, the social planner needs to respect the downward IC of ω_3 workers. For any $\alpha_p > 1$, the first-and-half-best tax for ω_3 workers is positive, while the tax for ω_2 workers is negative. If the tax on ω_3 workers is large, then allocation $(\tilde{y}, \tilde{\delta})$ violates the IC constraint of ω_3 workers. To slacken this constraint, the labor supply of ω_2 workers has to be distorted *downwards* at the intensive margin. But then, the planner would have to deviate even further from the inverse elasticity rule. In general, it is unclear whether the optimal allocation involves upwards distortions or downwards distortions at the intensive margin in skill group 2.

By part (ii) of Proposition 1, however, a negative marginal tax is unambiguously optimal if the population share f_3 and the productivity ω_3 of high-skill workers are large enough. For the intuition behind this result, assume that the social planner wants to provide an average transfer of A to the low-skill workers. To finance these transfers, each ω_3 worker has to pay some tax $T(y_3)$. The larger the share f_3 of high-skilled worker is, the smaller is the required tax payment $T(y_3)$. Hence, the first-and-half-best bundle assigned to ω_3 workers gets more attractive relative to the bundle of ω_2 workers. Moreover, the larger the productivity ω_3, the higher are the gross income y_3 and the net income $y_3 - T(y_3)$ of high-skill workers. Again, this makes the bundle of ω_3 workers relatively more attractive. Put differently, the larger are f_3 and ω_3, the less increasing is the first-and-half-best tax between y_2 and y_3. If bth parameters are large enough to exceed the thresholds ϕ_f and ϕ_ω, respectively, the first-and-half-best tax is sufficiently flat to satisfy the downward IC of ω_3 workers. Less formally, the larger are f_3 and ω_3, the smaller is the average marginal tax between income levels y_2 and y_3 and the smaller are the induced intensive-margin distortions. Eventually, an EITC is optimal because its extensive-margin benefits dominate against the intensive-margin costs.

4.1.4 Weak desire for redistribution at the bottom

In the next step, I generalize the previous results to the case where the welfare weight of the very poor is strictly larger than the welfare weight on the poor, $\alpha_1 > \alpha_2 > 1$ (i.e., the planner has a desire for redistribution from the poor to the very poor). The next proposition clarifies that upwards distortions at the intensive margin can still be optimal if desire for redistribution at the bottom is sufficiently weak.

Proposition 2. *Under Assumption 1, there is a function $\beta_U : (\alpha_1, \omega_2) \mapsto \beta_U(\alpha_1, \omega_2)$ such that, at the intensive margin, optimal labor supply in skill group* 2

(i) can only be upwards distorted if $\alpha_2 > \beta_U(\alpha_1, \omega_2)$;

(ii) is upwards distorted if $\alpha_2 > \beta_U(\alpha_1, \omega_2)$ and, additionally,

$$\begin{aligned} \omega_3/\omega_2 &> \phi_\omega(\alpha_2) \; , \\ f_3 &> \phi_f(\alpha_2, \omega_3/\omega_2) \left(1 + f_1 \frac{\alpha_1 - \alpha_2}{\alpha_2 - 1}\right) \; , \end{aligned}$$

where ϕ_ω and ϕ_f are the same functions as in Proposition 1.

In case (ii)*, labor supply is upwards distorted at the extensive margin in skill groups* 1 *and* 2*. Function β_U is strictly increasing in α_1 and satisfies $\beta_U(\alpha_1, \omega_2) < \alpha_1$ if and only if $\alpha_1 > \bar{\beta}(\omega_2)$.*

Part (i) of Proposition 2 shows that the optimal allocation can only involve upwards distortions at the intensive margin if α_1 and α_2 are close enough to each other and both larger than $\bar{\beta}$, the critical value from Proposition 1 above. In particular, the formal proof shows that the first-and-half-best allocation violates the upward IC of ω_1 workers as long as α_2 is above a threshold $\beta_U(\cdot)$. If and only if $\alpha_1 > \alpha_p$, the threshold $\beta_U(\cdot)$ is smaller than α_1. In this case, the upward IC of ω_1 workers can be violated even with a weak desire for redistribution from the poor to the very poor, $\alpha_1 > \alpha_2 > \beta_U(\alpha_1, \omega_2)$.

Part (ii) of Proposition 2 generalizes the sufficient conditions for the optimality of negative marginal taxes to the case where $\alpha_1 > \alpha_2$. In particular, the condition on the high-skill productivity ω_3 is just the same as in Proposition 1, with the argument of ϕ_ω being α_2 instead of α_p. The condition on the high-skill population share f_3 uses the same threshold function ϕ_f, buts 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 and, hence, needs to collect higher taxes from the high-skilled workers. Hence, the tax schedule gets steeper, which tends to increase the cost of intensive-margin distortions due to an EITC.

4.2 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 only impose some structure on the skill set by assuming that the ratio ω_{j+1}/ω_j is equal to the same number $a > 1$ for all $j \in \{1, \ldots, n-1\}$, i.e., that the log difference between each pair of adjacent skill types is identical. Second, I do not specify the functional forms of the utility function (1) and the type distribution. In particular, I allow the fixed cost distributions G_1, G_2 etc. to vary across skill groups.

4.2.1 Assumptions

I only impose three assumptions on the labor supply elasticities at both margins. First, I put a restriction on the relative magnitudes of participation responses along the skill distribution, as measured by the semi-elasticity of participation.

Assumption 2. *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 $j \in J$.*

The empirical literature usually estimates the participation elasticity $\pi_j(c, y)$, unanimously finding that low-skill workers respond more elastically at the extensive margin than medium- and high-skill workers, both in the US and internationally (see, e.g., Juhn et al. 1991, 2002, Meghir & Phillips 2010, Bargain et al. 2014). According to the studies by Bastani et al. (2019) for Sweden and Miller et al. (2018) for the US, participation elasticities are also decreasing within the subset of low-skilled workers. In contrast, participation elasticities are flat at the bottom and only decreasing at higher income levels according to Bargain et al. (2014). Even if participation elasticities are constant along the skill distribution, however, the corresponding semi-elasticities are strictly decreasing.[15]

In general, the semi-elasticities of participation should be regarded as endogenous objects that vary with allocation (c, y). In particular, a uniform transfer to the workers in two skill groups j and $j+1$ can lead to variations in the semi-elasticities η_j and η_{j+1} as well as in their ratio $\hat{\eta}_{j,j+1} := \eta_j/\eta_{j+1}$. The following assumption imposes bounds on the variations in this ratio. Specifically, it puts a restriction on the semi-elasticity of the ratio $\hat{\eta}_{j,j+1}$ with respect to such a uniform transfer.[16]

Assumption 3. *Let $c_k = c_{k0} + c' - c_0$ for any $k \in J$. In any implementable allocation, the absolute value of the semi-elasticity of the relative participation response η_j/η_{j+1} with respect to c' is smaller than $\eta_j - \eta_{j+1}$.*

Under Assumption 3, uniform transfers to adjacent skill groups may increase or decrease the ratio of the semi-elasticities η_j/η_{j+1}, but need to have sufficiently small effects in absolute terms. While the semi-elasticity of $\hat{\eta}_{j,j+1}$ is in principle an observable object, I am not aware of any empirical results on the effects of tax reforms on participation elasticities (and their ratios). 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 3 is reasonably weak.[17]

Finally, I impose a condition that rules out degenerate labor supply responses at the intensive margin.

Assumption 4. *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 intensive-margin elasticity of income with respect to the retention rate is estimated by a large number of studies. While there is no clear consensus on the precise level of this elasticity, basically all papers estimate strictly positive elasticities, in line with the first part of Assumption 4. The second part is almost a tautology. Hence, Assumption 4

[15] The elasticity $\pi_T(\omega)$ and the semi-elasticity $\eta_T(\omega)$ of participation are related by $\eta_T(\omega) = \pi_T(\omega)/\left[y_T(\omega) - T^P(y_T(\omega))\right]$. The denominator is increasing in ω if the marginal tax T' is below 1.

[16] Formally, I define this semi-elasticity as $\varepsilon_{\hat{\eta},c} := \frac{\partial \hat{\eta}_{j,j+1}}{\partial c'} \frac{1}{\hat{\eta}_{j,j+1}}$, where $\hat{\eta}_j = \eta_j(\cdot)/\eta_{j+1}(\cdot)$.

[17] If the participation elasticities for low-skill workers with earnings \$1,000 and \$1,500 are both given by 0.5 as assumed by Saez (2002) and the US income tax is approximated as in Section 5, the ratio $\hat{\eta}_{j,j+1}$ is given by 1.5. Hence, the semi-elasticity of participation in the first group is 50% higher than in the second group. Assumption 3 requires this relative difference to remain between 32.3% and 67.7% after an (additional) uniform transfer of \$500 to both groups.

can be regarded as a weak regularity condition on the intensive-margin elasticity of labor supply.

4.2.2 Results

The following paragraphs generalize my results on the optimality of an EITC in the stylized model to the general model. For the stylized model, I have shown that, first, the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ violates the upward IC of ω_1 workers if the welfare weights at the bottom are close enough to satisfy the condition $\alpha_2 > \beta_U(\alpha_1)$, and satisfies the downward IC of ω_3 workers if the welfare weight of high-skill workers is large enough to satisfy the condition $\alpha_3 > \beta_D(\alpha_2)$. Second, I have shown that, if these conditions on the first-and-half-best are met, the optimal allocation involves upwards distortions at the intensive margin for skill group 2. Third, I have shown that the condition on the high-skill weight α_3 is satisfied if and only if both the population share f_3 and the productivity ω_3 of high-skill workers are large enough.

The derivation of sufficient conditions for the general model follows the same steps: I first clarify under which conditions allocation $(\tilde{y}, \tilde{\delta})$ violates or satisfies the local IC constraints between the workers with adjacent skill types ω_j and ω_{j+1} in Ω. Second, I identify conditions on the first-and-half-best that ensure upwards distortions in the optimal allocation in all skill groups 2 to k, i.e., negative marginal taxes at the bottom. Third, I provide a condition on the skill distribution under which the conditions in step 2 can indeed be satisfied.

For the first step, consider two adjacent skill types ω_j and ω_{j+1} with welfare weights such that $\alpha_j \geq \alpha_{j+1}$. The following lemma identifies the conditions under which the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ violates the upward IC of ω_j workers, violates the downward IC constraint of ω_{j+1} workers, or satisfies both constraints.

Lemma 3. *Let Assumptions 2 and 3 be satisfied. If a is sufficiently close to 1, the following statements are true for any skill group $j \in J/\{n\}$:*

(i) There is a strictly increasing function $\beta_{Uj} : \mathbb{R} \to \mathbb{R}$ such that $(\tilde{y}, \tilde{\delta})$ 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 $(\tilde{y}, \tilde{\delta})$ 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 $\alpha_j > \underline{\beta}_j$.

Figure 3 illustrates the statements in Lemma 3 to make them more easily accessible. The shaded area below the 45° line comprises all relevant combinations of the welfare weights with $\alpha_j \geq \alpha_{j+1}$. By Lemma 3, there are two functions β_{Uj} and β_{Dj} that partition the shaded area into three regions. In particular, by part (i), allocation $(\tilde{y}, \tilde{\delta})$ violates the upward IC of ω_j workers for all pairs of (ω_j, ω_{j+1}) in the red-shaded area above function

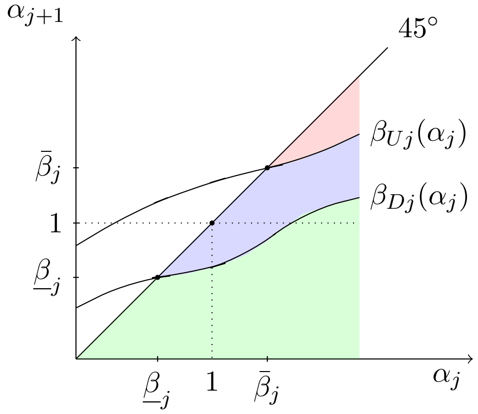

Figure 3: Local IC constraints in the first-and-half-best allocation.

β_{Uj}, where α_j and α_{j+1} are close to each other and above the threshold $\bar{\beta}_j$. As this threshold is located above the average weight of 1, the condition can only be met for pairs of low-skill groups (i.e., at the bottom of the skill distribution).

By part (ii), allocation $(\tilde{y}, \tilde{\delta})$ satisfies the downward IC of ω_{j+1} workers for all pairs of (ω_j, ω_{j+1}) in the red-shaded area and in the blue-shaded area above function β_{Dj}, where α_j and α_{j+1} are still close to each other and larger than a second threshold $\underline{\beta}_j$. As this second threshold $\underline{\beta}_j$ is located below the population average of 1, the condition can also be met for pairs of high-skill groups (i.e., at the top of the skill distribution). By contrast, allocation $(\tilde{y}, \tilde{\delta})$ only violates the downward IC constraint for (α_j, α_{j+1}) pairs in the green-shaded area below the function β_{Dj}, where either the difference between the welfare weights α_j and α_{j+1} is large or both welfare weights are low.

Lemma 3 generalizes the necessary condition in Proposition 2 (i): In the general as in the stylized model, the first-and-half-best allocation violates the upward ICs of low-skill workers if the welfare weights at the bottom are sufficiently flat. For a non-formal explanation, consider again the limit case where α_j and α_{j+1} are equal to the same number $\alpha_p > 1$. In this case, the planner wants to provide work subsidies to the workers in skill groups j and $j+1$. By Assumption 2, the semi-elasticity of participation is decreasing along the skill distribution, $\tilde{\eta}_{j+1} < \tilde{\eta}_j$ (as in the stylized model). Hence, the inverse elasticity rule implies that the planner should provide higher work subsidies to ω_{j+1} workers than to ω_j workers. Moreover, the larger α_p is, the larger should the work subsidies to both groups of workers be. In the stylized model, the difference between both work is increasing in α_p as well, because the ratio of $\tilde{\eta}_j$ and $\tilde{\eta}_{j+1}$ remains constant for all α_p. In the general model, the ratio η_j/η_{j+1} may vary with α_p. Assumption 3 ensures, however, that the variation in this ratio is small enough to ensure that the difference between the work subsidies to ω_j and ω_{j+1} workers is still growing with α_p. As a result, the first-and-half-best allocation violates the upward IC of ω_j workers if and only if the welfare weight α_p is above a unique threshold $\bar{\beta}_j > 1$. A minor difference to Proposition 2 is that Lemma 3 is restricted to cases where the distance between adjacent skill types

$a = \omega_{j+1}/\omega_j$ is sufficiently small.[18]

The second step provides conditions under which labor supply is upwards distorted at the intensive margin in the optimal (second-best) allocation. In the stylized model, labor supply in skill group 2 is upwards distorted in the optimum if the first-and-half-best allocation violates the upward IC of ω_1 workers and satisfies the downward IC of ω_3 workers. Labor supply in the two other skill groups is never upwards distorted in the optimal allocation. The following proposition shows that, in the general model, optimal labor supply can be upwards distorted at the intensive margin in many skill groups.

Proposition 3. *Let Assumptions 2 and 3 be satisfied. In the optimal allocation, labor supply is upwards distorted at both margins in 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* $j \in \{k, \ldots, n-1\}$.

Proposition 3 provides a sufficient condition for the optimal allocation to involve upwards distortions at both margins in the lowest k skill groups. To decentralize this allocation, the optimal income tax has to be given by an EITC with negative marginal and participation taxes at all income levels below y_k, the income level chosen by workers with skill type ω_k. The intuition behind Proposition 3 is the same as for the stylized model. First, if the conditions in (i) are met, the first-and-half-best allocation violates the upward ICs at the bottom. Hence, the planner finds it optimal to reduce the labor supply distortions at the extensive margin by means of negative marginal taxes below the income level y_k, even though this leads to (initially negligible) upwards distortions at the intensive margin. Second, the planner needs to set positive marginal taxes above y_k to collect taxes from the higher-skilled workers in order to finance the transfers to low-income workers. If the conditions in (ii) are met, however, these positive marginal taxes induce only small downward distortions among high-income workers.

In the stylized model, the optimality of an EITC is related to the share of high-skilled workers. Implicitly, Proposition 3 establishes a similar relation. To see this, note that that the sufficient conditions in Proposition 3 depend on the location of y_k, the endpoint of the phase-in range with negative marginal taxes. If y_k is close to y_n, the income of the highest-skilled workers, Proposition 3 verifies the optimality of an EITC with negative marginal taxes for a majority of the labor force. In this case, a small share of workers with incomes above y_k have to pay taxes to finance the net transfers to the majority of individuals with incomes below y_k. Then, the conditions in Proposition 3 are hard to satisfy. For the sake of concreteness, consider the case where $k = n - 1$. To meet the

[18] Even in the stylized model, the necessary condition gets harder to satisfy if the distance between ω_1 and ω_2 gets larger. In the general model, the upward IC and the downward IC might be satisfied given a large skill distance for all levels of α_P such that the optimal allocation can be identified using the first-order approach (see footnote 14).

conditions in part (i), the welfare weight associated to each skill group 1, 2, ..., $n-1$ has to be above some threshold $\bar{\beta}_j > 1$. Hence, the welfare weight α_n of the highest skill group has to be much lower than the average weight of 1. But then, it is impossible to satify condition in part (ii), which requires the welfare weights α_{n-1} and α_n to be close to each other (see Figure 3 and Lemma 3).

In contrast, if y_k is close to y_1, the income of the lowest-skilled workers, the EITC involves negative marginal taxes for a small share of the labor force only. In this case, a large majority of workers with incomes above y_k have to finance the net transfers to a minority of workers with very low incomes. Then, the conditions in Proposition 3 are much easier to satisfy. In particular, the conditions in (i) only require the welfare weights of a small fraction of the population to be above $\bar{\beta}_j$. If the welfare weights of all other skill groups are gradually decreasing, they can ensure an average welfare weight of 1 and satisfy the conditions in (ii) at the same time.

The following proposition confirms this intuition. In particular, it does so providing a condition on the population shares of the agents with skills above ω_k which ensures that an EICT with phase-in endpoint y_k can indeed be optimal. Thereby, it generalizes the sufficient condition in Proposition 2 (ii) in the stylized model, which also relates the optimality of an EITC to the share of high-skill agents.

Proposition 4. *Let Assumptions 2, 3 and 4 be satisfied. There are three numbers $\bar{a} > 1$, $m \in \{z \in \mathbb{N} : z \geq k+1\}$ and $\bar{z} \in (0,1)$ such that, if*

(a) $a = \frac{\omega_{j+1}}{\omega_j} < \bar{a}$,

(b) $n \geq m$ *and*

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

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

Proposition 4 provides three jointly sufficient conditions for the existence of welfare weights such that an EITC with negative marginal taxes for all incomes below a_k is optimal. Condition (a) requires the skill set to be sufficiently fine (i.e., the ratio ω_{j+1}/ω_j to be sufficiently small). This condition already appeared in Lemma 3; it ensures that a decreasing first-and-half-best tax at the bottom actually violates the upward ICs of low-skill workers. Otherwise, the optimal tax would still be decreasing at the bottom, but labor supply would not be upwards distorted at the intensive margin. Conditions (b) and (c) require that the share of high-skill workers is large enough, in line with the sufficient conditions in Propositions 1 and 2. In particular, there is a cutoff skill type $\omega_m > \omega_k$ such that the population share of workers with skill types ω_m and above must exceed some threshold $\bar{z}$.

Again, the intuition behind this sufficient condition is the same as in the stylized model with three skill types. An EITC with negative marginal 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. This intuition goes through as long as, first, agents with larger skill types indeed choose higher incomes and, second, the participation elasticity of agents with higher incomes does not grow without bounds. Both conditions are ensured by Assumption 4.

5 Numerical simulations

My theoretical analysis has shown that an EITC can be optimal, but has not provided information about the optimal size of an EITC (i.e., eligibility threshold, phase-in end-point, 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. This 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 C.5, I additionally provide a calibration for single parents, who benefit from a more generous EITC under the current US tax system.

5.1 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 re-distributive 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 using quasi-experimental estimates of extensive-margin elasticity. Additionally, I consider the studies by Chetty, Friedman & Saez (2013) on the elasticities on EITC recipients with different marital status and by Bargain et al. (2014) on the elasticities of childless singles (see Online Appendix C.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.[19]

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

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 (2002). 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 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.

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 singes in the US. Moreover, a number of studies consistently find that participation elasticities are decreasing along the skill distributions (e.g., Juhn et al. 1991, 2002, Meghir & Phillips 2010, Bargain et al. 2014 and Bastani et al. 2019). 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

$$G_j(\delta) = \frac{\exp\left(-\psi_j + \rho_j \delta\right)}{1 + \exp\left(-\psi_j + \rho_j \delta\right)} \,. \tag{23}$$

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

$$\pi_j = \bar{\pi} - (\underline{\pi} - \bar{\pi}) \left(\frac{\omega_j - \omega_1}{\omega_n - \omega_1} \right)^{1/3} \,, \tag{24}$$

where $\underline{\pi}$ and $\bar{\pi}$ are the participation elasticities in the lowest and the highest skill group, respectively. For my benchmark calibration, I assume that participation elasticities decrease from $\underline{\pi} = 0.4$ to $\bar{\pi} = 0.18$ in the group of childless singles. These parameters ensure that the average elasticity equals the preferred estimate of 0.25 in Chetty, Guren, Manoli & Weber (2013). In the sensitivity analysis, I consider both higher elasticities (decreasing from 0.5 to 0.4 as in the calibration by Jacquet et al. 2013) and lower elasticities (decreasing from 0.3 to 0.1). Additionally, I consider a model version in which the participation elasticity is equal to 0.25 in all skill groups. Similarly, I target skill-specific employment shares given by

$$L_j = \underline{L} + (\bar{L} - \underline{L}) \left(\frac{\omega_j - \omega_1}{\omega_n - \omega_1} \right)^{1/3} \,, \tag{25}$$

where $\underline{L}$ and $\bar{L}$ 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.[20] Using the OECD tax database, I approximate the US income tax in 2015 for both population groups by a linear tax function with marginal tax rate 29.3% (OECD 2017, p. 54).[21] 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 $\omega_{j+1}/\omega_j = 1.05$. Compared to most previous papers, this represents a relatively fine skill set. For example, Saez (2002) uses a model with only 17 occupation groups. In Online Appendix C.3, however, I additionally provide simulation results for models with a larger and a smaller number of skill groups (n equal to 192 and 48, respectively). The workers with the lowest and highest skill types receives annual incomes of \$500 and \$206,942, corresponding to wages of $\omega_1 = \$129$ and $\omega_n = \$13,300$ per unit of work, 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 imposed by (25) 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 two sequences of exogenous welfare weights that depend only on the agents' skill types. Graphical illustrations and details on the construction of both sequences are provided in Online Appendix C.2.

Specifically, sequence α^A assigns a constant welfare weight of 1.04 to all working agents with annual incomes below \$11,000 under the current US tax system (i.e., to the agents in the lowest 49 skill groups). The welfare weights for higher-income workers are gradually decreasing in such a way that the average welfare weight equals 1. In particular, sequence

[20]In particular, I compute for each worker his average income for each 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.

[21]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).

α^A is constructed to satisfy the sufficient conditions for an EITC with negative marginal taxes for all incomes below $\$11,000$. 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.

Sequence α^B assigns a constant welfare weight of 1.04 to the working agents with annual incomes below $\$26,000$ (i.e., to the 26.3% of agents in the lowest 63 skill groups). Furthermore, it assigns a constant welfare weight to all higher-income workers. In particular, the welfare weight of the high-skill agents is set equal to 0.986 in order to get an average welfare weight of 1. Importantly, sequence α^B satisfies the necessary conditions for an EITC with negative marginal taxes for incomes below $\$26,000$, but violates the sufficient conditions. In this case, the numerical simulation allows to determine whether the optimal tax is nevertheless given by an EITC with negative marginal taxes.

Both sequences α^A and α^B imply that the social planner has no concerns for redistribution among the working poor, i.e., from low-income earners to even-lower-income earners.[22] In the sensitivity analysis, I additionally consider a case in which the welfare weights are steeply falling at the bottom of the skill distribution (as in the simulations by Saez 2002 and Jacquet et al. (2013)).

5.2 Benchmark results

I start with the results of my numerical simulation for the benchmark calibration. Figure 4 illustrates the optimal tax schedules for weight sequences α^A and α^B. In particular, it depicts the optimal participation taxes $T_A^P(y) = T_A(y) - T_A(0)$ and $T_B^P(y) = T_B(y) - T_B(0)$ for annual incomes below $\$60,000$. To make the discreteness of the skill set transparent, I mark the simulated income-consumption bundles for workers in all skill groups by dots and triangles, respectively. As a reference, the red dotted line in Figure 4 illustrates 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.[23]

First, I report the properties of the optimal tax schedule for weight sequence α^A, which was constructed to satisfy the sufficient conditions in Proposition 3. Thus, I already know that an EITC with a negative marginal tax at all income levels below $\$11,000$ 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 solid blue line in Figure 4). Put differently, labor supply is upwards distorted at the intensive margin in the lowest 51 skill groups and at the extensive margin in the lowest 65 skill groups.

[22] In the working paper version (Hansen 2018), I show that the optimal tax has the same qualitative properties if the welfare weights are strictly decreasing with a small skill gradient at the bottom.

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

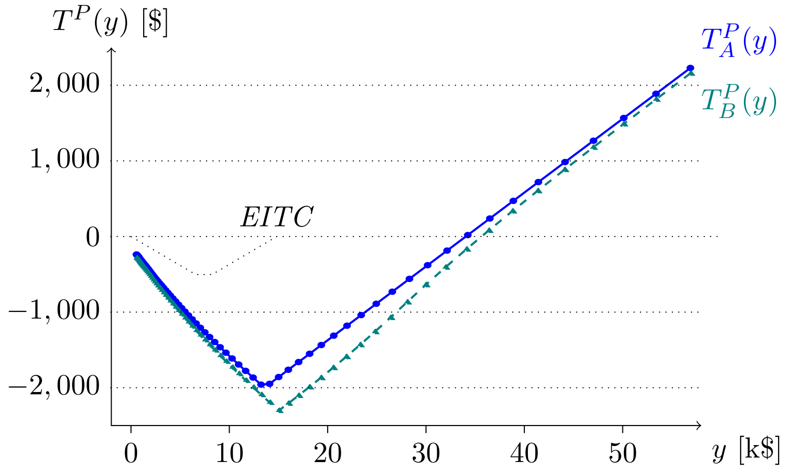

Figure 4: Optimal participation taxes for welfare weights α_A and α_B, benchmark case.

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 pre-tax incomes, which is sometimes referred to as the participation tax rate, is around −40% for very low incomes such as $y_1^A = \$555$.[24] The ratio subsequently diminishes to levels around −15% at the phase-in endpoint. The average marginal tax in the phase-in range is given −13,6%, and the average marginal tax in the phase-out range is given by 9.4%.[25]

Second, I report the simulation results for weight sequence α^B, which does not satisfy the sufficient conditions identified in Proposition 3. The simulation shows that the optimal tax is given by an EITC, nevertheless (see dashed teal line in Figure 4). In particular, the optimal tax for weight sequence α^B even involves negative marginal taxes for incomes up to $15,020 (in the lowest 53 skill groups) and negative participation taxes for incomes up to $34,100 (in the lowest 66 skill groups). In this case, 31.3% of agents benefit from participation subsidies and 24.7% of EITC recipients face negative marginal taxes.

Regarding the details of the optimal tax function, the maximum tax credit is given by $2.303, while the optimal transfer to non-working agents is equal to $1,923. The participation tax rate is below −50% at very low incomes and around −15% at the phase-in endpoint. The average marginal tax rate in the phase-in range is equal to −13.9%, and the average marginal tax in the phase-out range is given by 11.2%. The optimal marginal tax in the phase-in range is hence smaller for weight sequence α_B than for sequence α_A,

[24] 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).

[25] 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.

while the phase-in range, the EITC range and the maximum tax credit are larger.

Summing up, the numerical simulations show that the effects of the mechanism studied in this paper - the tradeoff between intensive efficiency and extensive 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. The simulations also show that an EITC can be optimal in cases where the welfare weights fail to satisfy the sufficient conditions derived in Proposition 3. Interestingly, the simulated tax schedule closely resembles the one 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).

5.3 Sensitivity analysis

To investigate the robustness of my results, I perform additional simulations in which I vary the key parameters of my calibrated model. Graphical illustrations of the optimal tax schedules for all parameter constellations are provided in Online Appendix C.3.

First, I simulate the optimal income tax for a lower intensive-margin elasticity of 0.1 and a higher elasticity of 0.5, respectively (benchmark: 0.3). In both cases, the qualitative properties of the optimal tax remain constant: In all cases, it is given by an EITC with negative marginal taxes for incomes below $\$12,500$ and negative participation taxes for incomes below $\$30,000$. The larger the intensive-margin elasticity, however, the flatter are the tax functions, both in the phase-in and in the phase-out range. For a higher elasticity, this implies both smaller marginal subsidies and a smaller maximum credit. This makes sense: A higher value of σ implies 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. Besides, a higher elasticity at the intensive margin also seems to be related to a slightly smaller phase-in endpoint (see Figure A2).

Second, I consider alternative assumptions on the average levels and skill gradients of participation elasticities. In the benchmark, I assumed participation elasticities to fall from 0.4 in the lowest skill groups to 0.18 in the highest ones. In the sensitivity analysis, I consider a case with higher participation elasticities (falling from 0.5 to 0.4 as in Jacquet et al. 2013) and a case with lower elasticities (falling from 0.3 to 0.1). Additionally, I consider a case where the participation elasticity is equal to 0.25 in all skill groups. The effects of variations in the participation elasticities on the optimal tax schedule are very limited and go in different directions. In the case with higher participation elasticities, e.g., the optimal EITC for welfare weights α^A is slightly smaller, while the optimal EITC for welfare weights α^B is slightly larger than in the benchmark case (see A3).

Third, I simulate the optimal income tax for versions of the model with a finer skill set (192 skill types) and a coarser skill set (48 skill types) than in the benchmark model with 96 skill types. Specifically, I vary the discretization of the skill set by jointly adjusting the number of skill types n and the relative distance ω_{j+1}/ω_j between adjacent skill types, while holding the level of the lowest skill type ω_1 and the highest skill type ω_n constant. Variations in the number of skill types affect the optimal tax schedule in a similar way as variations in the level of the intensive-margin elasticity: The larger the number of skill types, the flatter is the optimal income tax and the smaller is the maximum tax credit (see Figure A4).

Fourth, I simulate the optimal tax schedule for an alternative sequence α^C of welfare weights that are convexly decreasing over the skill distribution (as the welfare weights considered by Saez 2002 and Jacquet et al. 2013). These welfare weights represent a social planner with a strong desire for redistribution among low-skilled workers (i.e., from the poor to the very poor). All other parameters are kept as in the benchmark calibration. The variation of welfare weights changes the simulation results drastically: The optimal tax is given by a *Negative Income Tax* with strictly positive marginal taxes and strictly positive participation taxes at all income levels (see Figure A5).

To sum up, variations in the levels of intensive-margin and extensive-margin elasticities, in the skill gradient of participation elasticities and in the number of skill types have some effects on the quantitative properties of the optimal income tax, but not on the qualitative properties. Specifically, the simulated EITC for all parameter constellations is more than twice as larger as the current EITC for childless workers in the US. In contrast, the qualitative properties of the optimal tax schedule depend crucially on whether the skill gradient of welfare weights is small or large at the bottom, i.e., on whether the planner has a weak or a strong concern for redistribution from the poor to the very poor.

6 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 and 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 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 labor supply distortions at the intensive margin and at the extensive margin, 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. (2019), 'The Anatomy of the Extensive Margin Labor Supply Response', *Scandinavian Journal of Economics* (forthcoming).

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.

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', *The 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*.

Hansen, E. (2018), Optimal income taxation with labor supply responses at two margins: When is an Earned Income Tax Credit optimal?, Working paper. URL: *http://cmr.uni-koeln.de/de/team/junior-faculty/hansen/*.

Holtz-Eakin, D., Joulfaian, D. & Rosen, H. S. (1993), 'The Carnegie Conjecture: Some empirical evidence', *The 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.

Lockwood, B. B. (2017), Optimal income taxation with present bias, Working paper. URL: *https://bepp.wharton.upenn.edu/profile/benlo/*.

Lorenz, N. & Sachs, D. (2012), Optimal participation taxes and efficient transfer phase-out, Working Paper Series of the Department of Economics, University of Konstanz 2012-37. URL: *http://ideas.repec.org/p/knz/dpteco/1237.html*.

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.

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, Technical report. 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 Series No. 21211. URL: *http://www.nber.org/papers/w21211*.

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

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', *The 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.

Appendix

A Proofs for Sections 3 and 4.1

The following preliminary result clarifies that every implementable and second-best Pareto efficient allocation involves the same kind of pooling across two-dimensional types. It is well-known, if all welfare weights are strictly positive, the optimal allocation is second-best Pareto efficient. Hence, I can restrict the following analysis to a specific class of pooling allocations.

Lemma 4. *Every implementable and second-best Pareto efficient allocation is characterized by two functions $c : \Omega \to \mathbb{R}$ and $y : \Omega \to \mathbb{R}$ and a number c_0 such that,*

- *all agents with skill type $\omega \in \Omega$ and fixed cost type $\delta \leq \hat{\delta}(\omega) := c(\omega) - h(y(\omega), \omega) - c_0$ receive bundle $(c(\omega), y(\omega))$, and*
- *all other agents receive bundle $(c_0, 0)$.*

Proof. Consider two types (ω^A, δ^A) and (ω^B, δ^B) in $\Omega \times \Delta$ and assume that type $k \in \{A, B\}$ is allocated bundle (c^k, y^k). In the following, I study the implications of incentive compatibility (9) and Pareto efficiency for (c^A, y^A) and (c^B, y^B). Pareto-efficiency requires that (c^k, y^k) minimizes the net transfer $c^k - y^k$ over the set of bundles (c', y') such that, first, $u(c', y', \omega^k, \delta^k) \geq u(c^k, y^k, \omega^k, \delta^k)$ and, second, (c', y') satisfies (9). Otherwise, the government can replace (c^k, y^k) by some other bundle (c', y') that makes agent k better off, but frees up resources that can be redistributed uniformly to all agents in the population.

First, if $y^A = y^B = 0$, then (9) requires that both types receive the same consumption $c^A = c^B$. I henceforth denote the consumption level of all non-working agents by c_0.

Second, if both y^A and y^B are strictly positive and $\omega^A = \omega^B = \omega'$, then (9) directly requires that there is a number z' such that $c^A - h(y^A, \omega') = c^B - h(y^B, \omega') = z'$. Second-best Pareto efficiency requires that, more specifically, $(c^A, y^A) = (c^B, y^B)$. To prove the last statement by contradiction, assume that $y^B > y^A > 0$ and that $c^k = z' + h(y^k, \omega')$ for each $k \in \{A, B\}$. Consider further a bundle (c', y') such that $y' \in (y^A, y^B)$ and $c' = z' + h(y', \omega')$. Note that both types A and B are indifferent between the three bundles (c^A, y^A), (c^B, y^B) and (c', y'). If the government provides an agent with bundle (c', y'), this involves a net resource transfer of $c' - y' = z' + h(y', \omega') - y'$. Due to $h_{yy} > 0$, this net transfer is strictly convex in y' and has a unique minimizer on $\left[y^A, y^B\right]$. Consequently, at least one of the bundles (c^A, y^A) or (c^B, y^B) does not satisfy the Pareto condition explained above.

Third, assume that both agents have the same skill type, $\omega^A = \omega^B = \omega'$, but $y^A > 0$ and $y^B = 0$. Then, (9) requires that $c_B = c_0$ and

$$
\begin{aligned}
c^A - h(y^A, \omega') - \delta^B \leq c_0 \leq c^A - h(y^A, \omega') - \delta^A \\
\Leftrightarrow \delta^A \leq c^A - h(y^A, \omega') - c_0 \leq \delta^B \ .
\end{aligned}
$$

Put differently, an agent with skill type ω' provides positive output if and only if his fixed cost type δ is below the threshold $\hat{\delta}(\omega') = c^A - h(y^A, \omega') - c_0$, where (c^A, y^A) is the bundle allocated to all workers with skill type ω'. $\square$

Proof of Lemma 1

Proof. By Lemma 4, welfare (3) can be rewritten as

$$\begin{aligned} E_\omega\left[\alpha(\omega)\left\{\int_{\underline{\delta}}^{\hat{\delta}(\omega)}\left[c(\omega)-h\left(y(\omega),\omega\right)-\delta\right]dG_\delta(\delta\mid\omega)+\left[1-G_\delta(\delta\mid\omega)\right]c_0\right\}\right] &= \\ E_\omega\left[\alpha(\omega)\left\{\int_{\underline{\delta}}^{\hat{\delta}(\omega)}\left[c_0+\hat{\delta}(\omega)-\delta\right]dG_\delta(\delta\mid\omega)+\left[1-G_\delta(\delta\mid\omega)\right]c_0\right\}\right] &= \\ c_0+E_\omega\left[\alpha(\omega)\left\{G_\delta(\hat{\delta}(\omega)\mid\omega)\hat{\delta}(\omega)-\int_{\underline{\delta}}^{\hat{\delta}(\omega)}\delta dG_\delta(\delta\mid\omega)\right\}\right] &, \qquad (26) \end{aligned}$$

where the expectation is taken over $\omega\in\Omega$. For the last equality, recall that the average welfare weight equals $E_\omega[\alpha(\omega)]=1$. Lemma 4 also allows to rewrite the feasilibity condition (8) as

$$\begin{aligned} E_\omega\left[G_\delta(\hat{\delta}(\omega)\mid\omega)\left(y(\omega)-c(\omega)\right)-\left(1-G_\delta(\hat{\delta}(\omega)\mid\omega)\right)c_0\right] &\geq 0 \\ \Leftrightarrow E_\omega\left[G_\delta(\hat{\delta}(\omega)\mid\omega)\left(y(\omega)-h\left[y(\omega),\omega\right]-\hat{\delta}(\omega)\right)\right] &\geq c_0\,, \qquad (27) \end{aligned}$$

using that the participation threshold is given by $\hat{\delta}(\omega)=c(\omega)-h(y(\omega),\omega)-c_0$. In the optimal allocation, (27) holds with equality. Hence, we can use (27) to replace c_0 in expression (26). This directly gives the objective (10) in Lemma 1, which already takes into account incentive compatibility among any pair of agents with identical skill types and among any pair of non-working agents.

I can also use the threshold condition $\hat{\delta}(\omega)=c(\omega)-h(y(\omega),\omega)-c_0$ to rewrite the incentive compatibility constraint between any pair of working agents with skill types ω' and ω'' as

$$\begin{aligned} c(\omega')-h\left[y(\omega'),\omega'\right] &\geq c(\omega'')-h\left[y(\omega''),\omega'\right] \\ \Leftrightarrow \hat{\delta}(\omega') &\geq \hat{\delta}(\omega'')+h\left[y(\omega''),\omega''\right]-h\left[y(\omega''),\omega'\right]\,, \end{aligned}$$

which is constraint (11) in Lemma 1. This condition also ensures incentive compatibility between any non-working agent with type (ω',δ') and any working agent with skill type (ω'',δ'') because $c_0\geq c(\omega')-h\left[y(\omega'),\omega'\right]-\delta'\geq c(\omega'')-h\left[y(\omega''),\omega'\right]-\delta'$. □

Proof of Lemma 2

Proof. Lemma 1 implies that $(\tilde{y},\tilde{\delta})$ can be derived by point-wise maximization of (10) with respect to $y(\omega)$ and $\hat{\delta}(\omega)$. The first-order condition with respect to $y(\omega)$ requires that $\tilde{y}(\omega)$ satisfies (12) and represents the unique maximizer of $y-h(y,\omega)$.

The first-order condition with respect to $\hat{\delta}(\omega)$ requires that

$$\alpha(\omega)-1+\eta(\tilde{\delta}(\omega)\mid\omega)\left[\tilde{y}(\omega)-h\left(\tilde{y}(\omega),\omega\right)-\tilde{\delta}(\omega)\right]\stackrel{!}{=}0\,, \qquad (28)$$

where $\eta(\cdot\mid\omega):=g_\delta(\cdot\mid\omega)/G_\delta(\cdot\mid\omega)$ denotes the semi-elasticity of participation among agents with skill type ω. Rearranging this equation gives (13). □

Proofs of Propositions 1 and 2

I proof Propositions 1 and 2 in two steps. First, I provide conditions under which the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ satisfies or violates each of the local IC constraints in the general case where $\alpha_1 \geq \alpha_2$. The conditions for the special case where $\alpha_1 = \alpha_2$ follow as a corollary. Second, I show that the labor supply distortions in the optimal (second-best) allocation can be determined based on the status of the IC constraints in the first-and-half-best allocation.

Lemma 5. *Let Assumption 1 hold and $\Omega = \{\omega_1, \omega_2, \omega_3\}$. Define $a_3 = \omega_3/\omega_2$. The first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ satisfies*

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

$$\alpha_2 \leq \beta_U(\alpha_1, \omega_2) := 2 - \frac{\omega_2^{1+\sigma}}{(2-\alpha_1)^{-1} + \sigma\omega_2^{1+\sigma}(\omega_2^{1+1/\sigma} - 1)} , \tag{29}$$

where $\beta_U(\alpha_1, \omega_2) < \alpha_1$ if and only if $\alpha_1 > \bar{\beta}(\omega_2) := 2 - \frac{a_2^{1+\sigma}-1}{\sigma a_2^{1+\sigma}\left(a_2^{1+1/\sigma}-1\right)} \in (1, 2)$.

(ii) the upward IC constraint of ω_2 workers whenever $\alpha_1 \geq \alpha_2 \geq \alpha_3$;

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

$$\alpha_3 \geq \beta_D(\alpha_2, a_3) := 2 - \frac{a_3^{1+\sigma}}{(2-\alpha_2)^{-1} + \sigma(1 - a_3^{-1-1/\sigma})} . \tag{30}$$

There is a number $\phi_\omega(\alpha_2) > 1$ such that condition (30) *is met if and only if $a_3 > \phi_\omega(\alpha_2)$ and*

$$f_3 \quad > \quad \frac{\alpha_2 - 1 + f_1(\alpha_1 - \alpha_2)}{\alpha_2 - \beta_D(\alpha_2, a_3)} . \tag{31}$$

Proof. I start by showing that the first-half-best allocation is defined in closed form by the equations (18) and (19). To see this, note that Lemma 2 provides implicit definitions of $\tilde{y}_j$ and $\tilde{\delta}_j$ for any $j \in \{1, 2, 3\}$. Under Assumption 1, the skill-specific 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 $\tilde{y}_j - h(\tilde{y}_j, \omega_j) = \omega_j^{1+\sigma}/(1+\sigma)$. Moreover, with uniformly distributed fixed costs, the semi-elasticity $\tilde{\eta}_j = g_j(\tilde{\delta}_j)/G_j(\tilde{\delta}_j)$ equals $\tilde{\delta}_j^{-1}$ as long as $\tilde{\delta}_j < \bar{\delta}$. Plugging these expressions into the implicit definition (28) of $\tilde{\delta}_j$ and solving for $\tilde{\delta}_j$ gives $\tilde{\delta}_j = \frac{\omega_j^{1+\sigma}}{(1+\sigma)(2-\alpha_j)}$, which equals (19). We are thus equipped to verify incentive compatibility in the first-and-half-best allocation.

First, the upward IC constraint of ω_1 workers is given by inequality (11) with $\omega' = \omega_1$ and $\omega'' = \omega_2$. Using the closed-form expressions (18) and (19), the constraint is satisfied in the first-and-half-best allocation if and only if

$$\begin{aligned} \frac{\omega_1^{1+\sigma}}{(1+\sigma)(2-\alpha_1)} - \frac{\omega_2^{1+\sigma}}{(1+\sigma)(2-\alpha_2)} \quad &\leq \quad \frac{\sigma}{1+\sigma}\omega_2^{1+\sigma}\left[1 - \left(\frac{\omega_2}{\omega_1}\right)^{1+1/\sigma}\right] \\ \Leftrightarrow \alpha_2 \leq \beta_U(\alpha_1, \omega_2) \quad &:= \quad 2 - \frac{\omega_2^{1+\sigma}}{(2-\alpha_1)^{-1} + \sigma\omega_2^{1+\sigma}(\omega_2^{1+1/\sigma} - 1)} . \end{aligned}$$

It is straightforward to show that $\beta_U(\alpha_1, \omega_2) > \alpha_1$ if and only if α_1 is larger than $\bar{\beta}(\omega_2)$, the unique fixed point of $\beta_U(\ , \omega_2)$. For the limit case where $\alpha_1 = \alpha_2 = \alpha_p$, this implies that the upward IC of ω_1 workers is satisfied if and only if $\alpha_p \leq \bar{\beta}(\omega_2)$. It remains to show that $\bar{\beta}(\omega_2)$ is between 1 and 2. It is easy to see that $\bar{\beta}(\omega_2)$ is strictly smaller than 2 if $\omega_2 > 1$ and $\sigma > 0$. Finally, for $\alpha_j = 1$, $\tilde{\delta}_j$ equals $\tilde{y}_j - h(\tilde{y}_j, \omega_j)$ by Lemma 2. For $\alpha_1 = \alpha_2 = 1$, hence, $(\tilde{y}, \tilde{\delta})$ satisfies the upward IC of ω_1 with equality,

$$\begin{aligned} \tilde{\delta}_1 - \tilde{\delta}_2 = \tilde{y}_1 - h(\tilde{y}_1, \omega_1) - \tilde{y}_2 + h(\tilde{y}_2, \omega_2) \quad &> \quad h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_1) \\ \Leftrightarrow \tilde{y}_1 - h(\tilde{y}_1, \omega_1) \quad &> \quad \tilde{y}_2 - h(\tilde{y}_2, \omega_1) , \end{aligned}$$

because $\tilde{y}_j$ is by construction the unique maximizer of $y - h(y, \omega_j)$. Hence, $\bar{\beta}(\omega_2)$ must be located in the interval $(1, 2)$.

Second, it is straightforward to show that allocation $(\tilde{y}, \tilde{\delta})$ violates the upward IC constraint of ω_2 workers if and only if $\alpha_3 > \beta_U(\alpha_2, \omega_3/\omega_2)$, where β_U is the function defined in part (i). This condition requires that either $\alpha_3 > \alpha_2$ or $\alpha_3 > \bar{\beta}(\omega_3) \in (1, 2)$. With $\alpha_1 \geq \alpha_2 \geq \alpha_3$ and an average welfare weight of 1, α_3 is below 1: the condition $\alpha_3 > \beta_U(\alpha_2, \omega_3/\omega_2)$ cannot be satisfied.

Third, the downward IC constraint of ω_3 workers is given by inequality (11) with $\omega' = \omega_3$ and $\omega'' = \omega_2$. Using again the closed-form expressions (18) and (19), the constraint is satisfied if and only if

$$\begin{aligned} \frac{\omega_3^{1+\sigma}}{(1+\sigma)(2-\alpha_3)} - \frac{\omega_2^{1+\sigma}}{(1+\sigma)(2-\alpha_p)} \quad &\geq \quad \frac{\sigma}{1+\sigma} \omega_2^{1+\sigma} \left[1 - \left(\frac{\omega_2}{\omega_3} \right)^{1+1/\sigma} \right] \\ \Leftrightarrow \alpha_3 \geq \beta_D(\alpha_p, a_3, \sigma) \quad &:= \quad 2 - \frac{a_3^{1+\sigma}}{(2-\alpha_p)^{-1} + \sigma \left(1 - a_3^{-1-1/\sigma} \right)} . \end{aligned}$$

Finally, by the normalization of welfare weights, the welfare weight of ω_3 workers satisfies

$$\begin{aligned} f_1 \alpha_1 + (1 - f_1 - f_3) \alpha_2 + f_3 \alpha_3 = 1 \\ \Leftrightarrow \alpha_3 = \alpha_2 - \frac{\alpha_2 - 1 + f_1(\alpha_1 - \alpha_2)}{f_3} . \end{aligned}$$

This implies that the condition $\alpha_3 \geq \beta_D(\alpha_2, a_3)$ is equivalent to inequality (31) in Lemma 5. However, as the share f_3 must be smaller than $1 - f_1$, inequality (31) requires that

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

For any $\sigma > 0$, $\alpha_2 \in (1, \alpha_1]$ and $\alpha_1 < 2$, $Z(1)$ is smaller than $-(\alpha_1 - 1) f_1 / (1 - f_1) \leq 0$. Moreover, Z is strictly increasing in a_3 if and only if $a_3^{2+\sigma+1/\sigma} > 1 + (\alpha_1 - 1) \frac{f_1}{1-f_1} \geq 1$. For $a_3 \to \infty$, $Z(a_3)$ converges to infinity. Hence, the threshold $\phi_\omega(\alpha_2)$ is uniquely defined for any combination of α_1, α_2, σ and $f_1 < 1$. □

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 $\alpha_1 = \alpha_2$ (i.e., the social planner has no concerns for redistribution at the bottom). As they follow directly from the conditions in Lemma 5, I state the corollary without further proof.

Corollary 1. *Let Assumption 1 hold, $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $\alpha_1 = \alpha_2\alpha_p > 1$. Define $a_3 = \omega_3/\omega_2$. The first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ satisfies*

(i) the upward IC constraint of ω_1 workers if and only if $\alpha_p > \bar{\beta}(\omega_2)$;

(ii) the downward IC constraint of ω_3 workers if and only if $f_3 > (\alpha_p - 1)/[\alpha_p - \beta_D(\alpha_p, a_3)]$ and $a_3 > \phi_\omega(\alpha_p)$.

and satisfies the downward IC of ω_3 workers.

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

Proof. The optimal allocation maximizes the objective (10) subject to the subset of the local IC constraints 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. In the following, I study which IC constraints can be binding if one of the conditions $\alpha_1 = \alpha_2 \leq \bar{\beta}(\omega_2)$ and $\alpha_2 \leq \beta_U(\alpha_1, \omega_2)$ holds. By Lemma 5, these conditions ensure that the first-and-half-best allocation $(\tilde{y}, \tilde{\delta})$ does not violate the upward IC of ω_1 workers. Besides, it never violates the upward IC of ω_2 workers.

There remain three possible cases. First, if $(\tilde{y}, \tilde{\delta})$ 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 $(\tilde{y}, \tilde{\delta})$ 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 $(\tilde{y}, \tilde{\delta})$ only violates the downward IC of ω_3 workers, this IC is unambiguously binding in the optimal allocation. In this case, the upward IC of ω_1 workers may also be binding, but y_2* cannot be upwards distorted nevertheless. Using Lemma 1, the Lagrangian for this case is given by

$$
\begin{aligned}
\mathcal{L} \;=\; & \sum_{j=1}^{3} f_j \left[G_j(\hat{\delta}_j) \left(y_j - h\left[y_j, \omega_j\right)\right] + \hat{\delta}_j \left[\alpha_j - 1\right] \right) - \alpha_j \int_{\underline{\delta}}^{\hat{\delta}_j} \delta dG_j(\delta) \right] \\
& + \mu_1^U \left[\hat{\delta}_1 - \hat{\delta}_2 - h(y_2, \omega_2) + h(y_2, \omega_1) \right] \\
& + \mu_3^D \left[\hat{\delta}_3 - \hat{\delta}_2 - h(y_2, \omega_2) + h(y_2, \omega_3) \right] \;, \qquad (32)
\end{aligned}
$$

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. As $G_j(\delta) = \delta/\bar{\delta}$ and $g_j(\delta)/G_j(\delta) = \delta$, the first-order conditions with respect to $\hat{\delta}_2$ and $\hat{\delta}_1$ require that

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

where the inequalities follow because $\mu_1^U > 0$, $\mu_3^D > 0$, $\tilde{y}_2 = \arg\max y - h(y, \omega_2)$ and $\tilde{\delta}_j$ is defined by (13). As the downward IC of ω_3 workers is binding by assumption, we also know that

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

The assumption of part (i) ensures that $(\tilde{y}, \tilde{\delta})$ satisfies the upward IC of ω_1, i.e., $\tilde{\delta}_1 - \tilde{\delta}_2 \geq h(\tilde{y}_2, \omega_2) - h(\tilde{y}_2, \omega_1)$. Hence, we must 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 in this case. □

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

Proof. By Lemma 5, the conditions $\alpha_1 = \alpha_2 > \bar{\beta}(\omega_2)$ and $\alpha_2 > \beta_U(\alpha_1, \omega_2)$ ensure that the first-and-half-best allocation violates the upward IC of ω_1 workers. The conditions on ω_3 and f_3 ensure that the first-and-half-best allocation satisfies the downward IC of ω_3 workers. Recall that the upward IC of ω_2 workers is generally satisfied.

These conditions jointly ensure that the upward IC of ω_1 workers is binding in the second-best allocation. With respect to the other IC constraints, there are two possible cases: (a) no other IC is binding, (b) the downward IC of ω_3 workers is binding. 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. Still, it can arise for some parameter constellations. Using Lemma 1, the Lagrangian for this case is again given by equation (32). For $j \in \{1, 3\}$, 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. As $G_j(\delta) = \delta/\bar{\delta}$ and $g_j(\delta)/G_j(\delta) = \delta$, 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 $\tilde{\delta}_j$ is defined by (13). As the downward IC of ω_3 workers is binding by assumption, we 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 $(\tilde{y}, \tilde{\delta})$ 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, labor supply in skill group y_2 is unambiguously upwards distorted in case (b).

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, the participation threshold of ω_2 workers satisfies $\delta_2^* = \delta_1^* + h(y_2^*, \omega_1)$ because the upward IC of ω_1 workers is binding. Note further that $\delta_1^* > y_1^* - h(y_1^*, \omega_1) > y_2^* - h(y_2^*, \omega_1)$, where the second inequality holds because $y_1^* = \tilde{y}_1$ is the unique maximizer of $y - h(y, \omega_1)$, while y_2^* is upwards distorted and, hence, larger than $\tilde{y}_2 > \tilde{y}_1$ as shown above. Combining these inequalities, we find that $\delta_2^* > y_2^* - h(y_2^*, \omega_2)$. Hence, labor supply by ω_2 workers is upwards distorted at the extensive margin as well. □