

Pareto-improving tax reforms and the Earned Income Tax Credit*

Felix J. Bierbrauer[†] Pierre C. Boyer[‡] Emanuel Hansen[§]

September 24, 2022

Abstract

We develop a new approach for the identification of Pareto-improving tax reforms. This approach yields necessary and sufficient conditions for the existence of Pareto-improving reform directions. A main insight is that “Two brackets are enough”: When the system cannot be improved by altering tax rates in one or two income brackets, then there is no continuous reform direction that is Pareto-improving. We also show how to check whether a given tax reform is Pareto-improving. We use these tools to study the introduction of the Earned Income Tax Credit (EITC) in the US in 1975. A robust finding is that, prior to the EITC, the US tax-transfer system was not Pareto-efficient. Under plausible assumptions about behavioral responses, the 1975 reform was not Pareto-improving. Qualitatively, though, it had the right properties: A similar reform with earnings subsidies made available to a broader range of incomes would have been Pareto-improving.

Keywords: Tax reforms; Non-linear income taxation; Optimal taxation; Earned Income Tax Credits; Pareto efficiency.

JEL classification: C72; D72; D82; H21.

*We are grateful to the editor, Alessandro Lizzeri, and three anonymous referees for many insightful comments and suggestions. We thank Thomas Gaube, Ulrich Glogowsky, Roger Guesnerie, Bas Jacobs, Laurence Jacquet, Etienne Lehmann, Jonas Löbbing, Andreas Peichl, Dominik Sachs, Florian Scheuer, Laurent Simula, Alain Trannoy, Nicolas Werquin and seminar audiences in Frankfurt, Brussels, Marseille, Lyon, Cologne, Munich, Paris, Florence, and Osnabrück for helpful comments. Susanna Bolz and Marvin Immesberger provided excellent research assistance. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC2126/1 – 390838866, the Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047), ANR-19-CE41-0011-01 and ANR-20-CE41-0013-01 are gratefully acknowledged.

[†]CMR - Center for Macroeconomic Research, University of Cologne, Albert-Magnus-Platz, 50923 Köln, Germany. E-mail: bierbrauer@wiso.uni-koeln.de

[‡]CREST, École polytechnique, Institut Polytechnique de Paris, 5, avenue Henry Le Chatelier, 91120 Palaiseau, France. E-mail: pierre.boyer@polytechnique.edu

[§]Center for Economic Studies, Ludwig-Maximilians-Universität Munich, Schackstr. 4, 80539 Munich, Germany, and CMR, University of Cologne, Albert-Magnus-Platz, 50923 Köln, Germany. E-mail: emanuel.hansen@econ.lmu.de

1 Introduction

We develop a new approach for the identification of Pareto-improving tax reforms. We prove two theorems that, respectively, give necessary and sufficient conditions for the existence of Pareto-improving reform directions. Our approach also enables us to assess whether a given historical or hypothetical reform direction is Pareto-improving. We, moreover, provide a characterization of reform directions that are optimal in the sense of yielding the maximal “free lunch”, defined as the tax revenue in excess of what is needed to make sure that no one is made worse off. This approach does not only tell whether *something* should be done by answering the yes-or-no question about the existence of a Pareto-improving direction for reform. It also says *what* can be done to improve the system: It identifies the relevant brackets and says for each such bracket whether marginal tax rates should go up or down. Making use of our results in applications is straightforward when sufficient-statistics formulas for the revenue implications of tax reforms are available. We therefore foresee a wide range of potential applications.

We use this approach in an empirical analysis of an important US tax reform: the introduction of the Earned Income Tax Credit (EITC), a system of earnings subsidies for low-income households with dependent children in 1975. It involved changes of marginal tax rates in two brackets of incomes. Marginal tax rates were reduced in the lower bracket, the phase-in range of the EITC. They were increased in the higher bracket, the phase-out range. We find that the US tax-transfer system was not Pareto-efficient prior to the introduction of the EITC. Under plausible assumptions about behavioral responses, the introduction of the EITC was not Pareto-improving. Qualitatively, however, it had the right properties. The optimal reform would also have introduced earnings subsidies, but would have made them available to a broader range of incomes. That is, the actual reform had phase-in and phase-out ranges that were too narrow.

A theory of Pareto-improving tax reforms. Our theoretical analysis is motivated by two observations. First, reforms of the EITC typically involved *two*

brackets, a phase-in range with lower marginal tax rates and a phase-out range with higher tax rates. Second, the literature that uses perturbation methods for a characterization of optimal tax systems focuses on reforms that affect marginal tax rates in *one bracket*: Optimality requires that no such reform can raise social welfare.¹ This contrast between one-bracket and two-bracket reforms spurs a more general question: If one seeks to improve the tax system, how many brackets does one actually need? To get at this question, we study reforms that involve changes of marginal tax rates in an arbitrary number of income brackets.

We obtain the following insights: Suppose that a given tax system is “one-bracket-efficient” in the sense that there does not exist a Pareto-improving one-bracket reform. Can there be reforms with two brackets that make everyone better off? We show that the answer is “yes”, i.e., there are Pareto improvements that require a second bracket. Now, suppose that the scope for Pareto-improving two-bracket reforms has been exhausted. Can there be reforms with three or even more brackets that make everyone better off? We show that the answer is “no”, i.e., if there is no Pareto-improving reform involving one or two brackets, then there is no Pareto-improving element in the set of continuous reform directions. Broadly summarized: “Two brackets are enough.” When there is nothing in the special classes of one- or two-bracket reforms, then there is nothing at all.

These findings are derived from a generic static model of income taxation: Individuals derive utility from consumption and the generation of income requires costly effort. They face a budget constraint that is shaped by a predetermined non-linear tax-transfer system. We focus on reforms that are “small” in the sense that tax rates stay close to what they are in the status quo. The range of incomes subject to tax rate changes is unrestricted; it can be “large”.

Detecting inefficiencies. How can one figure out whether a given tax-transfer system admits a Pareto-improving reform? Our answer involves an object that we refer to as the *revenue function*. This function assigns to every level of income y

¹See Piketty (1997), Saez (2001), Golosov, Tsyvinski and Werquin (2014), or Jacquet and Lehmann (2021).

the additional tax revenue $\mathcal{R}(y)$ that becomes available when marginal tax rates are raised in a small bracket containing that income level, thereby increasing the tax burden at all higher levels of income. The test for Pareto efficiency then makes use of the following insights:

1. There is no Pareto-improving one-bracket reform if and only if the function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0 and from above by 1. These bounds admit an interpretation as Laffer conditions that, respectively, indicate whether marginal tax rates are inefficiently high or inefficiently low: When $\mathcal{R}(y) < 0$, a one-bracket tax cut for incomes close to y is Pareto-improving. By contrast, when $\mathcal{R}(y) > 1$, then raising tax rates yields so much revenue that every one can be made better off, even those who face a higher tax burden after the reform.
2. There is no Pareto-improving two-bracket reform if and only if the function $y \mapsto \mathcal{R}(y)$ is non-increasing. A violation of this monotonicity condition implies that the tax system can be Pareto-improved by an EITC-like reform – i.e., a two-bracket reform with a phase-in and a phase-out range.

Evaluating tax reforms. When a tax system is found to be inefficient, this raises the question what to do about it. Typically, there is a set of Pareto-improving directions for reform. We provide a formula that allows to check, for any given continuous reform direction, whether it belongs to this set. This test requires the same information as the test for the Pareto efficiency of a tax system sketched above: An estimate of the revenue function $y \mapsto \mathcal{R}(y)$ needs to be available.

Optimal reforms – Quantifying inefficiencies. Can one Pareto-improving direction be singled out as the optimal one? Generally, what an optimal selection from the set of Pareto improvements looks like cannot be answered without specifying an objective function. Intuitively, however, one often thinks of a Pareto improvement as a “free lunch”, a gain that is available to every one. A money-metric

measure of the size of the free lunch is the tax revenue that can be redistributed lump sum after those whose tax burden goes up (if any) have been compensated. We provide a characterization of the reform direction that is optimal in the sense of yielding the maximal free lunch. In particular, when the revenue function is known, we can pin down the optimal brackets for tax rate cuts and hikes.

The value function of this optimization problem – i.e., the maximum free lunch that the status quo tax system leaves on the table – can be interpreted as a measure of how inefficient a given tax system is. With this measure, one can study whether inefficiencies increase or decrease over time, or whether the inefficiencies for one subgroup of the population are larger or smaller than those for another subgroup.

The introduction of the EITC. We use these tools to analyze the introduction of the EITC for low-income households with dependent children in the US. Its introduction in 1975 was a substantial policy change for many low-income households, see, e.g., Bastian (2020). It was meant as a response to excessively high marginal tax rates for families that depended on welfare. Specifically, the reform decreased marginal taxes by 10 percentage points in the phase-in range between 0 and 4,000 USD, and increased them by 10 percentage points in the adjacent phase-out range between 4,000 and 8,000 USD.

Our analysis of this reform is based on a model with behavioral responses at the intensive and the extensive margin. We calibrate the corresponding revenue function using empirical estimates of labor supply elasticities and the income distribution among single parents, as well as information on the tax-and-transfer schedule. We first check whether the US tax-transfer system for single parents was Pareto-efficient prior to the introduction of the EITC. We find that it was not: Specifically, our test indicates that there were Pareto-improving reforms both with one and with two brackets.² In the second step, we investigate whether the 1975 EITC reform had a Pareto-improving direction. For empirically plausible elasticities, the answer is “no”. In a third step, we show that the optimal reform – i.e., the reform yielding the maximal “free lunch” – was qualitatively similar

²By contrast, the tax-transfer system for childless singles only involved negligible inefficiencies.

to the 1975 reform. Specifically, the optimal reform also involved lower marginal tax rates in a phase-in range and higher marginal tax rates in a phase-out range, but it covered a wider range of incomes; plausibly, annual incomes up to 10,250 USD under the optimal reform versus annual incomes up to 8,000 USD under the actual reform. We show that these conclusions survive various robustness checks involving alternative elasticity estimates, alternative estimates of the income distribution, or alternative representations of the US tax-transfer schedule system for single parents.

In an extension of this analysis, we look into subsequent reforms of the EITC. We show that the 1979 reform, which expanded the range of incomes covered by the EITC, had a Pareto-improving direction. We also find that the introduction of more generous EITC schedules for taxpayers with more children (as implemented in the 1990s and 2000s) would have allowed to realize additional efficiency gains already in the mid-1970s. We conclude that both the introduction of an EITC and its subsequent expansion can, through the lens of our framework, be rationalized as efficiency-enhancing reforms.

Related literature on Pareto-efficient income taxation. We build on a previous literature that has identified distinct necessary conditions for the Pareto efficiency of a tax system. A contribution of this paper is to provide a unified treatment by making use of the revenue function $y \mapsto \mathcal{R}(y)$. One branch of the literature has generalized the notion of a Laffer bound, i.e., of an upper bound for marginal tax rates, to non-linear tax schedules, see Stiglitz (1982), Brito, Hamilton, Slutsky and Stiglitz (1990) and, more recently, Badel and Huggett (2017).³ This corresponds to the condition in this paper that the revenue function $y \mapsto \mathcal{R}(y)$ must be non-negative. Bierbrauer, Boyer and Peichl (2021) show that there is also a lower Pareto bound for marginal tax rates. As we show here, this corresponds to the condition that the revenue function $y \mapsto \mathcal{R}(y)$ must be bounded from above

³There is a literature deriving the second-best Pareto frontier for a two-type Mirrlees model with contributions by Stantcheva (2014) and Bastani, Blomquist and Micheletto (2020). See, for reviews, Stiglitz (1987) and Boadway and Keen (2000).

by 1. Werning (2007) and Lorenz and Sachs (2016) develop a test for Pareto efficiency. We find that this test is essentially equivalent to the condition that the revenue function $y \mapsto \mathcal{R}(y)$ must be non-increasing.⁴

We advance this literature by showing that, taken together, all these conditions on the revenue function are sufficient for the non-existence of a Pareto-improving reform direction. The previous literature on sufficient conditions for Pareto efficiency is scarce. An exception is Werning (2007). He analyses a *Pareto problem*, an allocation problem where net resources are maximized subject to reservation utilities that stem from a status quo tax policy. The analysis invokes regularity conditions to ensure that first-order conditions are necessary and sufficient for Pareto efficiency.⁵ Thus, Werning’s analysis goes beyond necessary conditions for Pareto efficiency. This said, the regularity conditions merely serve as background assumptions that justify the focus on first-order conditions. Relative to this benchmark, our approach is more parsimonious in that there is no need to take these regularity conditions on board. The revenue function $y \mapsto \mathcal{R}(y)$ exists under weaker conditions. Moreover, it can be used for a direct test of whether a given tax system satisfies the sufficient conditions for Pareto efficiency. One simply needs to check whether the revenue function is bounded and non-increasing.

Our paper is also related to a literature on the *inverse tax problem*.⁶ This literature takes a status quo tax system as given and then asks whether there is a social welfare function for which this system appears to be welfare-maximizing.

⁴Lorenz and Sachs (2016) extend Werning (2007) by considering extensive margin responses. For related work, see also Blundell and Shephard (2012), Scheuer (2014), Koehne and Sachs (2022), or Hendren (2020).

⁵For instance, Werning (2007) imposes a specific additively separable utility function and considers behavioral responses to taxation only at the intensive margin. Similar conditions can be found in the literature on the validity of the first-order approach in principal-agent problems, e.g., Rogerson (1985), Jewitt (1988) or Mirrlees (1999). For our application of interest, the introduction of the EITC, postulating only intensive margin responses would be inappropriate.

⁶References include Blundell, Brewer, Haan and Shephard (2009), Bargain, Dolls, Neumann, Peichl and Siegloch (2011), Bourguignon and Spadaro (2012), Lockwood and Weinzierl (2016), Jacobs, Jongen and Zoutman (2017), or Hendren (2020). For the older literature on the inverse approach to indirect taxation, see Christiansen and Jansen (1978) or Ahmad and Stern (1984).

The formal analysis focuses on the first-order conditions of a Mirrleesian optimal income tax problem and treats the welfare weights that appear in these conditions as endogenous objects. If these *implicit welfare weights* are found to be negative, this is taken to indicate that the given tax system is incompatible with the maximization of a Paretian social welfare function, and hence not Pareto-efficient. Again, it is common in this literature to impose regularity conditions to ensure that first-order conditions actually identify a welfare maximum.⁷

We contribute to the literature on the *inverse tax problem* by adding a formal proof that non-negative implicit weights are both necessary and sufficient for the non-existence of a Pareto-improving reform direction. Specifically, we show that the implicit welfare weights are negative in parts of the income distribution if and only if the revenue function \mathcal{R} violates the above conditions for Pareto efficiency. We also show how the information on where negative implicit weights arise can be used for the design of Pareto-improving reforms.

Related literature on the EITC. The EITC has motivated an optimal tax literature on the desirability of earnings subsidies for the “working poor”, giving rise to negative marginal and participation tax rates. Such negative tax rates can only be consistent with utilitarian welfare maximization when there are labor supply responses both at the intensive and at the extensive margin, see Saez (2002), Jacquet, Lehmann and Van der Linden (2013), and Hansen (2021).⁸

We complement this literature by taking a tax reform perspective: We ask whether, starting from the prevailing US tax system in the mid-1970s, the *introduction* of an EITC can be rationalized as an efficiency-enhancing reform.⁹ Note, moreover, that the introduction of the EITC in 1975 did not bring effective tax rates below zero.¹⁰ Our empirical analysis shows that the US tax system

⁷For a detailed discussion, see Bourguignon and Spadaro (2012).

⁸Choné and Laroque (2011) study optimal taxation with labor supply responses at the extensive margin only. For a non-utilitarian assessment, see Saez and Stantcheva (2016).

⁹For earlier contributions to the analysis of tax reforms, see Feldstein (1976), Weymark (1981) and the review in Guesnerie (1995). For more recent contributions, see Piketty (1997), Saez (2001), Golosov et al. (2014), or Jacquet and Lehmann (2021).

¹⁰Negative effective tax rates were only implemented after subsequent expansions of the federal

was not well-designed, and that it was possible to introduce earnings subsidies in a Pareto-improving way. This finding holds even without extensive-margin responses. Bastian and Jones (2021) estimate the extent to which EITC expansions since the 1990s were self-financing, taking account of a wide range of potential fiscal externalities associated with the expansion of the EITC. While they focus on a different period and use a different approach, their work is related to ours: Self-financing reforms are Pareto-improving.

Finally, our analysis draws on an empirical literature estimating the behavioral responses to the EITC. Specifically, Bastian (2020) estimates the behavioral responses to the 1975 EITC introduction, the reform that we focus on.¹¹

Outline. Section 2 provides necessary and sufficient conditions for the existence of Pareto-improving tax reforms. It also develops tools for an evaluation of tax reforms and for the characterization of optimal reform directions. Section 3 presents our empirical analysis of the reform introducing the EITC in 1975. Formal proofs are relegated to the Online Appendix.

2 Pareto-improving tax reforms

We first introduce a formal framework for the analysis of tax reforms. Section 2.2 then contains two theorems that, respectively, give necessary and sufficient conditions for the existence of Pareto-improving reform directions. The implications of this characterization for the inverse tax problem are discussed in Section 2.3. Tools to determine whether a given tax reform has a Pareto-improving direction are introduced in Section 2.4. This section also discusses the Pareto improvement that is optimal in the sense of yielding the maximal “free lunch”.

EITC and the introduction of state EITCs in the 1990s and 2000s.

¹¹Further references include Eissa and Liebman (1996), Meyer and Rosenbaum (2001), Moffitt (2003), Eissa and Hoynes (2004), Blundell (2006), or Kleven (2021). For surveys, see Hotz and Scholz (2003), Nichols and Rothstein (2015), and Hoynes (2019). For single mothers, early papers such as Meyer and Rosenbaum (2001) estimated large participation elasticities (sometimes above 1), while Kleven (2021) finds much smaller ones.

2.1 The model

We consider an economy with a continuum of individuals. Individuals value consumption c and generate earnings y . The generation of earnings comes with effort costs that depend on a vector of individual characteristics $\theta \in \Theta \subset \mathbb{R}^n$. The cross-section distribution of θ is assumed to be atomless and represented by a cumulative distribution function F . Preferences are represented by the utility function $u : \mathbb{R}_+^2 \times \Theta \rightarrow \mathbb{R}$. Thus, $u(c, y, \theta)$ is the utility that a type θ individual derives from a bundle (c, y) . The function u is continuously differentiable and increasing in the first argument, with partial derivative denoted by u_c . It is decreasing in the second argument. We do not impose an assumption of continuity or differentiability in y so as to allow for fixed costs in moving from zero earnings to strictly positive earnings.

We assume that a single-crossing condition holds in one dimension of the type space, Θ_j : If type (θ_j, θ_{-j}) weakly prefers a bundle (c, y) to another bundle $(c', y') < (c, y)$, then type (θ'_j, θ_{-j}) with $\theta'_j > \theta_j$ strictly prefers (c, y) to (c', y') . This assumption implies that the individuals' earnings are increasing in θ_j .¹²

There is a status quo tax policy. It is represented by a parameter c_0 and a tax function T_0 , which jointly define the budget set $C_0(y) = c_0 + y - T_0(y)$ that individuals face. The parameter c_0 is the intercept of this consumption schedule. It is the transfer to individuals with no earnings. Without loss of generality, we let $T_0(0) = 0$.¹³ We assume that T_0 is continuous. Otherwise, it can be an arbitrary non-linear tax function, possibly with kinks. Before the reform, individuals solve

$$\max_{y \in \mathcal{Y}} u(C_0(y), y, \theta),$$

where $\mathcal{Y} = [0, \bar{y}]$ is a set of feasible earnings level.

¹²This assumption ensures that the revenue implications of tax reforms are well-defined, see footnote 15.

¹³In the literature, $T_0(y)$ is often referred to as the *participation tax*; see, e.g., Kleven (2014). This reflects that $T_0(y)$ is the additional tax payment of a person with earnings of y , relative to a person with no earnings. Alternatively, we could represent the status quo by a tax function \tilde{T}_0 so that $\tilde{T}_0(y) := -c_0 + T_0(y)$ with the implication that $\tilde{T}_0(0) = -c_0$. We find it more convenient to separate the transfer c_0 from the tax function.

A tax reform replaces T_0 by a new tax function T_1 so that $T_1 = T_0 + \tau h$. The scalar τ is a measure of the size of the tax reform and the function h gives the direction of the tax reform. Again, h is assumed to be a continuous function with $h(0) = 0$. For a given income y , the change in the tax burden due to the reform is therefore given by $T_1(y) - T_0(y) = \tau h(y)$. After the reform, individuals solve

$$\max_{y \in \mathcal{Y}} u(C_1(y), y, \theta) , \quad (1)$$

where $C_1(y) = c_1 + y - T_0(y) - \tau h(y)$, and c_1 is the intercept after the reform. We denote the reform-induced change in tax revenue by $R(\tau, h)$ and assume that it is absorbed by the intercept so that

$$c_1 = c_0 + R(\tau, h) .$$

Thus, any change in tax revenue is redistributed in a lump-sum fashion.

The change in tax revenue $R(\tau, h)$ is an endogenous object that depends on the behavioral responses to taxation. To see how it is determined, let $y^*(e, \tau, h, \theta)$ be the solution to (1), where

$$C_1(y) = c_0 + e + y - T_0(y) - \tau h(y) ,$$

and e is a source of income that is exogenous from the individual's perspective. Also, let $y_0(\theta) := y^*(0, 0, h, \theta)$ be a shorthand for income in the status quo.¹⁴ Then, $R(\tau, h)$ solves

$$R(\tau, h) = \mathbf{E} [T_1(y^*(R(\tau, h), \tau, h, \theta)) - T_0(y_0(\theta))] , \quad (2)$$

where the operator \mathbf{E} indicates that we compute a population average using the distribution F .¹⁵

¹⁴There may be types for whom the utility-maximization problem in (1) has multiple solutions. The function y^* is then taken to select one of them. How this selection is done is inconsequential for the analysis that follows.

¹⁵ Brouwer's fixed point theorem can be used to establish the existence of a solution to this fixed point equation, in combination with conditions that ensure that $\mathbf{E} [T_1(y^*(e, \tau, h, \theta))]$ is continuous in e . This continuity is not immediate when the function y^* may exhibit jumps due to extensive-margin responses or discontinuities in marginal tax rates. With a single-crossing condition on preferences and an atomless type distribution, such jumps can be shown to wash out in the aggregate and therefore do not upset the continuity of $\mathbf{E} [T_1(y^*(e, \tau, h, \theta))]$ in e .

We denote by $v(\tau, h, \theta)$ the indirect utility that a type θ individual realizes after a tax reform (τ, h) . We can use the analysis of “*Envelope theorems for arbitrary choice sets*” in Milgrom and Segal (2002) to describe how individuals are affected by marginal changes of the reform intensity τ . Specifically, fix some type θ . Then, by Corollary 4 in Milgrom and Segal (2002),

$$\frac{d}{d\tau} v(\tau, h, \theta) = u_c(\cdot, \theta) [R_\tau(\tau, h) - h(y^*(\cdot))] , \quad (3)$$

where the marginal consumption utility of type θ , $u_c(\cdot, \theta)$, is evaluated at point $(C_1(y^*(\cdot)), y^*(\cdot))$, and $R_\tau(\tau, h)$ is the marginal effect of an increase in the reform intensity τ on tax revenue.¹⁶ More formally, it is the Gateaux differential of tax revenue in direction h .¹⁷ The envelope theorem covers cases in which the marginal tax rates (either in the status quo or after the reform) exhibit discontinuous jumps. It also applies when there are fixed costs of labor market participation, so that the utility function is, at $y = 0$, not continuous in y .

Equation (3) makes it possible to decompose the set of taxpayers into winners and losers of a tax reform. For concreteness, fix a reform direction h and suppose that, starting from the status quo policy, a small reform step has a positive impact on tax revenue, $R_\tau(0, h) > 0$. A taxpayer with type θ benefits from the reform if and only if this revenue gain outweighs the additional tax payment $h(y_0(\theta))$.

¹⁶For a type θ such that the utility-maximization problem in (1) has multiple solutions, the right-hand derivative of v is relevant for increases of τ and the left-hand derivative is relevant for decreases of τ .

¹⁷Our notation for Gateaux differentials is inspired by the one for partial derivatives. Conventions in mathematics are different. To make this explicit, let tax revenue \bar{R} be a real-valued functional of the tax function T . Then, the Gateaux differential of tax revenue in direction h is formally defined as

$$\partial \bar{R}(T, h) := \lim_{\tau \rightarrow 0} \frac{\bar{R}(T + \tau h) - \bar{R}(T)}{\tau} ,$$

where the left-hand side is the “typical” notation in the literature. Our notation can now be more formally introduced as $R_\tau(0, h) := \frac{d}{d\tau} \bar{R}(T_0 + \tau h)|_{\tau=0} = \partial \bar{R}(T_0, h)$. In Online Appendix A, we lay down further assumptions which guarantee the linearity of the Gateaux differential in the direction h , a property that is used in the proofs of Theorems 1 and 2.

Thus, a reform in direction h is Pareto-improving if

$$R_\tau(0, h) - \max_{y \in y_0(\Theta)} h(y) > 0, \quad (4)$$

where $y_0(\Theta)$ is the image of the function y_0 . If (4) holds, then every taxpayer's indirect utility goes up if a small step in direction h is undertaken.

There is no Pareto-improving direction in a class of reforms H if, for all functions $h \in H$,

$$R_\tau(0, h) - \max_{y \in y^*(\Theta)} h(y) < 0. \quad (5)$$

If (5) holds, then, for every direction $h \in H$, some taxpayer's indirect utility goes down if a small step in direction h is undertaken.¹⁸

The set H will be expanded as we go along. We first analyze classes of reforms with one or two income brackets in which marginal tax rates are changed. We then extend the results to tax reforms with finitely many brackets and, finally, cover the entire set of continuous reform directions h .

Single-bracket reforms. A single-bracket reform is a pair (τ, h^s) , where the function h^s is such that

$$h^s(y) = \begin{cases} 0, & \text{if } y \leq \hat{y}, \\ y - \hat{y}, & \text{if } y \in (\hat{y}, \hat{y} + \ell), \\ \ell, & \text{if } y \geq \hat{y} + \ell. \end{cases}$$

for some threshold value of income \hat{y} , see Figure 1. Thus, a single-bracket reform is characterized by a triplet (τ, ℓ, \hat{y}) , where \hat{y} is the income level at which the bracket starts, ℓ is the length of the bracket and τ is the amount by which the marginal tax rate changes for incomes in the bracket.

¹⁸Reforms so that $R_\tau(0, h) - \max_{y \in y_0(\Theta)} h(y) = 0$ require a more nuanced discussion, depending on whether, given h , the indirect utility of taxpayers with $y_0(\theta) \in \operatorname{argmax} h(y)$ attains a local minimum or a local maximum at $\tau = 0$. In the following, we focus on conditions (4) and (5) to avoid such case distinctions.

After a one-bracket reform, the new tax schedule is given by

$$T_1(y) = T_0(y) + \tau h^s(y) = \begin{cases} T_0(y), & \text{if } y \leq \hat{y}, \\ T_0(y) + \tau(y - \hat{y}), & \text{if } y \in (\hat{y}, \hat{y} + \ell), \\ T_0(y) + \tau\ell, & \text{if } y \geq \hat{y} + \ell. \end{cases}$$

Hence, the reform increases tax liabilities for all earnings above \hat{y} , with a maximum increase of $\tau\ell$. The marginal tax rate changes by τ for earnings in the bracket $(\hat{y}, \hat{y} + \ell)$. It does not change for incomes above or below this bracket. Formally, the new schedule of marginal tax rates equals

$$T_1'(y) = T_0'(y) + \tau h^{s'}(y) = \begin{cases} T_0'(y), & \text{if } y \leq \hat{y}, \\ T_0'(y) + \tau, & \text{if } y \in (\hat{y}, \hat{y} + \ell), \\ T_0'(y), & \text{if } y \geq \hat{y} + \ell. \end{cases}$$

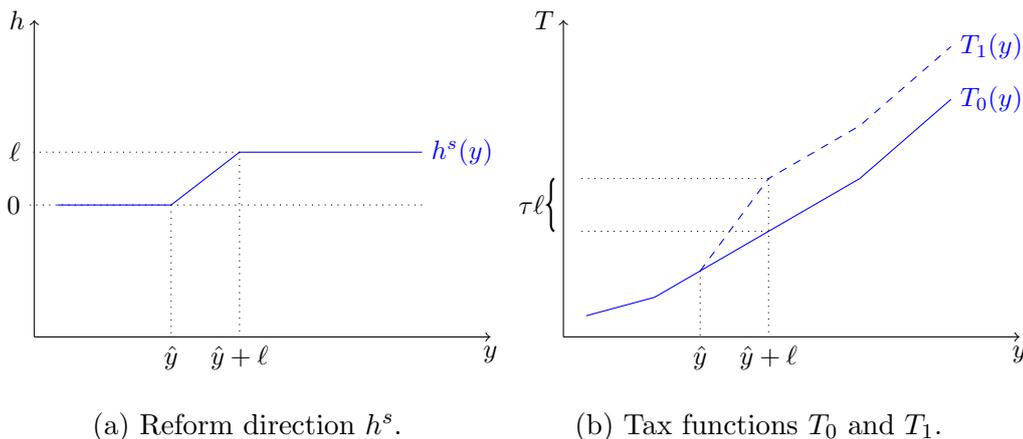


Figure 1: Income tax reforms with one bracket.

We will trace the welfare implications of multi-bracket reforms back to the properties of single-bracket reforms. It will prove convenient to have separate notation for the revenue implications of single-bracket reforms. For such reforms, we write $R^s(\tau, \ell, \hat{y})$ rather than $R(\tau, h^s)$. We write R_τ^s for the derivative of this function with respect to the first argument, and $R_{\tau\ell}^s$ for the cross-derivative with respect to the first and the second argument. It follows from first-order Taylor approximations that, for τ and ℓ close to zero,

$$\tau \ell R_{\tau\ell}^s(0, 0, y)$$

is a good approximation of $R^s(\tau, \ell, y)$, i.e., of the reform's revenue implications.¹⁹

Thus, the cross-derivative $R_{\tau\ell}^s$ can be interpreted as a measure of how much revenue can be raised by a small single-bracket reform. The function $y \mapsto R_{\tau\ell}^s(0, 0, y)$ is a recurrent theme in what follows. For a more concise notation, we will henceforth write $\mathcal{R}(y)$ rather than $R_{\tau\ell}^s(0, 0, y)$ and frequently refer to it as the *revenue function* $y \mapsto \mathcal{R}(y)$.

Two-bracket reforms. A two-bracket reform combines two single-bracket reforms. Formally, it is a pair (τ, h_2) , where the subscript of h_2 signifies a reform involving two brackets. The function h_2 is defined by

$$h_2(y) := \tau_1 h_1^s(y) + \tau_2 h_2^s(y) , \quad (6)$$

for

$$h_1^s(y) = \begin{cases} 0, & \text{if } y \leq y_1 , \\ y - y_1, & \text{if } y \in (y_1, y_1 + \ell \ell_1) , \\ \ell \ell_1, & \text{if } y \geq y_1 + \ell \ell_1 . \end{cases}$$

and

$$h_2^s(y) = \begin{cases} 0, & \text{if } y \leq y_2 , \\ y - y_2, & \text{if } y \in (y_2, y_2 + \ell \ell_2) , \\ \ell \ell_2, & \text{if } y \geq y_2 + \ell \ell_2 . \end{cases}$$

Thus, a two-bracket reform links two single-bracket reforms in a particular way: marginal tax rates change by $\tau \tau_1$ for incomes in the first bracket and by $\tau \tau_2$ for incomes in the second bracket. The first bracket has a length of $\ell \ell_1$, and the second bracket has a length of $\ell \ell_2$. The new tax schedule satisfies

$$T_1(y) = T_0(y) + \tau h_2(y) ,$$

and the new schedule of marginal tax rates equals

$$T_1'(y) = T_0'(y) + \tau h_2'(y) ,$$

¹⁹Note that $R^s(\tau, \ell, y) \simeq R^s(0, \ell, y) + \tau R_\tau^s(0, \ell, y) = \tau R_\tau^s(0, \ell, y)$ since $R^s(0, \ell, y) = 0$. Moreover, $\tau R_\tau^s(0, \ell, y) \simeq \tau \left(R_\tau^s(0, 0, y) + \ell R_{\tau\ell}^s(0, 0, y) \right) = \tau \ell R_{\tau\ell}^s(0, 0, y)$ since $R_\tau^s(0, 0, y) = 0$.

where

$$h'_2(y) = \begin{cases} \tau_1, & \text{for } y \in (y_1, y_1 + \ell \ell_1), \\ \tau_2, & \text{for } y \in (y_2, y_2 + \ell \ell_2), \\ 0, & \text{for } y \leq y_1, y \in [y_1 + \ell \ell_1, y_2], y \geq y_2 + \ell \ell_2. \end{cases}$$

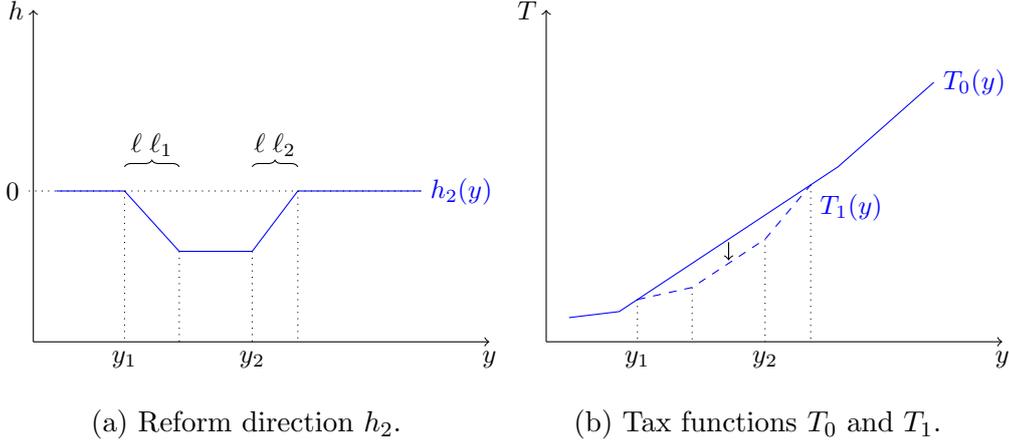


Figure 2: A two-bracket tax cut.

In what follows, two-bracket reforms with $\tau_1 < 0$, $\tau_2 > 0$ and $\tau_1 \ell_1 + \tau_2 \ell_2 = 0$ are of particular interest. We refer to them as *two-bracket tax cuts*. This terminology reflects that these reforms do not increase anyone's tax burden and that all people with an income between the endpoints of the two brackets get a tax cut. Moreover, they involve a phase-in range where marginal taxes are reduced, and a subsequent phase-out range where marginal taxes are increased, see Figure 2.

Our construction of two-bracket reforms facilitates an analysis of the limit case $\tau \rightarrow 0$ and $\ell \rightarrow 0$, see Figure 3 for the case of a small two-bracket tax cut. As τ goes to zero, the ratio of the marginal tax rate changes is kept constant at τ_1/τ_2 . Analogously, both brackets shrink when ℓ is sent to zero, while the ratio of their lengths is kept constant at ℓ_1/ℓ_2 .

Reforms with finitely many brackets. We extend the construction of two-bracket reforms to reforms with a finite number of brackets in the natural way: A reform (τ, h_m) with m brackets is given by a collection of m single-bracket reforms

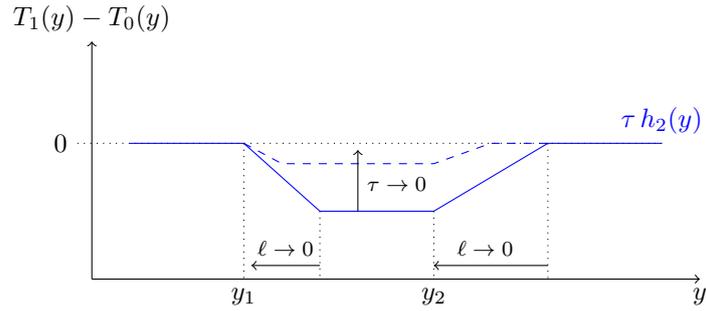


Figure 3: A small two-bracket tax cut.

such that $h_m(y) := \sum_{k=1}^m \tau_k h_k^s(y)$, where

$$h_k^s(y) = \begin{cases} 0, & \text{if } y \leq y_k, \\ y - y_k, & \text{if } y \in (y_k, y_k + \ell \ell_k), \\ \ell \ell_k, & \text{if } y \geq y_k + \ell \ell_k. \end{cases}$$

The reform-induced tax schedule is then given by $T_1 = T_0 + \tau h_m$.

2.2 On the existence of Pareto-improving reforms

Under what conditions is it possible to make everyone better off by increasing or lowering the marginal tax rates in a finite number of income brackets? Theorems 1 and 2 provide answers to this question.

Theorem 1 *If T_0 is a Pareto-efficient tax system, then the function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1, and non-increasing.*

Theorem 1 states necessary conditions for the Pareto efficiency of a tax system. The first condition is that $\mathcal{R}(y) \geq 0$ for all y . Hence, a one-bracket reform involving an increase of marginal tax rates must not lead to a loss of tax revenue. If the condition was violated, it would be possible to raise revenue by means of a tax cut, and such a reform would be Pareto-improving. The logic is familiar from analyses of the Laffer curve. The second condition is that $\mathcal{R}(y) \leq 1$ for all y . It is a mirror image of the first condition. If it was violated, it would be possible to raise so much revenue by increasing marginal tax rates that even those who suffer

most from the tax increase would be compensated. If T_0 is a Pareto-efficient tax system, there must be no scope for such a Pareto improvement.

The following proposition clarifies what reform options exist when the function $y \mapsto \mathcal{R}(y)$ is increasing over some range.

Proposition 1 *If there are two income levels y_1 and $y_2 > y_1$ such that $\mathcal{R}(y_2) > \mathcal{R}(y_1)$, then there exists a Pareto-improving two-bracket tax cut.*

In light of Proposition 1, Theorem 1 provides a characterization of two necessary conditions for Pareto efficiency: First, there must not be a Pareto improvement in the class of one-bracket reforms. Second, there must not be a Pareto improvement in the class of two-bracket tax cuts. As we show formally in the proof of Proposition 1, if $y \mapsto \mathcal{R}(y)$ is increasing, a two-bracket tax cut between incomes y_1 and y_2 is self-financing: The revenue loss due to a reduction of marginal tax rates in the first bracket is more than offset by the revenue gain from the increase of marginal tax rates in the second bracket. Thus, the condition that $y \mapsto \mathcal{R}(y)$ must be non-increasing is an analogue to the condition $\mathcal{R}(y) \geq 0$ for all y . The latter rules out self-financing tax cuts for one-bracket reforms. The former does so for two-bracket reforms.

Theorem 1 and Proposition 1 show that there may exist Pareto-improving two-bracket reforms, even when no Pareto-improving one-bracket reform can be found. Given this finding, one might conjecture that there is no hope to obtain a concise characterization of Pareto-efficient tax systems: Even if one found a condition ruling out Pareto-improving two-bracket reforms, there would still be the possibility of a Pareto-improving three-bracket reform. If one had eliminated those, one would still have to deal with four-bracket reforms, and so on. Theorem 2 shows that this is not the case: Ruling out Pareto-improving one- and two-bracket reforms is sufficient for Pareto efficiency.

Theorem 2 *If the function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1, and non-increasing, then there is no Pareto-improving direction in the class of reforms with finitely many brackets.*

According to Theorem 2, when there is no Pareto-improving direction in the sets of one-bracket or two-bracket reforms, then there is no Pareto-improving direction in the overall class of reforms where tax rates change in finitely many brackets. Put differently, a tax system that can be Pareto-improved by a tax reform that affects three or more brackets, can also be Pareto-improved by a tax reform that affects at most two brackets.

For an intuitive understanding of Theorem 2, consider the following problem: Try to find a tax reform that makes people with incomes in some arbitrary bracket ranging from an income level y_1 to some income level y_2 better off, without making any one else worse off. Since revenue is rebated lump sum, the reform must not yield a loss of overall tax revenue. Otherwise, people with zero incomes would be worse off.

Assume that there is no Pareto-improving reform with a single bracket, i.e., the revenue function is bounded from below by zero and from above by one. In this case, reforms that involve only tax rate cuts yield a revenue loss, so they cannot be Pareto-improving. Likewise, reforms that only involve tax rate increases neither fit the bill: They necessarily harm taxpayers further up in the income distribution. Thus, the only chance to make people with incomes between y_1 and y_2 better off, without making anyone else worse off, is to combine tax rate increases in some parts of the income distribution with tax rate reductions in other parts of the income distribution.

One candidate is a two-bracket tax cut that lowers marginal tax rates for incomes close to y_1 and increases them close to y_2 , so as to reduce the tax burden just for people with incomes between y_1 and y_2 . When y_2 is close to y_1 , a graph of the change in marginal tax rates looks like a sawtooth, see Figure 2a on page 15. If \mathcal{R} is decreasing, this reform comes with an overall revenue loss, violating again the constraint that no one must be made worse off. Reducing the tax burdens in further intervals – i.e., adding further sawteeth – will only aggravate the overall revenue loss. Therefore, if the function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1, and non-increasing, then there is nothing one can do

to make everyone better off. The formal proof of Theorem 2 in the Appendix is a generalization of this logic. It covers all combinations of tax rate increases and cuts that are logically conceivable with any given number m of brackets.

Continuous reform directions. So far, our results were restricted to the class of tax reforms with finitely many brackets, thereby excluding, e.g., continuously differentiable reform directions. The following corollary extends Theorem 2 to cover the entire class of continuous reform directions. It exploits that any continuous function $h : \mathcal{Y} \rightarrow \mathbb{R}$ can be approximated arbitrarily well by an m -bracket reform with m sufficiently large.

Corollary 1 *If $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1, and non-increasing, then there is no Pareto-improving direction in the class of continuous functions $h : \mathcal{Y} \rightarrow \mathbb{R}$.*

Thus, any tax system that can be Pareto-improved by a continuous reform, can also be Pareto-improved by a tax reform that affects at most two brackets.

An alternative characterization. As we discuss in part C of the Online Appendix, when earnings are bounded away from zero and bounded from above, it is possible to obtain a more parsimonious characterization of sufficient conditions for Pareto efficiency: The monotonicity condition on $y \mapsto \mathcal{R}(y)$ is then sufficient for the non-existence of Pareto-improving reform directions.²⁰ That said, for our application of interest, the introduction of the EITC, incentives for labor market participation play a key role. In this context, an assumption that everybody has strictly positive earnings would not be appropriate.

Tagging. Our analysis can be extended to allow for tagging.²¹ Suppose that the population can be divided into separate groups and that it is publicly observable to which group a person belongs. The tax-transfer system may then treat individuals who belong to different groups differently. For instance, transfers and earnings

²⁰We are grateful to an anonymous referee for pointing to this possibility.

²¹The seminal reference is Akerlof (1978). For a review, see Piketty and Saez (2013).

subsidies for lone mothers may be larger than those for childless individuals. The above analysis of Pareto-efficient tax reforms can then be applied separately for each group. This implies, in particular, that revenue changes due to a tax reform that affects one group are rebated lump sum in this group.²²

2.3 On the relation to the inverse tax problem

The literature on the “inverse tax problem” looks at observed tax policies using a revealed-preferences logic.²³ Specifically, it solves for “implicit welfare weights”, the weights for which an observed tax system satisfies the first-order conditions of an optimal tax problem. The following proposition clarifies the relation between an analysis of Pareto-improving tax reforms based on the revenue function and a welfare analysis that identifies implicit welfare weights. The proposition refers to the function $g : y \mapsto g(y)$ of implicit welfare weights. In Online Appendix A, we state the inverse tax problem explicitly and provide a formal definition of the function $y \mapsto g(y)$, see Equation (A.12).

Proposition 2 *The following statements are equivalent:*

- A. *The revenue function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1, and non-increasing for all $y \in \mathcal{Y}$.*
- B. *The function specifying the implicit welfare weights $y \mapsto g(y)$ is, almost everywhere, bounded from below by zero.*
- C. *There is no Pareto-improving direction in the set of continuous functions $h : \mathcal{Y} \mapsto \mathbb{R}$.*

Proposition 2 establishes that positive implicit welfare weights at all income levels are both necessary and sufficient for the non-existence of a Pareto-improving

²²This is without loss of generality. Redistributing the revenue gains from a self-financing tax reform among various groups can only make it more difficult to realize a Pareto improvement. Then, there are less resources that can be used to compensate those adversely affected by the tax reform.

²³See, e.g., Blundell et al. (2009), Bargain et al. (2011), Bourguignon and Spadaro (2012), Lockwood and Weinzierl (2016), Jacobs et al. (2017), or Hendren (2020).

reform direction ($B \Leftrightarrow C$). The previous literature has already argued that negative implicit weights indicate an inefficiency in the observed tax system, see Saez (2001) and Jacobs et al. (2017). Put differently, it has been anticipated that positive weights are necessary for Pareto efficiency ($C \Rightarrow B$). By contrast, it has not been anticipated that positive weights imply Pareto-efficiency ($B \Rightarrow C$). A direct proof of this statement faces the difficulty that, for a generic optimal tax problem with its incentive and resource constraints, it cannot be taken for granted that a tax system satisfying first-order conditions is actually a welfare maximum.²⁴

Our proof of Proposition 2 builds on Theorem 1 and Corollary 1, by which our conditions on the revenue function are both necessary and sufficient for the non-existence of Pareto-improving reform directions ($A \Leftrightarrow C$). We then proceed to show that implicit welfare weights are positive if and only if the revenue function is bounded from above and below, and non-increasing ($A \Leftrightarrow B$). This implies the equivalence of B and C .

Our proof that $A \Leftrightarrow B$ is instructive. It shows how negative weights map into Pareto-improving reforms and vice versa. Specifically, if $\mathcal{R}(y) < 0$ so that there is a Pareto-improving one-bracket tax cut for incomes close to y , then the implicit weights of people with incomes above y are, on average, negative. The converse implication also holds: If people richer than y have, on average, negative weights, then lowering marginal tax rates for incomes close to y is Pareto-improving. Analogously, we show that negative weights for incomes below y are equivalent to $\mathcal{R}(y) > 1$ so that there is a Pareto-improving reform that increases marginal tax rates close to y . Finally, there is an equivalence between \mathcal{R} being increasing on some income range, indicating that there is a Pareto improving two-bracket tax cut, and the average weight over this income range being negative.

The proof uses that the Gateaux differential of welfare in any reform direction h can be written as a functional of both the welfare weights $y \mapsto g(y)$ and the revenue function $y \mapsto \mathcal{R}(y)$, see Equation (A.12) in Online Appendix A. For a

²⁴See Appendices A and D in Bourguignon and Spadaro (2012) for a set of restrictions that ensure the consistency of the inverse tax problem. Note that Proposition 2 holds even if these restrictions are violated. It only requires that the revenue function exists.

welfare-maximum, it needs to be the case that this Gateaux differential is zero, for any direction h . Hence, this first-order condition implies a relationship between the revenue function and the implicit welfare weights that we exploit in the proof.

2.4 Evaluating tax reforms

Subsequently, we will use our approach to study the introduction of the EITC in the 1970s. We will find that there were indeed inefficiencies in the tax-transfer system prevailing at the time. Such a finding raises further questions: Did the 1975 EITC reform have a Pareto-improving direction? Relatedly, did the reform make the pre-existing inefficiency smaller? These questions require tools for an evaluation of tax reforms. We introduce them in this subsection. Proposition 3 below clarifies how one can check whether some given tax reform has a Pareto-improving direction. Further below, we introduce a money-metric measure of how inefficient a tax system is. This measure can be used to substantiate a statement such as “the inefficiency in the tax-transfer system for childless singles is small in comparison to the inefficiency prevailing in the one for single parents.”

Is a given reform direction Pareto-improving? Recall that a small reform in direction h is Pareto-improving if

$$R_\tau(0, h) > \max_{y \in y_0(\Theta)} h(y) ,$$

see inequality (4). It is not Pareto-improving if this condition is violated with a strict inequality, $R_\tau(0, h) < \max_{y \in y_0(\Theta)} h(y)$, see (5). The following proposition provides a characterization of $R_\tau(0, h)$ in terms of the revenue function \mathcal{R} .

Proposition 3 *The Gateaux differential of tax revenue in direction h is given by*

$$R_\tau(0, h) = \int_{\mathcal{Y}} h'(y) \mathcal{R}(y) dy . \tag{7}$$

By Equation (7), the Gateaux differential $R_\tau(0, h)$ is a weighted average of the revenue implications of one-bracket reforms, with the weights given by the function $y \mapsto h'(y)$. To interpret this function, note that a reform with direction h and

step size τ changes the marginal tax rate at income y by $\tau h'(y)$; hence, $h'(y)$ is the change in marginal tax rates per unit change of τ . Thus, with the function $y \mapsto \mathcal{R}(y)$ at hand, one can, for any continuous direction h , compute $R_\tau(0, h)$ using Proposition 3, and then use inequalities (4) and (5) to check whether or not it is Pareto-improving.

Measuring the size of inefficiencies. When a tax system (T_0, c_0) is not Pareto-efficient, there is a set of Pareto-improving reform directions. Is there a way to judge whether one of these reform directions is *better* than another one? One possibility is to order reform directions according to $R_\tau(0, h) - \max_{y \in y(\Theta)} h(y)$, the revenue gain in excess of what is needed to compensate the agents facing the largest tax increase.²⁵ We interpret this quantity as the *free lunch* that becomes available for everybody after a small reform in a Pareto-improving direction h . While some taxpayers may benefit more, no one gains less than $R_\tau(0, h) - \max_{y \in y(\Theta)} h(y)$.²⁶

With this measure, the *best* reform direction maximizes

$$\int_0^{\bar{y}} h'(y) \mathcal{R}(y) dy - \max_{y \in y_0(\Theta)} \int_0^y h'(z) dz$$

over the set of continuous functions $h : \mathcal{Y} \rightarrow \mathbb{R}$. For each income level y , this problem is linear in $h'(y)$, the reform-induced change in the marginal tax rate at income y . So, to ensure the existence of a solution, we impose the constraint that the changes in marginal tax rates are bounded in absolute value, $|h'(y)| \leq a$ for each $y \in \mathcal{Y}$. Henceforth, h^* denotes the solution to this problem and we refer to it as the *optimal reform direction*.

Proposition B.1 in Online Appendix B characterizes the optimal reform of a tax system that is inefficient because the revenue function $y \mapsto \mathcal{R}(y)$ is increasing

²⁵Recall that, by construction, the revenue gain $R_\tau(0, h)$ is distributed to the agents in a lump-sum fashion.

²⁶An alternative would be to evaluate Pareto-improvements according to their gross revenue implication $R_\tau(0, h)$. Since $R_\tau(0, h)$ is the additional transfer to people with no income, this is equivalent to an evaluation according to a Rawlsian social welfare function. A related approach is taken in Werning (2007) who proposes to measure the inefficiency of a tax system by the additional tax revenue that can be realized without making anyone worse off. Werning, however, does not specify how this tax revenue would be used.

between two income levels y_a and y_b . This is the relevant scenario in the context of our application. As we show, the reform direction h^* is then a two-bracket tax cut that affects marginal tax rates over a uniquely defined income range (y_s, y_t) .²⁷ In particular, h^* involves a reduction of marginal tax rates in a phase-in range going from the starting point y_s to the midpoint $(y_s + y_t)/2$ of the interval. Tax rates are increased in a phase-out range of the same length, going from the midpoint to the endpoint y_t . The impact of the optimal reform on the size of the free lunch can be written as

$$a I(T_0, c_0) = a \left[\int_{\frac{1}{2}(y_s+y_t)}^{y_t} \mathcal{R}(y) dy - \int_{y_s}^{\frac{1}{2}(y_s+y_t)} \mathcal{R}(y) dy \right]. \quad (8)$$

It equals the difference between the gain of tax revenue in the phase-out range and the loss of tax revenue in the phase-in range. Note that the term denoted by $I(T_0, c_0)$ does not depend on the parameter a that bounds the change in marginal tax rates. Hence, $I(T_0, c_0)$ is a scale-invariant measure of how inefficient the tax system is. Thus, we can say that tax system (T_A, c_A) is more inefficient than tax system (T_B, c_B) if $I(T_A, c_A) > I(T_B, c_B)$.

Formally, $I(T_0, c_0)$ is the Gateaux differential of the free lunch measure $R_\tau(0, h) - \max_{y \in y(\Theta)} h(y)$ in direction h^* . In the subsequent section, we will use this observation to give a sense of how big the free lunch could have become for small discrete tax reforms. Specifically, we will make use of the first-order Taylor approximation²⁸

$$R(\tau, h_R) - \max_{y \in y_0(\Theta)} \tau h_R(y) \simeq \tau a I(T_0, c_0), \quad (9)$$

and choose the parameters τ and a such that $\tau a = 0.01$. Thereby, we can approximate the revenue potential of tax reforms that change the marginal tax rates by at most one percentage point.

²⁷Specifically, the starting point y_s and the endpoint y_t are pinned down by the condition that $\mathcal{R}(y_s) = \mathcal{R}\left(\frac{y_s+y_t}{2}\right) = \mathcal{R}(y_t)$.

²⁸Higher-order Taylor approximations would involve higher-order Gateaux differentials.

2.5 Sufficient statistics for function \mathcal{R} : Examples

All previous results are expressed using the function $y \mapsto \mathcal{R}(y)$. Thus, given an empirical estimate of this function, our approach allows to check whether an observed tax system of interest is Pareto-efficient. Different models of taxation give rise to different versions of the function $y \mapsto \mathcal{R}(y)$. The concrete specification will depend on the application of interest and on a choice of what model to use for this application. We illustrate this with two examples.

First, consider the model of Diamond (1998) with $u(c, y, \theta) = c - \frac{1}{1+\frac{1}{\epsilon}} \left(\frac{y}{\theta}\right)^{1+\frac{1}{\epsilon}}$, where $\theta \in \Theta \subset \mathbb{R}_+$ is a measure of productivity and the parameter ϵ pins down the labor supply elasticity at the intensive margin. For this model, the revenue function is given by

$$\mathcal{R}(y) = 1 - F_y(y) - \varepsilon_0(y) y f_y(y) \frac{T'_0(y)}{1 - T'_0(y)} , \quad (10)$$

where F_y is the *cdf* of the earnings distribution, f_y is the corresponding *pdf*, and $\varepsilon_0 : y \mapsto \varepsilon_0(y)$ is a function that gives, for each level of y , the intensive-margin elasticity of earnings with respect to the retention rate $1 - T'_0(y)$.

Second, the literature on the desirability of earnings subsidies for the “working poor” suggests the use of a framework with taxpayers who differ both in the variable costs of productive effort and in the fixed costs of labor market participation.²⁹ For this framework, a version of the revenue function has first been derived by Jacquet et al. (2013); it also appears in Lorenz and Sachs (2016). For ease of exposition, we focus here on the case of quasi-linear preferences and iso-elastic effort costs. This is also the specification that we will use in our benchmark analysis of the EITC in the subsequent section. Hence, suppose that

$$u(c, y, \omega, \gamma) = c - \frac{1}{1+\frac{1}{\epsilon}} \left(\frac{y}{\omega}\right)^{1+\frac{1}{\epsilon}} - \gamma \mathbb{1}_{y>0} ,$$

where ω and γ are, respectively, interpreted as a taxpayer’s variable and fixed cost types. Thus, an individual’s type θ is now taken to be a pair $\theta = (\omega, \gamma)$ and

²⁹A similar framework is also used in the literature on optimal pension and retirement policies, see, e.g., Golosov, Shourideh, Troshkin and Tsyvinski (2013), Michau (2014), and Shourideh and Troshkin (2017).

$\Theta = \Omega \times \Gamma$. Then, the revenue function is given by

$$\mathcal{R}(y) = 1 - F_y(y) - \varepsilon_0(y) y f_y(y) \frac{T_0'(y)}{1 - T_0'(y)} - \int_y^\infty f_y(y') \pi_0(y') \frac{T_0(y')}{y' - T_0(y')} dy' , \quad (11)$$

where $\pi_0(y)$ is an extensive-margin (participation) elasticity. It measures the percentage of individuals with an income of y who leave the labor market when their after-tax income $y - T_0(y)$ is decreased by one percent. In the Supplementary Material, we also derive function $y \mapsto \mathcal{R}(y)$ for a more general framework (see Proposition E.1 in part E).³⁰

3 Empirical application: The introduction of the EITC

We now apply our insights on Pareto-improving tax reforms to study the 1975 introduction of the EITC and its subsequent expansion. After describing the 1975 EITC reform, we first use Theorem 1 in combination with the sufficient-statistics formula (11) to show that the US tax-transfer system was not Pareto-efficient prior to the introduction of the EITC. We then apply Proposition 3 to check whether the direction of the 1975 EITC reform was Pareto-improving, and compare it to the reform that would have maximized the “free lunch” for single parents at the time.

3.1 Background on the EITC

The introduction of the EITC in 1975 was a response to a “poverty trap”. In the 1960s, new welfare programs had been introduced as part of President Johnson’s “war on poverty”. On the one hand, the new programs provided more generous benefits to low-income households with children, especially to single mothers. On the other hand, these benefits were phased out in a way that implied high effective marginal tax rates for many single parents, often exceeding 70% (see Figure 4

³⁰This derivation is of stand-alone-interest in that it is based on a general specification of preferences, allowing for income effects at both margins, monetary or psychic fixed costs of labor market participation, and complementarities between consumption and leisure.

below). In the following decade, the share of welfare recipients increased substantially. By the early 1970s, finding ways out of the “poverty trap” by an increase of work incentives was considered a pressing concern.³¹

The US Congress enacted the EITC for the year 1975.³² As described in Bastian (2020), this reform was a substantial policy change that affected a large share of the population.³³ Initially, the program was restricted to working taxpayers with dependent children. It was set up as a refundable tax credit that was phased in at a marginal rate of 10% for taxpayers with less than 4,000 USD earned income per year, giving a maximum credit of 400 USD. The credit was then phased out at a marginal rate of 10% for earned incomes between 4,000 and 8,000 USD. Taxpayers with incomes above 8,000 USD were not eligible. Over the following decades, there were several expansions.³⁴

3.2 Calibration

We focus on two subgroups of the population, single parents and childless singles. In 1975, the EITC was introduced for the former, but not for the latter. Our analysis below will rationalize this policy choice: We will show that there was clearly scope for a Pareto-improving reform of the tax-transfer system for single parents, whereas no equally strong case can be made for childless singles. Our benchmark analysis is based on the formula for $y \mapsto \mathcal{R}(y)$ in Equation (11).

³¹Detailed reviews of the debates at the time can be found in Ventry (2000), Moffitt (2003), or Nichols and Rothstein (2015).

³²While the program was initially introduced as a temporary policy under the name *Earned Income Credit*, it was soon made permanent and relabeled to its current name *Earned Income Tax Credit*.

³³According to CPS data, 51.3% of single parents in the US had earned incomes in the EITC range (i.e., strictly positive and below 8,000 USD), and another 30.9% had no earned incomes. Similar figures applied at the state level, e.g., 52.2% of single parents with earned income in the EITC range and 25.8% with zero incomes in California.

³⁴For example, the eligibility thresholds were strongly increased in several steps, the first taking place in 1979. More generous credits were introduced for parents with two or more children in 1991, and for parents with three or more children in 2009. In 1994, US authorities also enacted a more modest EITC for childless workers. See Hoynes (2019) for a review.

We use data on the most important elements of the US tax-transfer system for the tax year 1974 and later. Specifically, we take account of the federal income tax and the two largest welfare programs, Aid for Families with Dependent Children (AFDC) and Supplementary Nutrition Assistance Programs (SNAP, also known as Food Stamps).³⁵ Some parameters of AFDC varied across states, so that a unified treatment for the US at large is not possible. Our benchmark analysis therefore focuses on California, the state with the largest population both in the 1970s and today.³⁶ Moreover, taxes and welfare transfers differed with respect to the number of children. In the benchmark analysis, we focus on the subgroup of single parents with two children.³⁷

Figure 4 shows effective marginal tax rates, $y \mapsto T'_0(y)$, and participation tax rates, $y \mapsto \frac{T_0(y)}{y}$, for single parents (left panel) and for childless singles (right panel) before the reform in 1974.³⁸ At low incomes, both marginal tax rates and participation tax rates were much higher for single parents than for childless singles. The reason is that, for single parents, the phasing-out of AFDC and SNAP transfers implied an income range with exceptionally high marginal tax rates well above 70% and participation tax rates above 60%. This was not the case for childless singles. The dotted vertical lines in both panels of Figure 4 indicate the income range that was affected by the introduction of the EITC in 1975: It reduced marginal taxes by 10 percentage points in the phase-in range between 0 and 4,000 USD (first dotted line) and raised them by 10 percentage points in the phase-out range between 4,000 and 8,000 USD (second dotted line).³⁹

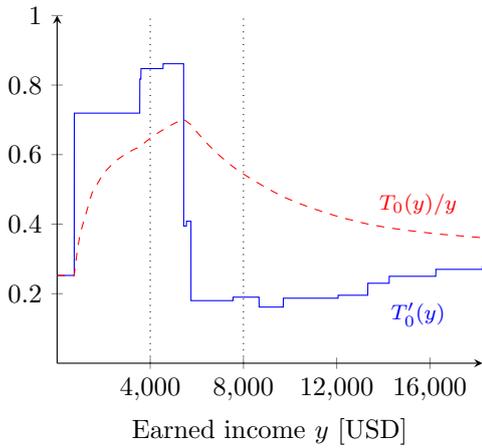
³⁵See Table S.2 in part D of the Supplementary Material for details and sources.

³⁶In the Supplementary Material, we also present results for the four largest US states next to California: New York, Texas, Pennsylvania, and Illinois.

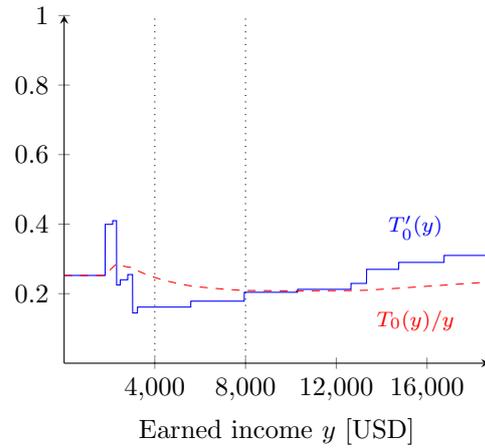
³⁷In our data, the median number of children in single-parent households was two, while the arithmetic mean was 2.2. As we show in the Supplementary Material, an analysis for single parents with less or more children yields similar conclusions.

³⁸Recall that we define $T_0(y)$ to capture the participation tax at income y , i.e., the tax payment at income y relative to the tax payment at zero income. In the literature, the ratio $T_0(y)/y$ is commonly referred to as the participation tax rate, see, e.g., Kleven (2014).

³⁹See Figure S.19 in the Supplementary Material for more detailed information on how effective marginal changed from 1974 to 1975.



(a) Single parents.



(b) Childless singles.

Figure 4: 1974 US tax-transfer schedules, single parents and childless singles.

Notes: Figure 4 shows the 1974 effective marginal tax $T'_0(y)$ (blue lines) and participation tax rate $T_0(y)/y$ (red lines) for single parents (left panel) and for childless singles (right panel) as functions of earned income in 1974 USD. The introduction of the EITC in 1975 decreased marginal taxes between 0 and 4,000 USD (first dotted vertical line) and increased them between 4,000 and 8,000 USD (second dotted vertical line).

Source: Authors' calculations (see part D of the Supplementary Material for details).

We estimate the 1974 income distributions in both subgroups based on data from the March 1975 Current Population Survey (CPS). For our benchmark scenario, we consider the sample of non-married individuals aged 25 to 60 who do neither co-habit with an unmarried spouse nor with another adult family member. We partition this sample into childless singles and single parents. In line with the EITC rules, we consider as earned income the sum of wage income and self-employment income. Single parents with strictly positive earned incomes below 8,000 USD were eligible for the EITC. For our benchmark analysis, we estimate the income distributions for both groups using a non-parametric kernel density estimation.

We draw on a rich literature providing empirical estimates of labor supply elasticities. Our benchmark analysis for childless singles is based on the elasticities suggested by Chetty, Guren, Manoli and Weber (2013): a participation elasticity that equals 0.25 on average, and an intensive-margin elasticity of 0.33. For single parents, we use an average participation elasticity of 0.58, as estimated by Bastian (2020) based on the 1975 EITC reform. For the intensive margin, we also use an elasticity of 0.33. For both groups, we assume, in line with the empirical evidence, that participation elasticities decline with income (see part D of the Supplementary Material for details).

3.3 Empirical results

In the following, we present our benchmark calibrations of the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ and $y \mapsto \mathcal{R}_{cs}(y)$ for single parents and childless singles, respectively. We then use these functions to investigate, first, whether the US tax-transfer system was Pareto-efficient prior to the EITC introduction and, second, whether the 1975 EITC introduction for single parents was a reform in a Pareto-improving direction. Third, we characterize the tax reform direction h^* that would have maximized the “free lunch”, i.e., the tax revenue in excess of what was needed to make sure that no one was made worse off.

Was the 1974 US tax-transfer system Pareto-efficient? Figure 5 plots the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ and $y \mapsto \mathcal{R}_{cs}(y)$ for our benchmark calibration of the 1974 US tax system. Specifically, the solid blue line depicts the revenue function $\mathcal{R}_{sp}(y)$ for single parents, while the teal line depicts the revenue function $\mathcal{R}_{cs}(y)$ for childless singles.

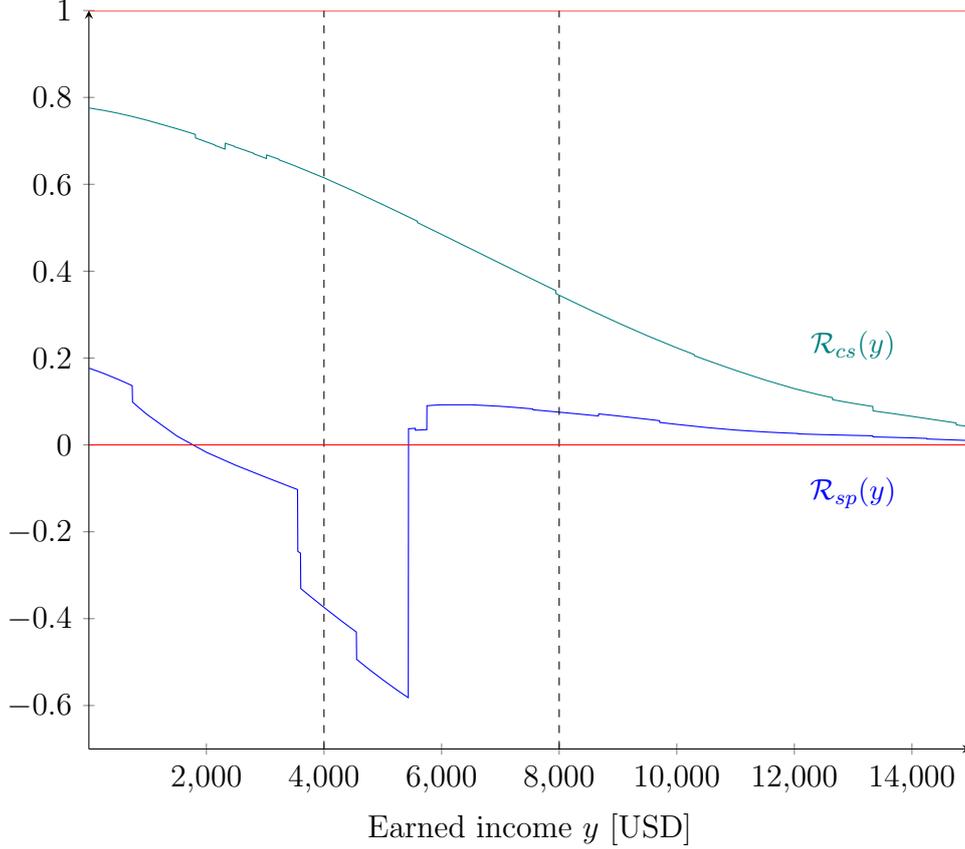


Figure 5: Revenue functions of the 1974 US tax-transfer system.

Notes: Figure 5 shows the revenue functions for single parents $y \mapsto \mathcal{R}_{sp}(y)$ (blue) and childless singles $y \mapsto \mathcal{R}_{cs}(y)$ (teal) in 1974 as functions of earned income for our benchmark calibration: intensive-margin elasticities of 0.33 for both groups, average participation elasticities of 0.58 for single parents and 0.25 for childless singles. The vertical dashed lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1974 EITC.

Source: Authors' calculations (see part D of the Supplementary Material for details).

For single parents, the function \mathcal{R}_{sp} does not satisfy the necessary conditions for Pareto efficiency in Theorem 1. First, it attains negative values for incomes

between approximately 1,500 and 5,400 USD. This implies that one-bracket tax cuts in this income range would have been self-financing and Pareto-improving. Second, \mathcal{R}_{sp} is increasing in the income range between 5,000 and 6,000 USD, thereby violating the monotonicity condition. Hence, there was room for Pareto-improving two-bracket tax cuts, resembling the EITC.

For childless singles, the revenue function \mathcal{R}_{cs} is throughout between 0 and 1, so that there was no scope for a Pareto-improving reform involving only a single bracket. By contrast, the monotonicity condition is violated as the teal line is slightly increasing in the range between 2,200 and 3,200 USD. Again, this indicates the possibility of Pareto-improving two-bracket reforms. That said, the visual impression is that the scope for such a Pareto improvement was more limited for childless singles than for single parents. Below, we confirm this conjecture using the inefficiency measure introduced in Subsection 2.4.

Sensitivity analysis. The finding that it was possible to realize a Pareto improvement by means of a two-bracket tax cut for single parents is robust in various dimensions. For instance, Figure 6 explores alternative assumptions about behavioral responses at the extensive margin: It shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for our benchmark calibration with an average participation elasticity of 0.58 (blue line), for a case with a larger participation elasticity of 0.9 (teal line), and for the limit case with a participation elasticity of zero (black line). The figure suggests that the scope for Pareto-improving reforms is larger, the more strongly labor supply responds at the extensive margin. But even with a participation elasticity of zero, such a reform would have been Pareto-improving. This observation is interesting in the light of the discussion about the EITC from an optimal tax perspective, where positive extensive-margin elasticities are often found to be necessary for the desirability of an EITC, see, e.g., Saez (2002) or Hansen (2021). As we show here, with a tax reform perspective applied to the tax-transfer system as of 1974, the introduction of the EITC can be rationalized even when there are no behavioral responses at the extensive margin.

In Part D.2 of the Supplementary Material, we provide further robustness

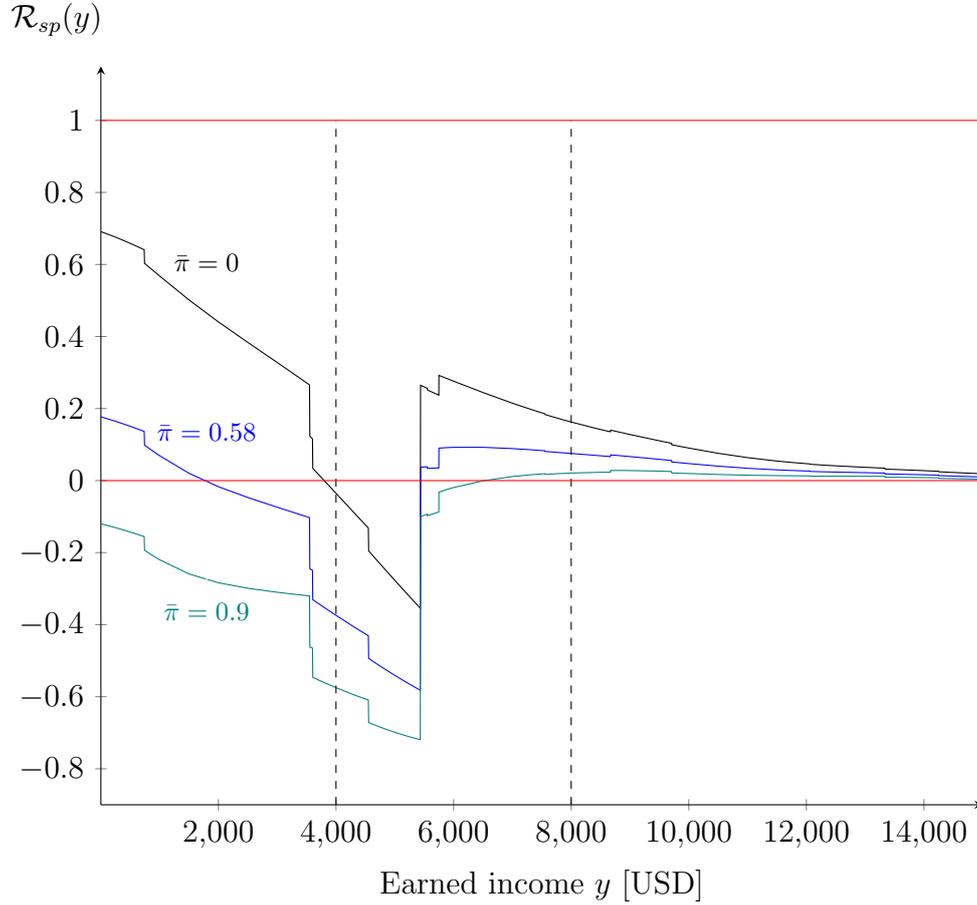


Figure 6: Revenue function of the 1974 US tax-transfer system for single parents, different participation elasticities.

Notes: Figure 6 shows the revenue function \mathcal{R}_{sp} for single parents in 1974, assuming an average participation elasticity of 0.58 (blue line, benchmark), a higher participation elasticity of 0.9 (teal line) and a case without extensive-margin responses (black line). The intensive-margin elasticity is held at the benchmark level of 0.33. The dashed vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see part D of the Supplementary Material for details).

checks with respect to all ingredients of the sufficient-statistics formula (11) for the revenue function: labor supply responses, income distributions, and tax-transfer rates. For example, we consider several alternative assumptions about intensive-margin elasticities and income effects. We also estimate the income distributions using different sample restrictions and other data sets. Finally, we consider a number of alternative representations of the tax-transfer system, e.g., including wealth tests or payroll taxes, focusing on single parents with one or three children, or using the transfer schedules from other US states. We find that, by and large, our main results remain valid across all these robustness checks.

Was the 1975 EITC reform Pareto-improving? Figure 5 shows that it was possible to Pareto-improve the US tax-transfer system by a two-bracket tax cut, i.e., by a reform with the same qualitative properties as the introduction of the EITC. A separate question is whether the EITC reform that *actually took place* in 1975 went into a Pareto-improving direction. To answer this question, we make use of the formula in Proposition 3, which combines an estimate of the revenue effects of single-bracket reforms with information on the direction of the 1975 EITC reform, referred to as \tilde{h}_{75} below. Specifically, the 1975 EITC reform reduced marginal taxes at all incomes below 4,000 USD by 10 percentage points, and increased marginal taxes at all incomes between 4,000 and 8,000 USD by the same magnitude. It did not increase tax liabilities at any income level. This reform had a Pareto-improving direction if the condition

$$R_{\tau}(0, \tilde{h}_{75}) = - \int_0^{4,000} \mathcal{R}_{sp}(y) dy + \int_{4,000}^{8,000} \mathcal{R}_{sp}(y) dy > 0 \quad (12)$$

is satisfied. Whether this inequality holds depends on the details of the calibration. For our benchmark scenario with an intensive-margin elasticity of 0.33 and an average participation elasticity of 0.58, the reform was not Pareto-improving. For a participation elasticity above 0.84, by contrast, it was Pareto-improving.⁴⁰

⁴⁰Some earlier studies indeed estimated participation elasticities in this range, see, e.g., Meyer and Rosenbaum (2001).

The optimal reform. What would have been the reform maximizing the “free lunch” in 1975 and how does it relate to the reform that was actually adopted? We find that, under our benchmark calibration, the former was a two-bracket tax cut that would have reduced marginal tax rates between 1,248 and 5,748 USD, and increased them between 5,748 and 10,248 USD. Thus, the optimal reform would have been a version of the EITC that involved a wider range of incomes, and also higher incomes than the actual 1975 EITC. Figure 7 illustrates this reform and its revenue implications graphically. With the revenue gain from this reform, it would have been possible to pay an additional lump-sum transfer of 12.6 USD (in 1975 values) per percentage-point change in marginal taxes to each single parent (corresponding to 71 USD in 2021).

For childless singles, the corresponding number is much smaller, namely about 1 cent per percentage-point change in marginal taxes. This confirms the conjecture above that the inefficiency in the tax-transfer system for childless singles was orders of magnitude smaller than the one for single parents.

Did subsequent reforms improve the EITC? Since its introduction, the EITC was repeatedly reformed and expanded in two major ways. First, the range of eligible incomes was enlarged in several steps, starting with the 1979 reform. Second, benefits were made dependent on family size, with larger benefits for families with more children introduced in 1991 and 2009. Did these reforms improve the design of the EITC or, put differently, was there progress in US tax policy for people with low incomes? Our approach to study Pareto-improving tax reforms can also be used to answer these questions. The following paragraphs summarize our findings; a more detailed analysis can be found in part D of the Supplementary Material.

First, we study the efficiency of the US tax-transfer system between the EITC introduction in 1975 and the first EITC expansion in 1979. As of 1974, the optimal reform would have reduced marginal taxes between 1,248 and 5,748 USD. The 1975 EITC only reduced marginal taxes below 4,000 USD, however, and increased them between 4,000 and 8,000 USD. Thus, the actual reform aggravated

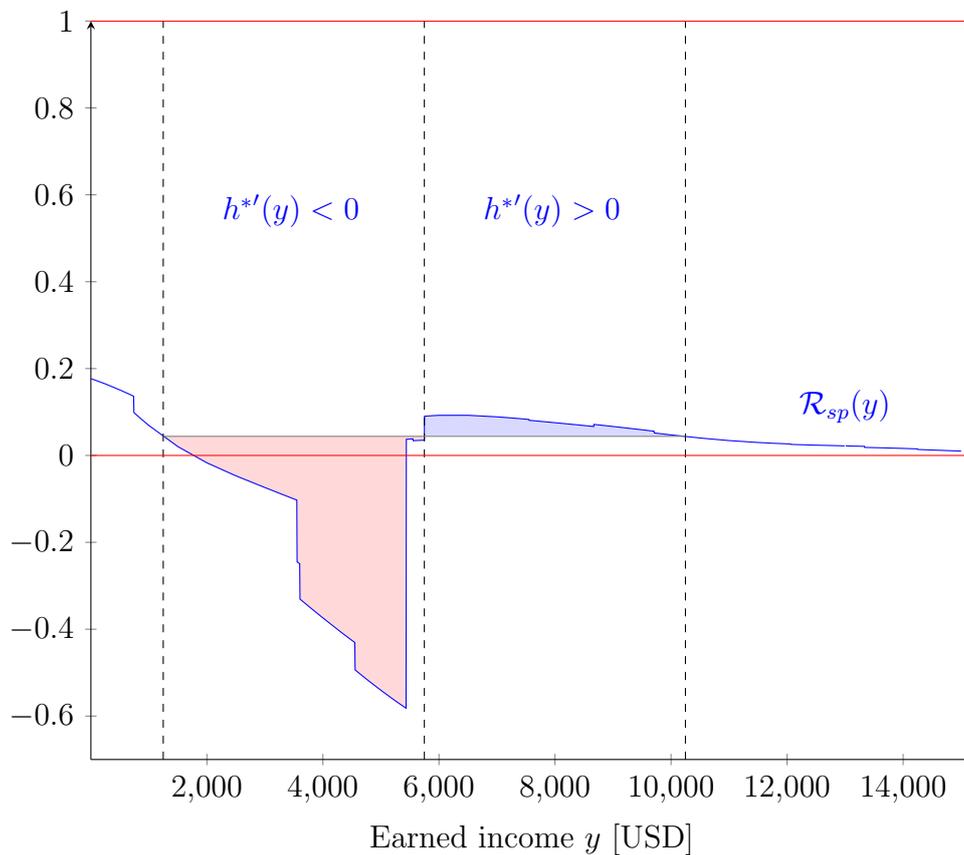


Figure 7: The optimal Pareto-improving reform for single parents.

Notes: Figure 7 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents given our benchmark calibration. The optimal tax reform would have reduced marginal taxes between 1,248 (first dashed line) and 5,748 USD (second dashed line) and increased them between 5,748 and 10,248 USD (third dashed line). The sum of both shaded areas represents the available revenue gain per single-parent family.

Source: Authors' calculations (see part D of the Supplementary Material for details).

inefficiencies at the bottom of the phase-out range. Correspondingly, the 1975 US tax-transfer system did not become Pareto-efficient, and there was further scope for Pareto-improving reforms at higher incomes. In 1979, the US government indeed extended the EITC to higher income levels – by means of a two-bracket tax cut affecting incomes between 4,000 and 10,000 USD. We find that the 1979 EITC reform actually had a Pareto-improving direction. This result is robust to alternative assumptions on labor supply elasticities. Thus, while the initial version of the EITC was suboptimal, its design was improved subsequently.

Second, we look into the desirability of making the EITC provisions dependent on the number of children. We find that, in 1975, the introduction of a hypothetical EITC schedule would have been Pareto-improving for each subgroup of single parents. The length and location of the optimal phase-in and phase-out ranges differed substantially, however, with wider income ranges for single parents with more children.⁴¹ Moreover, the size of the available free lunches was increasing in the number of children. Taken together, these results indicate, first, that the inefficiencies in the tax systems for larger families were more severe and, second, that the introduction of differentiated EITC schedules would have allowed to realized larger efficiency gains. The 1975 EITC did not condition on the number of children, by contrast. More than a decade later, the US government improved the design of the EITC in this dimension, introducing more generous tax credits for families with two and more children in 1991, and for families with three and more children in 2009.

4 Concluding Remarks

A key lesson from this paper is that tax reforms with two brackets – a lower one in which tax rates are lowered, and a higher one in which tax rates are increased – deserve particular attention. Our theoretical results show that such reforms can make every one better off, even if no simple one-bracket tax reform can. Moreover,

⁴¹We also find that the historical 1975 EITC reform would have been Pareto-improving if applied to the subgroup of single parents with one child only.

a tax system is Pareto-efficient if there is no Pareto improvement in the class of tax reforms that affect at most two brackets. We also show how Pareto-improving reforms can be identified in practice.

Our empirical analysis of the EITC introduction shows that such reforms have also been successfully used: Before the EITC was introduced, the US tax-transfer system for single parents was not Pareto-efficient. The historical 1975 EITC reform was not Pareto-improving, but a similar reform affecting marginal taxes in two larger income brackets would have been. The 1979 reform of the EITC then expanded the range of incomes covered, thereby realizing a Pareto improvement. Thus, while the initial version of the EITC was suboptimal, its design was improved subsequently.

References

- Ahmad, E. and N. Stern**, “The theory of reform and indian indirect taxes,” *Journal of Public Economics*, 1984, 25 (3), 259–298.
- Akerlof, G.A.**, “The economics of “tagging” as applied to the optimal income tax, welfare programs, and manpower planning,” *American Economic Review*, 1978, 68 (1), 8–19.
- Badel, A. and M. Huggett**, “The sufficient statistic approach: Predicting the top of the Laffer curve,” *Journal of Monetary Economics*, 2017, 87, 1–12.
- Bargain, O., M. Dolls, D. Neumann, A. Peichl, and S. Siegloch**, “Tax-Benefit Systems in Europe and the US: Between Equity and Efficiency,” *IZA Discussion Papers 5440*, 2011.
- , —, —, —, and —, “Comparing inequality aversion across countries when labor supply responses differ,” *International Tax and Public Finance*, 2014, 21 (5), 845–873.
- Bastani, S., S. Blomquist, and L. Micheletto**, “Pareto efficient income taxation without single-crossing,” *Social Choice and Welfare*, 2020.

- Bastian, J.E.**, “The Rise of Working Mothers and the 1975 Earned Income Tax Credit,” *American Economic Journal: Economic Policy*, 2020, 12 (3), 44–75.
- and **M.R. Jones**, “Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Public Assistance Spending,” *Journal of Public Economics*, 2021, 196, 104355.
- Bierbrauer, F.J. and P.C. Boyer**, “Politically feasible reforms of non-linear tax systems,” *CEPR Discussion Paper 13059*, 2018.
- , — , and **A. Peichl**, “Politically feasible reforms of nonlinear tax systems,” *American Economic Review*, 2021, 111 (1), 153–91.
- Blundell, R.**, “Earned income tax credit policies: Impact and optimality: The Adam Smith Lecture, 2005,” *Labour Economics*, 2006, 13 (4), 423–443.
- and **A. Shephard**, “Employment, hours of work and the optimal taxation of low-income families,” *Review of Economic Studies*, 2012, 79 (2), 481–510.
- , **M. Brewer, P. Haan, and A. Shephard**, “Optimal Income Taxation of Lone Mothers: An Empirical Comparison of the UK and Germany,” *Economic Journal*, 2009, 119, F101–F121.
- Boadway, R. and M. Keen**, “Redistribution,” in A.B. Atkinson and F. Bourguignon, eds., *Handbook of Income Distribution*, Vol. 1, Elsevier, 2000, pp. 677–789.
- Bourguignon, F. and A. Spadaro**, “Tax-Benefit Revealed Social Preferences,” *Journal of Economic Inequality*, 2012, 10 (1), 75–108.
- Brito, D.L., J.H. Hamilton, S.M Slutsky, and J.E. Stiglitz**, “Pareto Efficient Tax Structures,” *Oxford Economic Papers*, 1990, 42 (1), 61–77.
- Cesarini, D., E. Lindqvist, M.J. Notowidigdo, and R. Östling**, “The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries,” *American Economic Review*, 2017, 107 (12), 3917–46.

- Chetty, R., A. Guren, D. Manoli, and A. Weber**, “Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities,” *NBER Macroeconomics Annual*, 2013, *27* (1), 1–56.
- Choné, P. and G. Laroque**, “Optimal taxation in the extensive model,” *Journal of Economic Theory*, 2011, *146* (2), 425–453.
- Christiansen, V. and E. Jansen**, “Implicit Social Preferences in the Norwegian System of Indirect Taxation,” *Journal of Public Economics*, 1978, *10* (2), 217–245.
- Daly, M.C. and R.V. Burkhauser**, “The Supplemental Security Income Program,” in R.A. Moffitt, ed., *Means-Tested Transfer Programs in the United States*, University of Chicago Press, 2003, pp. 79–139.
- Diamond, P.A.**, “Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates,” *American Economic Review*, 1998, *88*, 83–95.
- Eissa, N. and H.W. Hoynes**, “Taxes and the labor market participation of married couples: the Earned Income Tax Credit,” *Journal of Public Economics*, 2004, *88*, 1931–1958.
- and **J.B. Liebman**, “Labor Supply Response to the Earned Income Tax Credit,” *Quarterly Journal of Economics*, 1996, *111* (2), 605–637.
- Feldstein, M.**, “On the theory of tax reform,” *Journal of Public Economics*, 1976, *6* (1-2), 77–104.
- Giannarelli, L.**, *An Analyst’s Guide to TRIM2: The Transfer Income Model, Version 2*, Washington: The Urban Institute Press, 1992.
- Golosov, M., A. Shourideh, M. Troshkin, and A. Tsyvinski**, “Optimal pension systems with simple instruments,” *American Economic Review: Papers and Proceedings*, 2013, *103*, 502–507.

- , **A. Tsyvinski**, and **N. Werquin**, “A Variational Approach to the Analysis of Tax Systems,” *NBER Working Paper 20780*, 2014.
- Guesnerie, R.**, *A Contribution to the Pure Theory of Taxation*, Cambridge University Press, 1995.
- Hansen, E.**, “Optimal Income Taxation with Labor Supply Responses at Two Margins: When Is An Earned Income Tax Credit Optimal?,” *Journal of Public Economics*, 2021, *195*, 104365.
- Hendren, N.**, “Measuring Economic Efficiency Using Inverse Optimum Weights,” *Journal of Public Economics*, 2020, *187*, 104198.
- Holtz-Eakin, D., D. Joulfaian, and H.S. Rosen**, “The Carnegie Conjecture: Some Empirical Evidence,” *Quarterly Journal of Economics*, 1993, *108* (2), 413–435.
- Hotz, V.J. and J.K. Scholz**, “The Earned Income Tax Credit,” in R.A. Moffitt, ed., *Means-Tested Transfer Programs in the United States*, University of Chicago Press, 2003, pp. 141–198.
- Hoynes, H.**, “The Earned Income Tax Credit,” *The Annals of the American Academy of Political and Social Science*, 2019, *686* (1), 180–203.
- Imbens, G.W., D.B. Rubin, and B.I. Sacerdote**, “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players,” *American Economic Review*, 2001, *91* (4), 778–794.
- Jacobs, B., E. Jongen, and F. Zoutman**, “Revealed social preferences of Dutch political parties,” *Journal of Public Economics*, 2017, *156*, 81–100.
- Jacquet, L. and E. Lehmann**, “Optimal Income Taxation with Composition Effects,” *Journal of the European Economic Association*, 2021, *19* (2), 1299–1341.

- , — , and **B. Van der Linden**, “Optimal redistributive taxation with both extensive and intensive responses,” *Journal of Economic Theory*, 2013, 148 (5), 1770–1805.
- Jewitt, I.**, “Justifying the First-Order Approach to Principal-Agent Problems,” *Econometrica*, 1988, 56 (5), 1177–1190.
- Juhn, C., K.M. Murphy, and R.H. Topel**, “Why Has the Natural Rate of Unemployment Increased over Time?,” *Brookings Papers on Economic Activity*, 1991, (2), 75–142.
- , — , and — , “Current Unemployment, Historically Contemplated,” *Brookings Papers on Economic Activity*, 2002, (1), 79–136.
- Kleven, H.**, “How Can Scandinavians Tax So Much?,” *Journal of Economic Perspectives*, 2014, 28 (4), 77–98.
- , “The EITC and the Extensive Margin: A Reappraisal,” NBER Working Paper No. 26405, 2021.
- Koehne, S. and D. Sachs**, “Pareto-Improving Reforms of Tax Deductions,” *European Economic Review*, 2022, p. 104214.
- Lockwood, B.B. and M. Weinzierl**, “Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions,” *Journal of Monetary Economics*, 2016, 14-119 (77), 30–47.
- Lorenz, N. and D. Sachs**, “Identifying Laffer Bounds: A Sufficient-Statistics Approach with an Application to Germany,” *Scandinavian Journal of Economics*, 2016, 118 (4), 646–665.
- Meghir, C. and D. Phillips**, “Labour Supply and Taxes,” in J. Mirrlees, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba, eds., *Dimensions of Tax Design: The Mirrlees Review*, Oxford University Press, 2010, pp. 202–274.

- Meyer, B.D. and D.T. Rosenbaum**, “Welfare, the Earned Income Tax Credit and the Labor Supply of Single Mothers,” *Quarterly Journal of Economics*, 10 2001, *116* (3), 1063–1114.
- Michau, J.-B.**, “Optimal redistribution: A life-cycle perspective,” *Journal of Public Economics*, 2014, *111*, 1–16.
- Milgrom, P. and I. Segal**, “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 2002, *70* (2), 583–601.
- Mirrlees, J. A.**, “The Theory of Moral Hazard and Unobservable Behaviour: Part I,” *The Review of Economic Studies*, 1999, *66* (1), 3–21.
- Moffitt, R.A.**, “Cumulative Effective Tax Rates and Guarantees in Low-Income Transfer Programs,” *Journal of Human Resources*, 1979, *14* (1), 122–129.
- , “The Negative Income Tax and the Evolution of U.S. Welfare Policy,” *Journal of Economic Perspectives*, 2003, *17* (3), 119–140.
- Nichols, A. and J. Rothstein**, “The Earned Income Tax Credit (EITC),” *NBER Working Paper Series*, 2015.
- Piketty, T.**, “La redistribution fiscale face au chômage,” *Revue française d’économie*, 1997, *12*, 157–201.
- and **E. Saez**, “Optimal Labor Income Taxation,” in A.J. Auerbach, R. Chetty, M. Feldstein, and E. Saez, eds., *Handbook of Public Economics*, Vol. 5, Elsevier, 2013, pp. 391–474.
- Rogerson, W.P.**, “The First-Order Approach to Principal-Agent Problems,” *Econometrica*, 1985, *53* (6), 1357–1367.
- Saez, E.**, “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, 2001, *68*, 205–229.
- , “Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses,” *Quarterly Journal of Economics*, 2002, *117* (3), 1039–1073.

- and **S. Stantcheva**, “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” *American Economic Review*, 2016, *106* (1), 24–45.
- , **J. Slemrod**, and **S.H. Giertz**, “The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review,” *Journal of Economic Literature*, 2012, *50* (1), 3–50.
- Scheuer, F.**, “Entrepreneurial taxation with endogenous entry,” *American Economic Journal: Economic Policy*, 2014, *6* (2), 126–63.
- Shourideh, A. and M. Troshkin**, “Incentives and Efficiency of Pension Systems,” Working Paper 2017.
- Stantcheva, S.**, “Optimal income taxation with adverse selection in the labour market,” *Review of Economic Studies*, 2014, *81* (3), 1296–1329.
- Stiglitz, J.**, “Self-Selection and Pareto-Efficient Taxation,” *Journal of Public Economics*, 1982, *17*, 213–240.
- , “Pareto Efficient and Optimal Taxation and the New New Welfare Economics,” in A.J. Auerbach and M. Feldstein, eds., *Handbook of Public Economics*, Vol. 2, Elsevier, 1987, pp. 991–1042.
- Ventry, D.J.**, “The Collision of Tax and Welfare Politics: The Political History of the Earned Income Tax Credit, 1969–99,” *National Tax Journal*, 2000, *53* (4), 983–1026.
- Werning, I.**, “Pareto Efficient Income Taxation,” Working paper 2007.
- Weymark, J.A.**, “Undominated directions of tax reform,” *Journal of Public Economics*, 1981, *16* (3), 343–369.
- Zorn, M.A.**, “Characterization of analytic functions in Banach spaces,” *Annals of Mathematics*, 1945, *46* (4), 585–593.

Online Appendix

A Proofs

A.1 Proof of Theorem 1 and Proposition 1.

To prove Theorem 1, we proceed in two steps. We first show that the non-existence of a Pareto-improving one-bracket reform implies that $\mathcal{R}(y) = R_{\tau\ell}^s(0, 0, y) \in [0, 1]$ for all y . We then show that the non-existence of a Pareto-improving two-bracket reform implies that $y \mapsto \mathcal{R}(y)$ is non-increasing.

Reforms with one bracket. Adapting inequality (4) to the case of a one-bracket reform, we find that a small reform with $\tau > 0$ that increases marginal tax rates from the status quo is Pareto-improving if, for some $\ell > 0$,

$$R_{\tau}^s(0, \ell, \hat{y}) - \ell > 0, \tag{A.1}$$

i.e., if marginal revenue gains are so large that even those agents are made better off whose tax bill increases by the maximal amount of $\max_y h^s(y) = \ell$. For a one-bracket reform with $\tau < 0$, we have to compare minus one times the derivative $R_{\tau}^s(0, \ell, \hat{y})$ with $\max_y [-h^s(y)] = 0$. Consequently, a small one-bracket reform that reduces marginal tax rates is Pareto-improving if

$$R_{\tau}^s(0, \ell, \hat{y}) < 0, \tag{A.2}$$

so that a tax cut leads to larger tax revenues, a logic familiar from analyses of the Laffer curve.

Below, we exploit the fact that a one-bracket reform on a bracket of length zero does not affect tax revenue, i.e., $R_{\tau}^s(0, 0, \hat{y}) = 0$ for any \hat{y} . To see this, recall that the new tax schedule satisfies $T_1(y) = T_0(y)$ for any $y \leq \hat{y}$, and $T_1(y) = T_0(y) + \tau \ell$ for any $y \geq \hat{y} + \ell$. For a one-bracket reform on a bracket of length $\ell = 0$, the new tax schedule is thus identical to the status quo tax schedule, $T_1(y) = T_0(y)$ for all y , independent of the step size τ . Hence, a variation in τ neither affects the budget set C_1 nor individual behavior y^* , nor tax revenue $R^s(\tau, 0, \hat{y})$.

Lemma A.1

- (i) If $R_{\tau\ell}^s(0, 0, \hat{y}) - 1 > 0$ for some $\hat{y} \in \mathcal{Y}$, there exists a Pareto-improving one-bracket reform with $\tau > 0$ and $\ell > 0$.
- (ii) If $R_{\tau\ell}^s(0, 0, \hat{y}) < 0$ for some $\hat{y} \in \mathcal{Y}$, there exists a Pareto-improving one-bracket reform with $\tau < 0$ and $\ell > 0$.
- (iii) If there is no Pareto-improving reform, then $R_{\tau\ell}^s(0, 0, \hat{y}) \in [0, 1]$ for all \hat{y} .

Proof of Lemma A.1. As explained above, we have $R_{\tau}^s(0, 0, \hat{y}) = 0$. If $R_{\tau\ell}^s(0, 0, \hat{y}) > 1$, this implies that $R_{\tau}^s(0, \ell, \hat{y}) - \ell$ turns positive if, starting from $\ell = 0$, the length of the interval is slightly increased. This proves (i). Analogously, if $R_{\tau\ell}^s(0, 0, \hat{y}) < 0$, this implies that $R_{\tau}^s(0, \ell, \hat{y})$ turns negative if, starting from $\ell = 0$, the length of the interval is slightly increased. This proves (ii). Thus, necessary conditions for the non-existence of a Pareto-improving one-bracket reform are $R_{\tau\ell}^s(0, 0, \hat{y}) \leq 1$ and $R_{\tau\ell}^s(0, 0, \hat{y}) \geq 0$. This proves (iii).

Reforms with two brackets. Lemma A.1 above gives sufficient conditions for the existence of Pareto-improving reforms with a single bracket. Proposition 1 gives the analogue for the case of two-bracket reforms. In particular, it shows that, if $y \mapsto \mathcal{R}(y) = R_{\tau\ell}^s(0, 0, y)$ is increasing, the combination of two reforms – each of which would not be Pareto-improving on a stand alone basis – yields a Pareto improvement. For this purpose, we denote by $R(\tau, h_2)$ the change in tax revenue due to a joint reform with two brackets, where $h_2 = h_1^s + h_2^s$ is composed of two single-bracket reforms.

Proof of Proposition 1. Fix two income levels y_1 and y_2 such that $y_2 > y_1$ and $R_{\tau\ell}^s(0, 0, y_2) > R_{\tau\ell}^s(0, 0, y_1)$. We now construct a Pareto-improving two-bracket reform with the parameters $\{(y_1, \tau_1, \ell_1, y_2, \tau_2, \ell_2, \tau, \ell)\}$. In particular, let $\tau_1 < 0$, $\tau_2 > 0$, and $\tau_1 \ell_1 + \tau_2 \ell_2 = 0 > \tau_1 \ell_1$. This implies that $\max_y h_2(y) = 0$. By the

linearity of the Gateaux differential,⁴² we moreover find that

$$\begin{aligned} R_\tau(0, h_2) &= \tau_1 R_\tau^s(0, \ell \ell_1, y_1) + \tau_2 R_\tau^s(0, \ell \ell_2, y_2), \text{ and} \\ R_{\tau\ell}(0, h_2) &= \tau_1 \ell_1 R_{\tau\ell}^s(0, 0, y_1) + \tau_2 \ell_2 R_{\tau\ell}^s(0, 0, y_2) \\ &= \tau_2 \ell_2 [R_{\tau\ell}^s(0, 0, y_2) - R_{\tau\ell}^s(0, 0, y_1)] > 0. \end{aligned}$$

Hence, there exists $\hat{\ell} > 0$ such that $R_\tau(0, h_2) - \max_y h(y) > 0$ for all $\ell \in (0, \hat{\ell})$. Finally, by Equation (4), this implies that, for a reform as constructed above with $\ell \in (0, \hat{\ell})$, a small increase in τ is Pareto-improving.

Suppose that $R_{\tau\ell}^s(0, 0, y_1)$ and $R_{\tau\ell}^s(0, 0, y_2)$ are between 0 and 1. Then, there is no Pareto-improving one-bracket reform for incomes close to y_1 or close to y_2 . If $R_{\tau\ell}^s(0, 0, y_1) < R_{\tau\ell}^s(0, 0, y_2)$, however, there is still scope for a Pareto improvement that involves two brackets.

A.2 Proof of Theorem 2

A reform with an arbitrary number m of brackets can be characterized as a collection

$$\{(y_k, \tau \tau_k, \ell \ell_k)\}_{k=1}^m$$

of one-bracket reforms, where the marginal tax in the k th bracket is changed by $\tau \tau_k$ and length of the k th bracket is given by $\ell \ell_k$. As before, the parameters (τ, ℓ) determine the size of the reform and the overall revenue is denoted by $R(\tau, h_m)$. The following lemma states sufficient conditions for the existence of a Pareto-improving reform with m brackets.

⁴² Gateaux differentials are not linear in general. To clarify the conditions under which they are, for $0 < \bar{y}, a < \infty$, let $\tau \in [-a, a]$ and $h \in \mathcal{H} := (\mathcal{C}[0, \bar{y}], \|\cdot\|_\infty)$, where $\|\cdot\|_\infty$ denotes the sup norm. We define the operator

$$\bar{R} : \mathcal{H} \rightarrow \mathcal{K} : h \mapsto \bar{R}(\tau, h),$$

where $\mathcal{K} := (\mathcal{C}_b([-a, a]), \|\cdot\|_\infty)$ and $\mathcal{C}_b([-a, a])$ denotes the set of bounded continuous real functions defined on $[-a, a]$. Note that \mathcal{H} and \mathcal{K} are Banach spaces. In this setting, the Gateaux differential of $R_\tau(\tau, \cdot)$ is linear (Zorn (1945); Theorem 2.3).

Lemma A.2 Consider a collection $\{(y_k, \tau_k, \ell_k)\}_{k=1}^m$ of simple reforms. Let $\tau_0 l_0 = 0$. There is a reform (τ, h_m) with $\tau > 0$ and $\ell > 0$ that is Pareto-improving if

$$\sum_{k=1}^m \tau_k \ell_k \mathcal{R}(y_k) - \max_{j \in \{0,1,\dots,m\}} \sum_{k=0}^j \tau_k \ell_k > 0.$$

Proof The linearity of the Gateaux differential implies that

$$\begin{aligned} R_\tau(0, h_m) &= \sum_{k=1}^m \tau_k R_\tau^s(0, \ell_k, y_k), \text{ and} \\ R_{\tau\ell}(0, h_m) &= \sum_{k=1}^m \tau_k \ell_k \mathcal{R}(y_k). \end{aligned}$$

Moreover,

$$\max_y h_m(y) = \ell \max_{j \in \{0,1,\dots,m\}} \sum_{k=0}^j \tau_k \ell_k.$$

As shown above, $R_\tau(0, h_m)$ equals zero for a reform with $\ell = 0$ such that all brackets have length zero. Hence, if the condition in the lemma is satisfied, there exists $\hat{\ell} > 0$ such that $R_\tau(0, h_m) - \max_y h_m(y) > 0$ for all $\ell \in (0, \hat{\ell})$. By Equation (4), this implies that, for such an m -bracket reform (τ, h_m) with $\ell \in (0, \hat{\ell})$, a small increase in the step size τ is Pareto-improving.

Lemma A.2 states sufficient conditions for the existence of Pareto-improving reforms. If we limit attention to small reforms these conditions are also necessary, i.e., if they do not hold, there is no small reform that is Pareto-improving. The following lemma shows that, if the conditions in Theorem 2 hold, the condition in Lemma A.2 is violated for any collection of m single bracket reforms. Consequently, there is no small Pareto-improving m -bracket reform.

Lemma A.3 Suppose that the function $y \mapsto \mathcal{R}(y)$ is bounded from below by 0, bounded from above by 1 and non-increasing. Let $\tau_0 l_0 = 0$. Then,

$$\sum_{k=1}^m \tau_k \ell_k \mathcal{R}(y_k) - \max_{j \in \{0,1,\dots,m\}} \sum_{k=0}^j \tau_k \ell_k \leq 0 \tag{A.3}$$

for any collection $\{(y_k, \tau_k, \ell_k)\}_{k=1}^m$, and for any $m \geq 1$.

Proof of Lemma A.3. Let j^* be a bracket in which the function h_m achieves a maximum, $j^* := \operatorname{argmax}_j \sum_{k=0}^j \tau_k \ell_k$. Note that this implies that $\sum_{k=z}^{j^*} \tau_k \ell_k \geq 0$ for any $z \in \{0, \dots, j^*\}$ and $\sum_{k=j^*+1}^z \tau_k \ell_k \leq 0$ for any $z \in \{j^*+1, \dots, m\}$; otherwise j^* would not be a maximizer.

Step 1. We verify the following claim: Suppose that $j^* > 0$ and that

$$\sum_{k=z}^{j^*} \tau_k \ell_k \mathcal{R}(y_k) \leq \mathcal{R}(y_z) \sum_{k=z}^{j^*} \tau_k \ell_k \quad (\text{A.4})$$

holds for some $z \in \{1, \dots, j^*\}$. Then, if $z > 1$, we also have

$$\begin{aligned} \sum_{k=z-1}^{j^*} \tau_k \ell_k \mathcal{R}(y_k) &= \tau_{z-1} \ell_{z-1} \mathcal{R}(y_{z-1}) + \sum_{k=z}^{j^*} \tau_k \ell_k \mathcal{R}(y_k) \\ &\leq \tau_{z-1} \ell_{z-1} \mathcal{R}(y_{z-1}) + \underbrace{\mathcal{R}(y_z)}_{\leq \mathcal{R}(y_{z-1})} \underbrace{\sum_{k=z}^{j^*} \tau_k \ell_k}_{\geq 0} \\ &\leq \mathcal{R}(y_{z-1}) \sum_{k=z-1}^{j^*} \tau_k \ell_k . \end{aligned}$$

Condition (A.4) is obviously satisfied for $z = j^*$. Hence, a repeated application of the preceding argument yields

$$\sum_{k=1}^{j^*} \tau_k \ell_k \mathcal{R}(y_k) \leq \underbrace{\mathcal{R}(y_1)}_{\in [0,1]} \underbrace{\sum_{k=1}^{j^*} \tau_k \ell_k}_{\geq 0} \leq \sum_{k=1}^{j^*} \tau_k \ell_k = \sum_{k=0}^{j^*} \tau_k \ell_k . \quad (\text{A.5})$$

Step 2. An analogous argument implies that

$$\sum_{k=j^*+1}^m \tau_k \ell_k \mathcal{R}(y_k) \leq \mathcal{R}(y_m) \underbrace{\sum_{k=j^*+1}^m \tau_k \ell_k}_{\leq 0} \leq 0 . \quad (\text{A.6})$$

Step 3. Together (A.5) and (A.6) imply that, if $j^* \in \{1, \dots, m-1\}$,

$$R_{\tau\ell}(0, h_m) = \sum_{k=1}^{j^*} \tau_k \ell_k \mathcal{R}(y_k) + \sum_{k=j^*+1}^m \tau_k \ell_k \mathcal{R}(y_k) \leq \sum_{k=0}^{j^*} \tau_k \ell_k , \quad (\text{A.7})$$

which proves (A.3). Note that the cases $j^* = 0$ and $j^* = m$ are also covered. With $j^* = 0$, $\sum_{k=1}^{j^*} \tau_k \ell_k \mathcal{R}(y_k)$ does not enter the chain of inequalities and (A.6)

directly implies (A.3). With $j^* = m$, $\sum_{k=j^*+1}^m \tau_k \ell_k \mathcal{R}(y_k)$ does not enter and (A.5) directly implies (A.3).

While the previous arguments only refer to reforms with small brackets, $\ell \rightarrow 0$, they can easily be adjusted to cover reforms with large brackets as well. In particular, we note that the revenue implications of a reform with a bracket of length $\ell \ell_k > 0$, starting at income y_k , can be approximated arbitrarily well by the revenue implications of a reform with m evenly spaced small brackets in the interval $(y_k, y_k + \ell \ell_k)$, if m is chosen sufficiently large. The following proof of Corollary 1 makes this approximation argument explicit.

A.3 Proof of Corollary 1 and Proposition 3

Take any continuous reform direction h on $[0, \bar{y}]$. We approximate h with a piecewise linear reform direction h_m that involves m one-bracket reforms $(\tau_k, h_k^s)_{k=1}^m$, so that $h_m(y) = \sum_{k=1}^m \tau_k h_k^s(y)$. Throughout, we let $\ell = 1$ and divide the domain $[0, \bar{y}]$ into m brackets of equal length $\ell_k = \frac{1}{m}\bar{y}$, starting at incomes $y_1 = 0$, $y_2 = \frac{1}{m}\bar{y}, \dots$, and $y_m = \frac{m-1}{m}\bar{y}$. Thus, we have m adjacent brackets – a special case of our general formalism, which also allows for gaps between the brackets where marginal tax rates change. For any k , we then let

$$\tau_k = \frac{h(y_{k+1}) - h(y_k)}{\ell_k}, \quad \text{where we set } y_{m+1} = \bar{y}.$$

This yields an approximation of h by a piecewise linear function. The construction is illustrated in Figure A.1. By choosing m sufficiently large, the piecewise linear function h_m approximates h in the sense that, for any $\varepsilon > 0$, there exists $\hat{m}(\varepsilon)$ so that for any $m > \hat{m}(\varepsilon)$,

$$\sup_{y \in \mathcal{Y}} |h(y) - h_m(y)| < \varepsilon.$$

For later reference, we note that this implies in particular that, for any y^* that maximizes $h(y)$ over \mathcal{Y} , we have

$$h(y^*) - h_m(y^*) < \varepsilon. \tag{A.8}$$

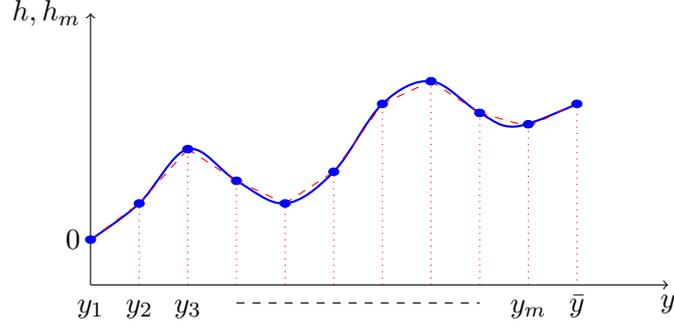


Figure A.1: Approximation of function h (solid, blue) by piecewise linear function h_m (dashed, red).

The Gateaux differential is linear in the reform direction h and hence continuous. We therefore have

$$\lim_{m \rightarrow \infty} R_\tau(0, h_m) = R_\tau(0, h), \quad (\text{A.9})$$

i.e., the Gateaux differential for direction h_m converges to the Gateaux differential for direction h . We now provide a characterization of $\lim_{m \rightarrow \infty} R_\tau(0, h_m)$. By the linearity of the Gateaux differential, we have

$$R_\tau(0, h_m) = \sum_{k=1}^m \tau_k R_\tau(0, h_k^s) = \sum_{k=1}^m \tau_k R_\tau^s(0, \ell_k, y_k).$$

For m large and $\ell_k = \frac{\bar{y}}{m}$ close to zero, a first-order Taylor approximation moreover gives

$$R_\tau^s(0, \ell_k, y_k) \approx \ell_k R_{\tau\ell}^s(0, 0, y_k) = \ell_k \mathcal{R}(y_k).$$

The approximation is perfect in the limit case $m \rightarrow \infty$ or, equivalently, $\ell_k = \frac{\bar{y}}{m} \rightarrow 0$. Therefore,

$$\begin{aligned} \lim_{m \rightarrow \infty} R_\tau(0, h_m) &= \lim_{m \rightarrow \infty} \sum_{k=1}^m \tau_k \ell_k \mathcal{R}(y_k) \\ &= \lim_{m \rightarrow \infty} \sum_{k=1}^m [y_{k+1} - y_k] \tau_k \mathcal{R}(y_k) \\ &= \int_{y \in \mathcal{Y}} h'(y) \mathcal{R}(y) dy, \end{aligned} \quad (\text{A.10})$$

where the last term is the Riemann integral that gives the marginal revenue effect of a reform in direction h , and $h' : y \mapsto h'(y)$ is the change of the marginal tax rate

at income y due to a unit increase in τ . To see this, note first that the term in the second line is the limit of a Riemann sum, the latter involving the step function $y \mapsto \tau_k \mathcal{R}(y_k)$ for $y \in [y_k, y_k + \ell_k]$. Second, note that both h and h_m are continuous functions on a compact interval. Hence, they have a bounded variation and are therefore differentiable almost everywhere. Moreover, we have $\lim_{m \rightarrow \infty} \ell_k = 0$ and therefore, for any y in the interior of bracket k ,

$$\tau_k = h'_m(y) \quad \xrightarrow{m \rightarrow \infty} \quad h'(y) .$$

This completes the derivation of Equation (7) in Proposition 3.

We now show that the conclusion of Theorem 2 extends to all continuous reform directions. We first note that (A.9) implies that for any $\varepsilon > 0$, there is an $\tilde{m}(\varepsilon) \in \mathbb{R}$ such that, for $m > \tilde{m}(\varepsilon)$

$$R_\tau(0, h) - R_\tau(0, h_m) < \varepsilon . \tag{A.11}$$

To complete the proof of Corollary 1, we proceed by contradiction. Suppose that the conditions of Theorem 2 apply and that there is a continuous reform direction h that is Pareto-improving, i.e., that satisfies

$$R_\tau(0, h) - \max_{y \in \mathcal{Y}} h(y) = 2\varepsilon ,$$

for some $\varepsilon > 0$. Then, by (A.8) and (A.11), for $m > \max\{\hat{m}(\varepsilon), \tilde{m}(\varepsilon)\}$, there is also an m -bracket reform such that $R_\tau(0, h_m) - \max_{y \in \mathcal{Y}} h_m(y) > 0$, i.e., that is Pareto-improving. But this is impossible by Theorem 2. The assumption that there is a Pareto-improving direction h in the class of continuous functions on $[0, \bar{y}]$ has therefore led to a contradiction and must be false.

A.4 Proof of Proposition 2

The equivalence of statements A and C in Proposition 2 follows from Theorems 1 and 2 and from Corollary 1. To complete the proof, we establish the equivalence of A and B .

Preliminaries First, as a preliminary step, we define a *solution to the inverse tax problem*. The inverse tax problem is based on the assumption that there is a social welfare function, $\mathcal{W}(\tau, h) := \mathbf{E}[G(v(\tau, h, \theta), \theta)]$, where $G : \mathbb{R} \times \Theta \rightarrow \mathbb{R}$ is increasing in its first argument. Upon using (3), the Gateaux differential of social welfare in direction h can be written as

$$\mathcal{W}_\tau(0, h) = \bar{g} \left\{ R_\tau(0, h) - \frac{1}{\bar{g}} \mathbf{E}_y [g(y) h(y)] \right\}, \quad (\text{A.12})$$

where the operator \mathbf{E}_y indicates the computation of a population average using the status-quo income distribution, which we take to be represented by the *cdf* F_y , function $y \mapsto g(y)$ is defined by $g(y) := \mathbf{E}[G'(v(0, h, \theta), \theta) u_c(\cdot) \mid y_0(\theta) = y]$, and \bar{g} is the population average $\bar{g} := \mathbf{E}_y [g(y)]$.

The interpretation is that $g(y)$ gives the marginal social benefit from increasing the consumption of people with earnings level y , whereas \bar{g} gives the marginal social benefit from increasing every one's consumption. The expression $\mathbf{E}_y [g(y) h(y)]$ interacts the marginal benefits with the way in which people's tax burdens change due to the reform. We say that the function $g : y \mapsto g(y)$ is a *solution to the inverse tax problem* if $\mathcal{W}_\tau(0, h) = 0$, for any reform direction h .

Equivalence of A and B. We show the equivalence of statements *A* and *B* by considering three types of reforms: One-bracket reforms with a tax cut or tax hike, and two-bracket tax cuts. First, consider a small single-bracket reform with a bracket of length ℓ , starting at income y_1 . To capture the welfare implications of this single-bracket reform, we write $\mathcal{W}^s(\tau, \ell, y_1)$ rather than $\mathcal{W}(\tau, h^s)$. Using this notation, the welfare implications of a single-bracket reform with τ and ℓ close to zero are given by the cross derivative

$$\mathcal{W}_{\tau\ell}(0, 0, y_1) = \bar{g} \left\{ \mathcal{R}(y_1) - \frac{1}{\bar{g}} [1 - F_y(y_1)] \mathbf{E}_y [g(y) \mid y > y_1] \right\}.$$

If the first-order condition for welfare-maximization is satisfied, $\mathcal{W}_{\tau\ell}(0, 0, y_1) = 0$, then we have

$$\mathbf{E}_y [g(y) \mid y > y_1] = \bar{g} \frac{\mathcal{R}(y_1)}{1 - F_y(y_1)}.$$

Hence, on average, the implicit welfare weights of people with incomes above y_1 are negative on average if and only if the revenue function is negative at income y_1 . In this case, a small one-bracket tax cut for incomes close to y_1 is Pareto-improving by Lemma A.1. Note that the formula above generalizes Equation (7) in Lorenz and Sachs (2016), according to which a negative average welfare weight above some income implies an inefficiency in the tax-transfer system.

Analogously, for a small one-bracket tax increase at income level y_1 , the first-order condition can be shown to imply that

$$\mathbf{E}_y [g(y) \mid y \leq y_1] = \bar{g} \frac{1 - \mathcal{R}(y_1)}{F_y(y_1)},$$

where we use that $\bar{g} = F_y(y_1) \mathbf{E}_y [g(y) \mid y \leq y_1] + [1 - F_y(y_1)] \mathbf{E}_y [g(y) \mid y > y_1]$. Hence, the implicit welfare weights of people with incomes below y_1 are negative on average if and only if the revenue function attains a value above 1 at income y_1 . In this case, a small one-bracket tax increase for incomes close to y_1 is Pareto-improving by Lemma A.1.

Finally, consider a small two-bracket tax cut between incomes y_1 and $y_2 > y_1$. For τ and ℓ close to zero, the welfare implications of such a two-bracket reform h_2 with $\tau_1 < 0$, $\tau_2 > 0$ and $\tau_1 \ell_1 + \tau_2 \ell_2 = 0$ are given by $\mathcal{W}_{\tau\ell}(0, h_2)$ where

$$\begin{aligned} \mathcal{W}_{\tau\ell}(0, h_2) &= \bar{g} \left\{ \tau_1 \ell_1 \mathcal{R}(y_1) + \tau_2 \ell_2 \mathcal{R}(y_2) \right. \\ &\quad \left. - \frac{1}{\bar{g}} \tau_1 \ell_1 [1 - F_y(y_1)] \mathbf{E}_y [g(y) \mid y > y_1] \right. \\ &\quad \left. - \frac{1}{\bar{g}} \tau_2 \ell_2 [1 - F_y(y_2)] \mathbf{E}_y [g(y) \mid y < y_2] \right\} \\ &= \bar{g} \tau_2 \ell_2 \left\{ \mathcal{R}(y_2) - \mathcal{R}(y_1) \right. \\ &\quad \left. + \frac{1}{\bar{g}} [F_y(y_2) - F_y(y_1)] \mathbf{E}_y [g(y) \mid y \in (y_1, y_2)] \right\}. \end{aligned}$$

If such a reform satisfies the first-order condition, $\mathcal{W}_{\tau\ell}(0, h_2) = 0$, then the solution to the inverse tax problem is defined by

$$\mathbf{E}_y [g(y) \mid y \in (y_1, y_2)] = \bar{g} \frac{\mathcal{R}(y_1) - \mathcal{R}(y_2)}{F_y(y_2) - F_y(y_1)}.$$

Hence, the implicit welfare weights of people with incomes between y_1 and y_2 are negative on average if and only if the revenue function $\mathcal{R} : y \mapsto \mathcal{R}(y)$ is increasing

so that $\mathcal{R}(y_2) > \mathcal{R}(y_1)$. By Proposition 1, this implies that a small two-bracket tax cut between incomes y_1 and $y_2 > y_1$ is Pareto-improving.

To sum up, if and only if one of the conditions listed in statement A is violated, the implicit welfare weights of people in some part of the income distribution are negative. If the average implicit weights above and below each income level, and in each subset of $[0, \bar{y}]$ are positive instead, then function $g(y)$ is bounded from below by 0.

Finally, note that, if the revenue function is continuously differentiable, the solution to the inverse tax problem can be characterized point-wise, with $g(y) = -\bar{g} \mathcal{R}'(y)/f_y(y)$ at each income level in the support of the *pdf* of the income distribution f_y . Differentiability of \mathcal{R} cannot be taken for granted, however. For instance, our analysis of the EITC below gives rise to discontinuous revenue functions. Our formal arguments above accommodate this possibility.

Equivalence of B and C . By Theorem 1, a Pareto-improving reform exists if one of the conditions in statement A is violated ($C \Rightarrow A$). By Theorem 2 and Corollary 1, there is no Pareto-improving reform direction if these conditions are satisfied ($A \Rightarrow C$). In combination with the proof given above, this implies that B and C are equivalent as well: Positive implicit welfare weights are both necessary and sufficient for the non-existence of Pareto-improving reform directions.

B The optimal reform direction

Fix a status quo tax-transfer system (T_0, c_0) and a corresponding revenue function $\mathcal{R}(y)$. By Proposition 3, a reform direction h is Pareto-improving if it generates a strictly positive *free lunch*

$$\Pi(h) := \int_0^{\bar{y}} h'(y) \mathcal{R}(y) dy - \max_{y \in y_0(\Theta)} \int_0^y h'(z) dz . \quad (\text{B.1})$$

Then, a small reform in direction h raises more tax revenue than what is needed to compensate the agents facing the largest tax increase. In the following, we solve for the reform direction h^* that maximizes the free lunch Π over the set of

functions such that $h' : [0, \bar{y}] \rightarrow [-a, a]$ for some fixed $a > 0$. To simplify the exposition, we impose further assumptions that are satisfied in the context of our application in Section 3 (see Figure 7).

Assumption 1 *Let $y_0(\Theta) = [0, \bar{y}]$. There is a unique triplet (y_s, y_t, r) with $0 < y_s < y_t < \bar{y}$ and $r \in (0, 1)$ with the following properties:*

- (i) \mathcal{R} is strictly decreasing on $(0, y_s)$ and on (y_t, \bar{y}) ,
- (ii) $\mathcal{R}(y_s) = \mathcal{R}(y_t) = r$,
- (iii) $\mathcal{R}(y) \in (r, 1)$ for each $y \in (0, y_s)$ and for each $y \in (y_m, y_t)$, for $y_m := \frac{y_s + y_t}{2}$,
- (iv) $\mathcal{R}(y) < r$ for each $y \in (y_s, y_m)$,
- (v) $\mathcal{R}(y) \in [0, r)$ for each $y \in (y_t, \bar{y})$,

Proposition B.1 *Fix $a > 0$. Under Assumption 1, the optimal reform direction h^* is, for any $a > 0$, given by a two-bracket tax cut with*

$$h^{*'}(y) = \begin{cases} 0 & \text{for } y \in [0, y_s) , \\ -a & \text{for } y \in (y_s, y_m) , \\ a & \text{for } y \in (y_m, y_t) , \\ 0 & \text{for } y \in (y_t, \bar{y}] . \end{cases} \quad (\text{B.2})$$

Moreover, $\Pi(h^*) = aI(T_0, c_0)$ with

$$I(T_0, c_0) = \int_{y_m}^{y_t} \mathcal{R}(y)dy - \int_{y_s}^{y_m} \mathcal{R}(y)dy . \quad (\text{B.3})$$

Proof To solve for the optimal reform direction, we proceed in three steps. First, we show that $\max_{y \in y_0(\Theta)} h^*(y)$ equals zero. Second, we solve for the reform that maximizes $\Pi(h)$ subject to (i) $\max_{y \in y_0(\Theta)} h(y) = 0$, (ii) $h'(y) \in [-a, a]$ for all $y \in [y_s, y_t)$, and (iii) the additional restriction that $h'(y) = 0$ for all incomes below y_s and above y_t . Third, we show that the solution to this more restricted problem also solves the original maximization problem.

Step 1. The normalization that system $T(0) = 0$ for any tax system that we consider also implies $h(0) = 0$ for any reform direction that we consider. Therefore $\max_{y \in y_0(\Theta)} h^*(y) \geq 0$. To show that $\max_{y \in y_0(\Theta)} h^*(y) = 0$, we provide a proof by contradiction. For this purpose, assume that there is some $\varphi > 0$ such that $\max_{y \in y_0(\Theta)} h^*(y) = \varphi$ and denote by y^* the lowest income level such that $h^*(y) = \varphi$. Then, $y^* > 0$ and there must be an income $y' \in (0, y^*)$ such that $h^*(y) > 0$ for all $y \in (y', y^*)$. Consider a perturbed reform h_ε such that $h'_\varepsilon(y) = h^{*'}(y) - \varepsilon$ for all incomes (y', y^*) and $h'_\varepsilon(y) = h^{*'}(y)$ for all other incomes. The free lunch from this perturbed reform is

$$\Pi(h_\varepsilon) = \int_0^{\bar{y}} h'_\varepsilon(y) \mathcal{R}(y) dy - \varphi + \varepsilon(y^* - y') .$$

The derivative of $\Pi(h_\varepsilon)$ with respect to ε is

$$\frac{d\Pi(h_\varepsilon)}{d\varepsilon} \Big|_{\varepsilon=0} = - \int_{y'}^{y^*} \mathcal{R}(y) dy + (y^* - y') > 0 ,$$

where the positive sign follows because $\mathcal{R}(y) < 1$ for any $y \in (0, \bar{y})$ by Assumption 1. This contradicts that assumption that Π obtains a maximum at h^* .

Step 2. Consider the problem to maximize $\Pi(h)$ over the set of functions h such that (i) $h'(y) \in [-a, a]$ for all $y \in (y_s, y_t)$, (ii) $h'(y) = 0$ for all $y \leq y_s$ and all $y \geq y_t$ and (iii) $h(y) = \int_0^y h'(z) dz = \int_{y_s}^y h'(z) dz \leq 0$ for all $y \in (y_s, y_t)$. Note that the function given in (B.2) satisfies these constraints.

We now consider a Lagrangian for a more relaxed problem that takes only the constraint $h(y_t) \leq 0$ into account. We argue below that a solution to this relaxed problem satisfies (i)-(iii).

$$\mathcal{L}(t) = \int_{y_s}^{y_t} h'(y) \mathcal{R}(y) dy - \mu \int_{y_s}^{y_t} h'(y) dy ,$$

where μ is a Lagrange multiplier. The solution to this restricted problem is given by a function $\tilde{h} : (y_s, y_t) \rightarrow [-a, a]$ and a value $\tilde{\mu}$ of the multiplier. For any $y \in (y_s, y_t)$, the derivative of \mathcal{L} with respect to $h'(y)$ is given by

$$\frac{\partial \mathcal{L}}{\partial h'(y)} = \mathcal{R}(y) - \tilde{\mu} .$$

As the Lagrangian is linear in each $h'(y)$, the solution involves $\tilde{h}'(y)$ equal to the lower bound $-a$ for all y such that $\mathcal{R}(y) < \tilde{\mu}$, and $\tilde{h}'(y)$ equal to the upper bound a for all y such that $\mathcal{R}(y) > \tilde{\mu}$. Under Assumption 1, this is only consistent with $\tilde{h}(y_t) = \int_{y_s}^{y_t} \tilde{h}'(y) dy = 0$ if $\tilde{\mu} = r = \mathcal{R}(y_s)$. Then, $\tilde{h}'(y) = -a$ for all $y \in (y_s, y_m)$ and $\tilde{h}'(y) = a$ for all $y \in (y_m, y_t)$. Hence, \tilde{h} equals the function given in (B.2), so it satisfies (i)-(iii). Consequently,

$$\Pi(\tilde{h}) = \int_{y_s}^{y_m} -a\mathcal{R}(y)dy + \int_{y_m}^{y_t} a\mathcal{R}(y)dy = aI(T_0, c_0) ,$$

with $I(T_0, c_0)$ given in (B.3). We also note that, as $\tilde{\mu}$ is strictly positive, the constraint $\tilde{h}(y_t) \leq 0$ is binding.

Step 3. It remains to show that we cannot increase Π further by allowing $h'(y) \in \{-a, a\}$ for incomes below y_s and above y_t , while respecting the constraint $\int_0^y h(z)dz \leq 0$ for all $y \in [0, \bar{y}]$. A repeated application of the arguments in Step 2, once for incomes below y_s , and once for incomes above y_t , exploiting the monotonicity of $y \mapsto \mathcal{R}(y)$ over these income ranges, shows that any candidate solution to this problem will take the form

$$h^{*'}(y) = \begin{cases} 0 & \text{for } y < y_\alpha , \\ -a & \text{for } y \in (y_\alpha, y_m) , \\ a & \text{for } y \in (y_m, y_\beta) , \\ 0 & \text{for } y > y_\beta . \end{cases}$$

where $y_\alpha \leq y_s$ and $y_\beta \geq y_t$. The constraint $h(\bar{y}) = \int_0^{\bar{y}} h^*(y)dy \leq 0$ is only satisfied if $y_\beta - y_m \leq y_m - y_\alpha$. Finally, choosing y_α and y_β to maximize $\Pi(h)$ subject to $y_\alpha \leq y_s$ and $y_\beta \geq y_t$ and $y_\beta - y_m \leq y_m - y_\alpha$, shows that $y_\alpha = y_s$ and $y_\beta = y_t$ is an optimal choice.

C Pareto efficiency when earnings are bounded away from zero and bounded from above

Suppose that $y \mapsto \mathcal{R}(y)$ is non-increasing over \mathcal{Y} , so that there is no scope for a Pareto improvement by a two-bracket reform. Then, by Theorem 2 and Corollary

1, the status quo tax system is Pareto-efficient if \mathcal{R} is bounded from above by 1 for the lowest incomes, and bounded from below by 0 for the highest incomes. Hence, our Pareto test only requires to verify one-bracket-efficiency at the very top and the very bottom. Under more restrictive assumptions, even the conditions for one-bracket-efficiency at the extremes become dispensable as we show in the following corollary. For this purpose, we denote by $y_{min} := \inf y_0(\Theta)$ the infimum of the income levels, and by $y_{max} := \sup y_0(\Theta)$ the supremum of the income levels chosen under the status quo tax policy.

Corollary 2 *If $y_{min} > 0$, $y_{max} < \bar{y}$, and $y \mapsto \mathcal{R}(y)$ is non-increasing, then there is no Pareto-improving direction in the class of continuous functions.*

Under the conditions of Corollary 2, the monotonicity of $y \mapsto \mathcal{R}(y)$ is sufficient for the non-existence of a Pareto-improving reform direction, i.e., there is no need to invoke the requirements that this function must be bounded from below by 0 and from above by 1. The conditions are that all types choose their incomes in some interior subset of \mathcal{Y} . In this case, as we show formally below, we have $\mathcal{R}(y) = 1$ for all $y \in [0, y_{min})$, and $\mathcal{R}(y) = 0$ for all $y \in (y_{max}, \bar{y}]$. Consequently, all three sufficient conditions in Theorem 2 are satisfied if $y \mapsto \mathcal{R}(y)$ is non-increasing.⁴³

The conditions in Corollary 2 hold, for instance, in a Mirrleesian model of income taxation with only intensive-margin responses when Inada conditions ensure positive and bounded incomes for everybody.⁴⁴ By contrast, there is typically a positive mass of taxpayers with zero income in models with an intensive and an extensive margin. In this case, the requirement of boundedness does not follow from the requirement of monotonicity. Put differently, with a mass of non-working people, there can exist Pareto-improving reforms with one bracket even if there is no Pareto-improving reform with two brackets.

⁴³Note that Corollary 2 provides a more compact expression of our sufficient conditions, but it should not be interpreted as showing that two-bracket reforms can achieve strictly more than one-bracket reforms. In particular, a two-bracket reform with one bracket below y_{min} or above y_{max} is economically equivalent to a one-bracket reform.

⁴⁴Technically, this also requires that the parameter \bar{y} in $\mathcal{Y} = [0, \bar{y}]$ is chosen so large that this upper bound does not interfere with individual choices.

C.1 Proof of Corollary 2

Recall that $y_0(\Theta)$ is the image function of y_0 , i.e., the set of income levels that are individually optimal for some type in Θ given the status quo tax system. The infimum of this set is denoted by y_{min} and the supremum by y_{max} .

Lemma C.1 *If $y_{min} > 0$, then $\mathcal{R}(\hat{y}) = 1$ for any $\hat{y} \in [0, y_{min})$.*

Proof Fix a one-bracket reform (τ, ℓ, \hat{y}) such that $\hat{y} \geq 0$, $\hat{y} + \ell < y_{min}$ and assume that $\tau \ell$ is close to zero. This implies that, at any income level $y \geq y_{min}$, the tax burden increases by $\tau \ell$. We now argue that there are no behavioral responses to such a reform. More specifically, we show that “no behavioral responses” is consistent with both utility-maximizing behavior and the government budget constraint. When there are no behavioral responses and $y_0(\theta) > \hat{y} + \ell$ for all θ , this implies that the change of aggregate tax revenue equals $R^s(\tau, \ell, \hat{y}) = \tau \ell$. Since this additional tax revenue is rebated lump-sum, this also implies that all taxpayers receive additional transfers of $\tau \ell$. Hence, for any income in $y^*(\Theta)$, the additional tax payment and the additional transfer cancel each other out, implying that taxpayers face the same budget set before and after this reform, $C_1(y) = C_0(y)$. Moreover, for $\tau \ell$ sufficiently small, incomes smaller than y_{min} remain dominated by $y_0(\theta) \geq y_{min}$ for each θ . From $R^s(\tau, \ell, \hat{y}) = \tau \ell$, we obtain $R_{\tau \ell}^s(0, 0, \hat{y}) = 1$, which completes the proof. \square

Figure C.1a illustrates these arguments: The solid blue line depicts the status quo budget set $C_0(y)$, the dashed blue line shows the upward shift by $\tau \ell$ in the post-reform budget set $C_1(y)$ for incomes below y_k . The red line, finally, shows an indifference curve of the lowest-earning type θ such that $y_0(\theta) = y_{min}$ before the reform. As apparent from Figure C.1a, the type continues to prefer y_{min} to any lower income and does not change her behavior as long as $\tau \ell$ is small enough.

Lemma C.2 *If $y_{max} < \bar{y}$, then $\mathcal{R}(y) = 0$ for any $y \in (y_{max}, \bar{y}]$.*

Proof Fix a one-bracket reform (τ, ℓ, \hat{y}) such that $\hat{y} \in (y_{max}, \bar{y})$ and $\tau \ell > 0$.

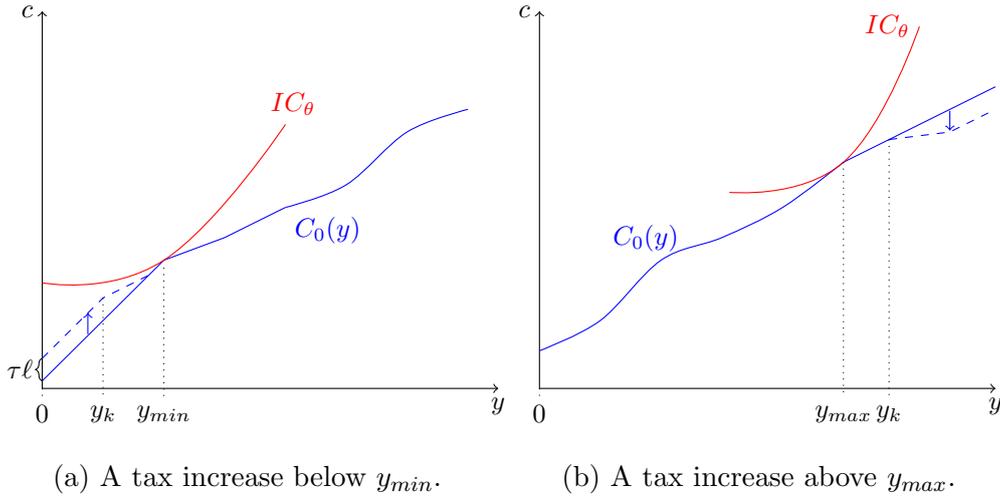


Figure C.1: Illustration of small one-bracket reforms below y_{min} and above y_{max} .

This reform does not change the tax burden at any income $y \leq y_{max}$. It increases the tax burden for incomes above \hat{y} that were already dominated by some income below y_{max} for each type θ prior to the reform. We now argue that there are no behavioral responses to such a reform. More specifically, we show that “no behavioral responses” is consistent with both utility-maximizing behavior and the government budget constraint. When there are no behavioral responses and $y_0(\theta) < \hat{y}$ for each θ , this implies that the aggregate tax revenue does not change, $R^s(\tau, \ell, \hat{y}) = 0$. This also implies that the base transfer is not changed, $c_1 = c_0$. Hence, for any income $y \leq \hat{y}$, taxpayers face the same budget set before and after this reform, $C_1(y) = C_0(y)$. Moreover, because of the tax increase for incomes larger than \hat{y} , these incomes remain dominated by some income below y_{max} for each type θ . From $R^s(\tau, \ell, \hat{y}) = 0$, we obtain $\mathcal{R}(\hat{y}) = R_{\tau\ell}^s(0, 0, \hat{y}) = 0$, which completes the proof. \square

Figure C.1b illustrates these arguments. Again, the solid blue line depicts the status quo budget set $C_0(y)$, and the dashed blue line depicts the downward shift in the budget set due to the reform. The red depicts an indifference curve of a type θ who chooses the highest income y_{max} before the reform. As apparent from Figure C.1b, the type continues to prefer y_{max} to any higher income and does not change her behavior after the reform.

Corollary 2 follows from combining Theorem 2 with Lemmas C.1 and C.2. Thus, if $y_{min} > 0$ and $y_{max} < \bar{y}$, monotonicity of \mathcal{R} is a sufficient condition for the non-existence of Pareto-improving tax reforms. The lower and upper bounds on \mathcal{R} become dispensable. Put differently, whenever there is a Pareto-improving reform with one bracket, there is also a Pareto-improving two-bracket tax cut.

Now suppose that the conditions in Corollary 2 are violated, so that either $y_{min} = 0$ or $y_{max} = \bar{y}$. In particular, the case with a mass of people with zero incomes is both empirically relevant and a typical outcome in the model with labor supply responses at the intensive and the extensive margin that we focus on in Section 3. In this case, it is possible that $\mathcal{R}(y)$ is monotonically decreasing over $(0, \bar{y})$ even though $\mathcal{R}(y) > 1$ for positive incomes close to 0. Then, there exists a Pareto-improving one-bracket reform, but no such reforms with two brackets. Hence, one-bracket efficiency at the bottom remains a substantive constraint.

By contrast, in the case $y_{max} = \bar{y}$, we can simply raise the upper threshold \bar{y} of the set of feasible incomes to make sure that all people choose incomes below \bar{y} . After this adjustment, $\mathcal{R}(y) = 0$ for all incomes between y_{max} and the adjusted level of \bar{y} by Lemma C.2. Then, $\mathcal{R}(y) < 0$ for some income level close to \bar{y} implies a violation of the condition that $y \mapsto \mathcal{R}(y)$ must be non-increasing over $[0, \bar{y}]$. Hence, whenever there is a Pareto-improving reform that reduces marginal taxes in one bracket of incomes below y_{max} , there is also a Pareto-improving two-bracket reform. In particular, any reform that combines the previously mentioned one-bracket tax cut with a tax increase in some bracket of incomes above y_{max} will do the job. It should be noted, however, that such a two-bracket reform is economically equivalent to a one-bracket reform by the arguments given in the proof of Lemma C.2: The second bracket has no effect on behavior as no type will choose an income in or above the second bracket.

Supplementary Material (not for publication)

D Empirical analysis

In this section, we first provide additional background information on the 1974 US tax-transfer system. Second, we conduct various robustness checks on the result that there was scope for a Pareto improvement by means of a two-bracket tax cut. Third, we refine our analysis by looking separately at single parents who differ in the number of children and discuss the desirability of tagging in this dimension. Finally, we analyze the US tax-transfer system for single parents in the post-reform years 1975 to 1978 and analyze the scope for further improvements of the EITC.

D.1 Data description and benchmark calibration

This section provides a description of our data and explains the choices for our benchmark calibration. We start with a description of the 1974 US tax-transfer system. Subsequently, we describe how we obtained estimates of the relevant income distributions and our benchmark assumptions on behavioral responses to tax reforms.

Status quo tax function: the US tax-transfer system in 1974. We take account of the federal income tax and two large welfare programs, Aid for Families with Dependent Children (AFDC) and Supplementary Nutrition Assistance Programs (SNAP, also called Food Stamps). The two latter were considered the most important welfare programs at the time and contributed most to high effective tax rates, see Moffitt (1979) and Moffitt (2003). As shown in Table S.1 based on CPS and PSID data, they also had much more recipients among and provided larger payments to low-income single parents than other programs. Moreover, they were eligible for all single parents, whereas other programs such as Social Security and the newly introduced Supplementary Security Income were restricted to aged, blind, and disabled individuals, see, e.g., Daly and Burkhauser (2003). AFDC was available only for single parents and varied to some extent across US

states. We focus on the AFDC rules in California. SNAP was a federal program that was available both for single parents and childless singles, but more generous for single parents. The transfer payments for single parents also depended on the number of children. In our benchmark analysis, we focus on the tax-transfer schedule for single parents with two children. In our data, the median number of children in single-parent households was two, and the arithmetic mean was about 2.2. Figure 4 in the main text depicts the effective tax rates for this subgroup of single parents and for childless singles. For the years 1975 and later, we also account for the Earned Income Tax Credit. Table S.2 depicts the sources that we use for computing the US tax-transfer systems for single parents and childless singles in the years 1974, 1975, and 1978.

Table S.1: Welfare transfers received by single parents in 1974.

| Transfer program | Sample | CPS | CPS | PSID | PSID |
|---|--------|------------|----------------|------------|----------------|
| | | any income | $y \leq 4,000$ | any income | $y \leq 4,000$ |
| Food stamps | Share | – | – | 49.3% | 71.0% |
| | Amount | – | – | \$369 | \$573 |
| AFDC | Share | 37.0% | 56.8% | 35.0% | 54.6% |
| | Amount | \$937 | \$1,501 | \$912 | \$1,561 |
| Supplementary Security Income (SSI) | Share | – | – | 3.5% | 5.6% |
| | Amount | – | – | \$58 | \$104 |
| Social security | Share | 12.4% | 17.4% | 14.6% | 19.4% |
| | Amount | \$363 | \$538 | \$411 | \$607 |
| Unemployment or workmen’s compensation | Share | – | – | 4.0% | 4.0% |
| | Amount | – | – | 70 | 99 |
| Other programs | Share | 10.2% | 11.3% | 8.9% | 14.3% |
| | Amount | \$155 | \$196 | \$210 | \$385 |

Notes: Table S.1 reports information on different welfare transfers to single parents in 1974, based on CPS-ASEC 1975 data and PSID 1975 data. Sample restrictions: Age 25-60, head of household, non-married, no partner or adult family member in household. Columns 2 and 4 show the shares of single parents receiving any positive transfers, and the average amounts of transfers received. Columns 3 and 5 show the shares of recipients and the average amounts received among single parents with earned income below 4,000 USD.

Source: Authors’ calculations (see Table S.2 for details).

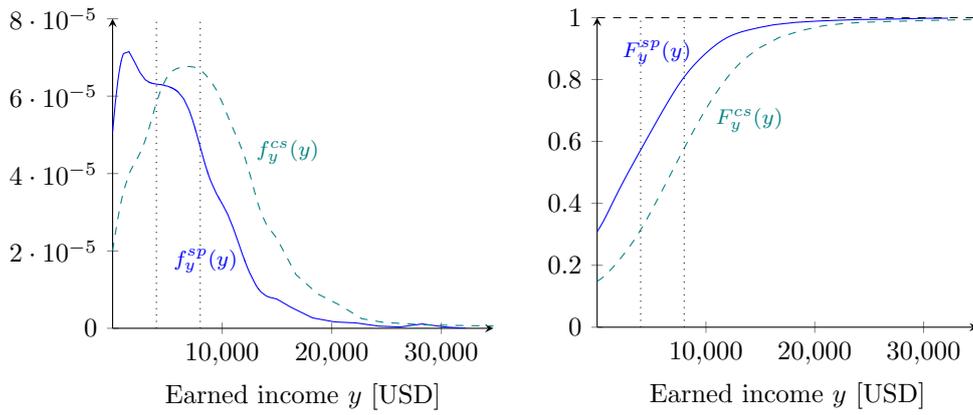
Income distribution. We estimate the 1974 income distributions based on Current Population Survey (CPS) data, using the Annual Social and Economic Supplement from the March 1975 wave (see Table S.3). We proceed in the same way for the years 1975 and 1978. For our benchmark estimates, we consider for each year the sample of non-married individuals aged 25 to 60 who do neither co-habit with an unmarried spouse nor with another adult family member. We estimate the earned income distributions for the subsamples of childless singles and single parents, respectively. For the benchmark analysis reported in the main text, we consider all singles with at least one qualifying child according to AFDC rules (i.e., either 15 years or younger, or 19 years and younger and in full-time education). In the subsequent sensitivity analysis, we also consider the subsets of single parents with one, two, and three children, and the somewhat larger set of singles with qualifying children according to EITC rules (i.e., either 19 years or younger, or 23 years and younger and in full-time education).

In line with the EITC rules, we consider as earned income the sum of (self-reported) wage income and self-employment income by all household members. In this sample, 30.9% of single parents and 14.8% of childless singles had zero or negative incomes, while 51.3% of single parents and 43% of childless singles had strictly positive incomes below 8,000 USD, the eligibility threshold of the EITC. For our benchmark analysis, we estimate the distributions of earned income for both groups using a non-parametric kernel density estimation with a Gaussian kernel. In the benchmark, we use bandwidths of 997 USD for single parents and 1,200 USD for childless singles, following Silverman’s rule. Figure S.1 shows the estimated pdf and the cdf of both income distributions.

Behavioral responses to taxation. We draw on a rich literature providing estimates of labor supply responses at the intensive and the extensive margin, see the discussions in Saez, Slemrod and Giertz (2012) or Chetty et al. (2013). Based on a meta-study and focusing on population-wide averages, Chetty et al. (2013) suggest an intensive-margin elasticity of labor supply with respect to the net-of-tax rate of 0.33, and an extensive-margin elasticity with respect to net labor income

Table S.2: Sources for US tax-transfer system, 1974-1978.

| Information | Years | Sources |
|-------------------|------------|--|
| Income tax | 1974-1975 | Internal Revenue Service, “Instructions for Form 1040”, years 1974, 1975. URLs: https://www.irs.gov/pub/irs_prior/i1040--1974.pdf ; www.irs.gov/pub/irs_prior/i1040--1975.pdf . |
| | 1974-1978 | Internal Revenue Service, “Statistics of Income, Individual income tax returns”, years 1974, 1975, 1978. URLs: https://www.irs.gov/pub/irs-soi/74inar.pdf ; https://www.irs.gov/pub/irs-soi/75inar.pdf ; www.irs.gov/pub/irs-soi/78inar.pdf . |
| EITC | 1975, 1978 | Tax Policy Center, “Earned Income Tax Credit Parameters, 1975-2021”. URL: https://www.taxpolicycenter.org/statistics/eitc-parameters . |
| AFDC | 1974-1978 | Office of the Assistant Secretary for Planning and Evaluation, “Aid to Families with Dependent Children. The Baseline”, 1998. URL: https://aspe.hhs.gov/basic-report/aid-families-dependent-children-baseline |
| | 1974 | US Department of Health, Education and Welfare, “Aid to Families with Dependent Children: Standards for Basic Needs, July 1974”, 1974. URL: https://hdl.handle.net/2027/mdp.39015088906634 . |
| | 1975, 1978 | TRIM3 project, “TRIM3 AFCD Rules”. Accessible at trim3.urban.org . |
| SNAP | 1974-1975 | US Bureau of the Census, “Characteristics of Households Purchasing Food Stamps. Current Population Reports”, Series 9-23, No. 61, 1976. URL: https://www2.census.gov/library/publications/1976/demographics/p23-061.pdf . |
| | 1978 | Federal Register Vol. 43, No. 95, May 16, 1978. Accessible at https://www.govinfo.gov/app/collection/fr . |
| SNAP, Payroll tax | 1974-1978 | Social Security Administration, “Annual Statistical Supplement. Section 2: Program Provisions and SSA Administrative Data”, 2010. URL: https://www.ssa.gov/policy/docs/statcomps/supplement/2010/2a1-2a7.html . |



(a) Probability density functions. (b) Cumulative distribution functions.

Figure S.1: Income distributions of single parents and childless singles, US 1974.

Notes: Figure S.1 shows the kernel estimates of the US income distributions among single parents (solid blue lines) and childless singles (dashed teal lines) in 1974. Panel (a) depicts the probability density functions; panel (b) depicts the cumulative distribution functions of both income distributions. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

Table S.3: Descriptive statistics for different samples.

| Sample | Sample size | Average age [yrs] | Share zero income [%] | Share EITC range [%] | Avg. earned income [USD] | Avg. capital income [USD] |
|---------------------------|-------------|-------------------|-----------------------|----------------------|--------------------------|---------------------------|
| 1: Mothers | 15,881 | 37.1 | 5.2 | 19.8 | 13,862 | 326 |
| 2: Single parents (BM) | 1,494 | 35.7 | 30.9 | 51.3 | 3,968 | 224 |
| 3: Single parents (EITC) | 1,875 | 37.5 | 26.3 | 53.2 | 4,432 | 248 |
| 4: Single parents (18-50) | 1,641 | 32.1 | 32.5 | 52.4 | 3,582 | 152 |
| 5: Single parents (Cal) | 187 | 35.7 | 25.8 | 52.2 | 4,341 | 101 |
| 6: Single parents (PSID) | 369 | 36.9 | 25.2 | 52.6 | 4,568 | 100 |
| 7: Childless singles | 3,407 | 43.4 | 14.8 | 43.1 | 7,355 | 425 |

Notes: Table S.3 shows descriptive statistics for different samples, based on CPS-ASEC 1975 data. Sample restrictions, unless otherwise stated: age 25-60, head of household, non-married, no partner or adult family member in household. Sample 1: any marital status, any household composition, at least one child qualifying for AFDC. Sample 2 (benchmark): at least one child qualifying for AFDC. Sample 3: at least one child qualifying for EITC. Sample 4: age 18-50. Sample 5: California resident, at least one child qualifying for AFDC. Sample 6: based on PSID 1975 data, at least one child aged 0-17. Sample 7: no dependent child.

Source: Authors' calculations (see Table S.2 for details).

of 0.25. Bargain, Dolls, Neumann, Peichl and Siegloch (2014) provide similar estimates for a sample of childless singles. For single parents, various studies find larger responses at the extensive margin. Specifically, Bastian (2020) estimates labor supply responses of single mothers to the 1975 EITC introduction, and finds an average participation elasticity of 0.58. Most earlier studies find similar or even larger participation responses by single mothers, e.g., Meyer and Rosenbaum (2001). By contrast, Kleven (2021) recently estimated a participation elasticity close to zero based on EITC reforms in the 1990s. Besides, several studies find that persons with little formal education and low incomes respond more strongly at the extensive margin than persons with higher education and higher incomes – see, e.g., Juhn, Murphy and Topel (1991), Juhn, Murphy and Topel (2002), Meghir and Phillips (2010). There is only limited empirical evidence on the relevance of income effects for labor supply. Recent evidence by Cesarini, Lindqvist, Notowidigdo and Östling (2017) based on Swedish lottery winners suggests (i) a marginal propensity to earn out of unearned income (MPE) of -0.08 , (ii) with about two thirds of

income effects arising at the intensive margin, one third at the extensive margin, and (iii) with little heterogeneity in income effects along the income distribution. Imbens, Rubin and Sacerdote (2001) report similar estimates with an MPE of -0.11 for the US, see also Holtz-Eakin, Joulfaian and Rosen (1993).

In our benchmark calibration, we assume an average participation elasticity of 0.58 for single parents and 0.25 for childless singles, and an intensive-margin elasticity of 0.33 for both subgroups, following Bastian (2020) and Chetty et al. (2013). Moreover, we assume that participation elasticities are decreasing with income in both groups, according to the function $\pi_0(y) = \pi_a - \pi_b (y/\tilde{y})^{1/2}$, where \tilde{y} equals 50,000 USD. Similar assumptions are employed by Jacquet et al. (2013) and Hansen (2021). For single parents, we assume that the participation elasticity falls from 0.67 at very low incomes to 0.4 at incomes above 50,000 USD (i.e., $\pi_a = 0.67$, $\pi_b = 0.27$), giving rise to an average value of 0.58. For childless singles, we assume π to fall from 0.4 to 0.1 (i.e., $\pi_a = 0.4$, $\pi_b = 0.3$), giving rise to an average value around 0.25. In the benchmark calibration, we leave out income effects. In the sensitivity analysis below, we consider a large range of alternative assumptions on labor supply elasticities.

D.2 Sensitivity analysis

In the following, we provide an extensive sensitivity analysis. We start by considering alternative assumptions on the behavioral responses to tax reforms. Then, we repeat our analysis using alternative estimates of the income distributions among single parents. Finally, we consider alternative representations of the US tax-transfer system. We find that our main results are robust to variations in all these dimensions: The 1974 US tax-transfer system for single parents was not Pareto-efficient, and there existed Pareto-improving tax reforms with two brackets similar to the 1975 EITC introduction.

Labor supply elasticities. Figure S.2 plots the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ under different assumptions on the extensive-margin elasticity π_0 . This

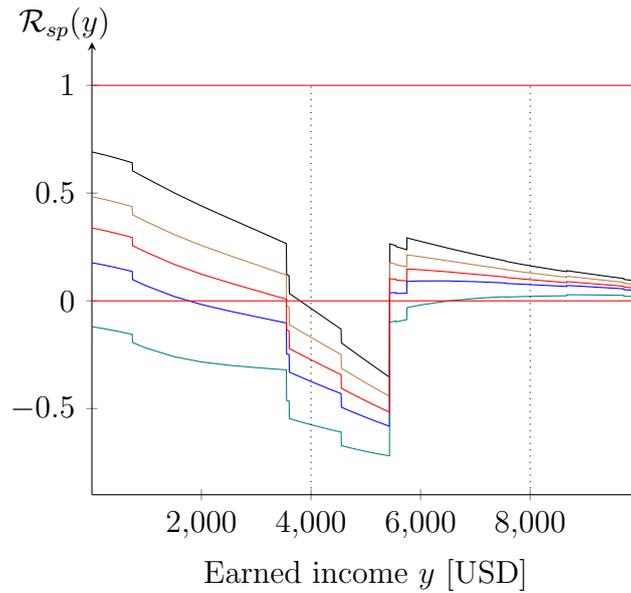


Figure S.2: Revenue function of the 1974 US tax-transfer system for single parents, different participation elasticities.

Notes: Figure S.2 shows the function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 1974, assuming different participation elasticities: constant at 0 (black line) or at 0.4 (red line), falling with income from 0.3 to 0.1 (brown line), falling from 0.67 to 0.4 (blue line), falling from 1 to 0.75 (teal line). The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

aspect deserves particular attention because the optimal tax literature has established the result that an EITC with negative marginal taxes can only be optimal with labor supply responses at the extensive margin, see Saez (2002) and Hansen (2021). Moreover, Kleven (2021) recently challenged the conventional view that the participation responses of single parents are particularly large. In Figure S.2, we therefore show the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for a range of different assumptions on participation elasticities. The blue line is our benchmark case in which the participation elasticity is falling from 0.67 at very low incomes to 0.4 at high incomes, giving an average elasticity of $\bar{\pi} = 0.58$. Additionally, the black line is a case where the elasticity is zero at all income levels; the brown line depicts a case where the elasticity is falling from 0.3 to 0.1 ($\bar{\pi} = 0.23$); the teal line shows a case where it is falling from 1 to 0.75 ($\bar{\pi} = 0.92$). Finally, the red line depicts a case with an elasticity of 0.4 at all income levels. In all these cases, the intensive-margin elasticity is held constant at $\varepsilon = 0.33$, as in our benchmark calibration.

The qualitative result is the same in all cases: for single parents, the 1974 US tax system was not Pareto-efficient. There existed Pareto-improving tax cuts with one and two brackets. This is even true in the limit case of vanishing labor supply responses at the extensive margin. Hence, the finding that an EITC-like two-bracket tax cut for low-income earners was Pareto-improving holds irrespective of what we assume about the strength of behavioral responses to taxation at the extensive margin. Quantitatively, though, they play a role. Higher participation elasticities imply larger revenue gains from a two-bracket tax cut; i.e., higher participation elasticities make an EITC-like reform more attractive (see Figure S.2).

Figure S.3 plots the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for different levels of the intensive-margin elasticity ε between 0 and 0.5, while holding the average participation elasticity at the benchmark level of 0.58. We find that the 1974 US tax system for single parents violated the monotonicity condition for Pareto efficiency for any level of the intensive-margin elasticity in this range. Again, this even re-

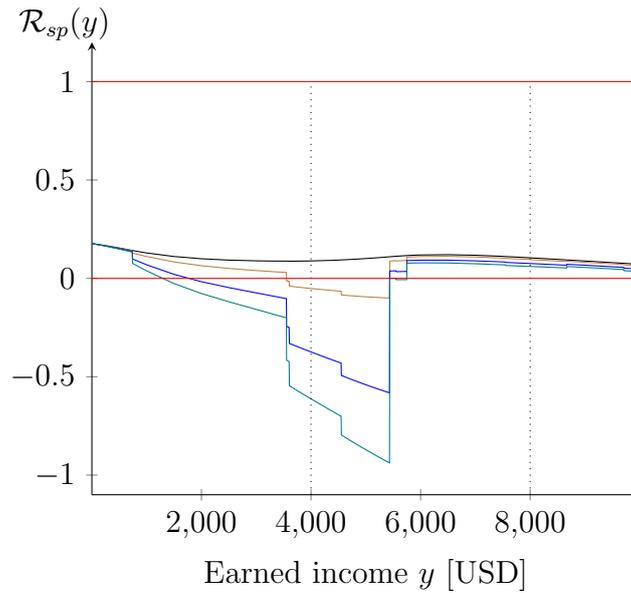


Figure S.3: Revenue function of the 1974 US tax-transfer system for single parents, different intensive-margin elasticities.

Notes: Figure S.3 shows the marginal revenue function $\mathcal{R}_{sp}(y)$ for single parents in 1974, assuming an intensive-margin elasticity of 0 (brown line), 0.1 (black line), 0.33 (blue line, benchmark) and 0.5 (teal line). The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

mains true in the limit case where labor supply does not respond at the intensive margin, i.e., for $\varepsilon = 0$. More specifically, we find the inefficiencies in the 1974 tax-transfer system to be quantitatively more pronounced with higher levels of the intensive-margin elasticity.

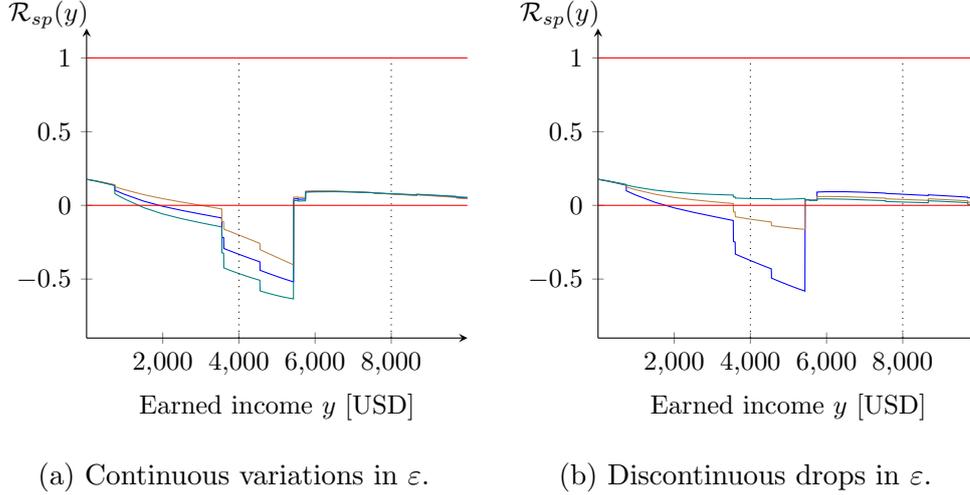


Figure S.4: Revenue function of the 1974 US tax-transfer system for single parents, heterogeneous intensive-margin elasticities.

Notes: Figure S.4a shows the marginal revenue function $\mathcal{R}_{sp}(y)$ for single parents in 1974, assuming that the intensive-margin elasticity ε varies with income: falling from 0.5 to 0.1 (brown line), or increasing from 0.1 to 0.5 (teal line), or equal to 0.3 at all income levels (blue line, benchmark). Figure S.4b shows function $\mathcal{R}_{sp}(y)$ under the assumption that the intensive-margin elasticity ε jumps in two steps from 0.13 over 0.33 to 0.73 (brown line), or from 0.03 over 0.33 to 0.93 (teal line), or equals 0.33 at all income levels (blue line, benchmark). In all cases, the average participation elasticity is held at the benchmark level of 0.58. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

In Figure S.4, we plot the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ under the assumption that the intensive-margin elasticity ε varies with income. The left panel compares the benchmark case where ε is equal to 0.3 at all income levels (blue line) with cases where ε is linearly increasing from 0.1 at the bottom to 0.5 at incomes above 15,000 USD (teal line), or decreasing from 0.5 at the bottom to 0.1 at incomes

above 15,000 USD (brown line). As can easily be seen, in all cases, the revenue function violates the monotonicity condition from Theorem 1 so that there is room for a Pareto improvement by means of a two-bracket tax cut. In an attempt to identify the limits of this robustness analysis, we also explore some extreme cases in which the elasticity ε varies with income in a discontinuous way. Specifically, the right panel compares the benchmark case where ε is equal to 0.33 at all income levels (blue line) with two extreme cases. The brown line shows a case with an elasticity of 0.13 at the bottom, an upward jump to 0.33 at income level 5,433 USD, and a further jump to 0.73 at 5,478 USD. Our qualitative results remain unchanged in this case. The teal line shows a case where both upward jumps are even more pronounced, from 0.03 at the bottom over 0.33 to 0.93 at the top. In this second case, the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ gets close to satisfying the monotonicity condition for a Pareto-efficient tax schedule (see Figure S.4b).

Income effects. Finally, we consider a model in which income effects in labor supply arise due to concave consumption utility

$$u(c, y, \omega, \gamma) = \psi(c) - k(y, \omega) - \gamma \mathbb{1}_{y>0},$$

with $\psi'(c) > 0$, $\psi''(c) < 0$ for all $c > 0$. With income effects, our sufficient-statistics formula for the revenue function $y \mapsto \mathcal{R}(y)$ becomes

$$\begin{aligned} \mathcal{R}(y) = & (1 - \mathcal{M}_0)^{-1} \left\{ 1 - F_y(y) - f_y(y) y \varepsilon(y) \frac{T'_0(y)}{1 - T'_0(y)} \right. \\ & \left. - \int_y^{\bar{y}} f_y(y') \pi(y') \frac{T_0(y')}{y' - T_0(y')} dy' - \int_y^{\bar{y}} f_y(y') \eta(y') T'_0(y') dy' \right\}, \end{aligned} \quad (\text{D.1})$$

where the multiplier term \mathcal{M}_0 is given by

$$\mathcal{M}_0 = \int_0^{\bar{y}} f_y(y) \eta(y) T'_0(y) dy + \int_0^{\bar{y}} f_y(y) \nu(y) T_0(y) dy. \quad (\text{D.2})$$

In this formula, $\eta(y)$ captures income effects at the intensive margin (i.e., the effect of an increase of the base transfer c_0 on the earnings of workers with pre-reform income y); $\nu(y)$ captures income effects at the extensive margin (i.e., the effect of an increase of the base transfer c_0 on the employment rate of agents who choose to earn y when becoming active on the labor market). With concave consumption

utility and additively separable preferences, $\eta(y)$ and $\nu(y)$ are both negative. Our benchmark case with quasi-linear utility is nested with $\eta(y) = \nu(y) = 0$. We formally derive a version of formula (D.1) in part E of the Supplementary Material below. Specifically, Proposition E.1 gives the revenue function $y \mapsto \mathcal{R}(y)$ in terms of the model's primitives based on a general framework that allows for complementarities between consumption and leisure.

By (D.1), income effects matter for the revenue effect of a one-bracket tax increase at income y in two ways. First, they increase labor supply at the intensive margin for all agents with income above y , for whom the reform has reduced net income directly. This effect is captured by the last term in (D.1). With positive marginal taxes, this implies an upward shift of the revenue function. Second, when the revenue gain is rebated lump sum, this spurs further income effects at both margins all across the income distribution. This is captured by the term \mathcal{M}_0 in (D.2). With positive marginal taxes and positive participation tax rates, income effects make \mathcal{M}_0 negative and decrease the multiplier $(1 - \mathcal{M}_0)^{-1}$. Consequently, relative to the benchmark without income effects, the revenue function $y \mapsto \mathcal{R}(y)$ shifts towards the horizontal axis, i.e., its absolute value is decreased.

Figure S.5 compares the revenue function $\mathcal{R}_{sp}(y)$ for our benchmark calibration to two scenarios with income effects. The teal line shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for a case with $\eta(y) = -0.08$ and $\nu(y) = -8 \times 10^{-6}$, corresponding to an MPE of -0.1 close to the estimates by Cesarini et al. (2017) and Imbens et al. (2001). The brown line shows $y \mapsto \mathcal{R}_{sp}(y)$ for a larger MPE of -0.4 , corresponding to $\eta(y) = -0.24$ and $\nu(y) = -2.4 \times 10^{-5}$. The blue line is the benchmark calibration without income effects. In all cases, we maintain our benchmark assumptions about substitution effects at both margins, i.e., $\varepsilon = 0.33$ and $\bar{\pi} = 0.58$. Figure S.5 demonstrates that our conclusion on the desirability of an EITC is robust to the inclusion of income effects.

Summing up, our sensitivity analysis demonstrates that our main result does not hinge on the specific assumptions on labor supply responses. The 1974 US tax-transfer system for single parents was not Pareto-efficient and there existed

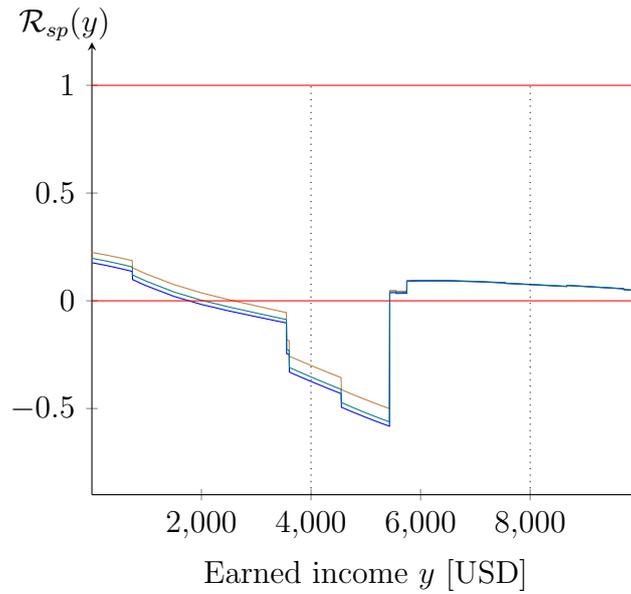


Figure S.5: Revenue function of the 1974 US tax-transfer system for single parents, accounting for income effects.

Notes: Figure S.2 shows the revenue function for single parents $y \mapsto \mathcal{R}_{sp}(y)$ in 1974 as a function of earned income, assuming income effects at both margins with an MPE of -0.08 (teal line) or -0.24 (brown line). For comparison, the blue line shows the benchmark case without income effects. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

Pareto-improving reforms akin to the introduction of the EITC. This conclusion is obtained for a wide range of empirically plausible assumptions about labor supply responses at the intensive and extensive margin.

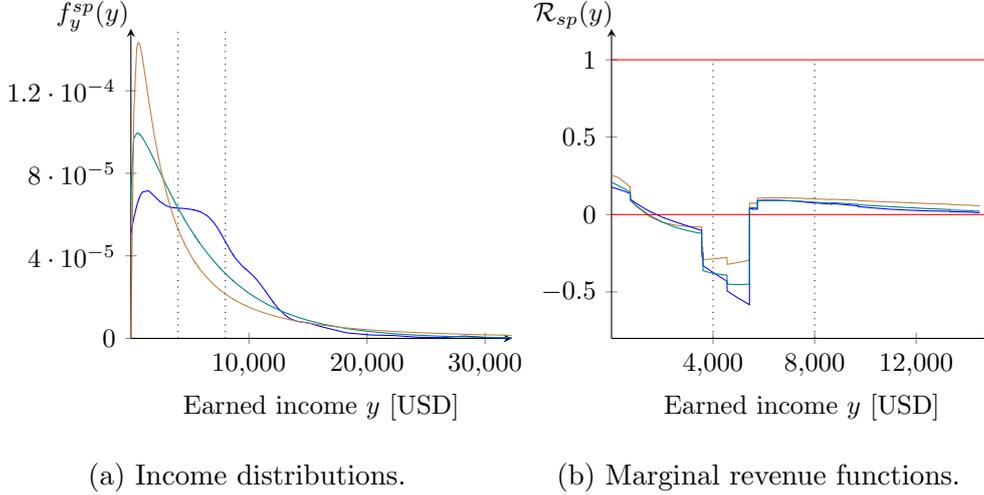


Figure S.6: Revenue function of 1974 US tax-transfer system for single parents, parametric income distributions.

Notes: Figure S.6a shows alternative estimates of the income distribution for single parents in 1974: kernel estimation (blue, benchmark), gamma distribution (teal), log-normal distribution (brown). Figure S.6b shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 1974, based on the kernel estimation (blue), gamma distribution (teal), and lognormal distribution (brown), using benchmark values for the elasticities at both margins. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

Alternative estimates of the income distribution. As alternatives to our benchmark calibration, we first verified that variations in the bandwidth choice for our kernel estimates do not substantially affect our results. We also fitted gamma and lognormal distributions to the earned income data in the March 1975 CPS. Figure S.6 shows versions of the revenue function that result with these parametric distributions. The differences to the benchmark case are small.

For the benchmark analysis reported in the main text, we estimate the earned

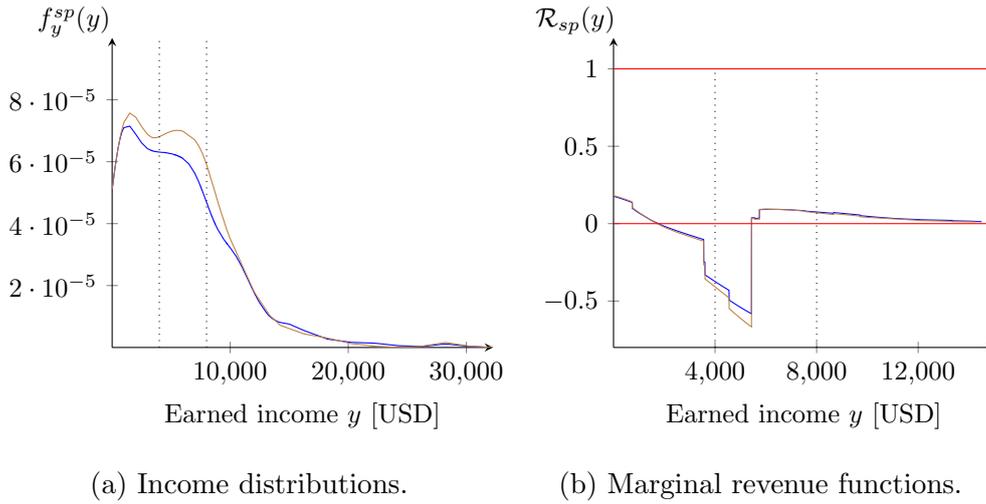


Figure S.7: Revenue function of 1974 US tax-transfer system for single parents, different samples.

Notes: Figure S.7a shows the pdfs of the distributions of earned income based on the samples of all single parents (blue, benchmark) versus single parents with exactly two children (brown line) on the basis of the March 1975 CPS. Figure S.7b gives the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 1974, based on the estimated income distributions for all single parents (blue) versus single parents with exactly two children (brown), using benchmark values for the elasticities at both margins. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

income distribution based on the set of all single parents ($N = 1,494$), with no conditioning on the number of children. As an alternative, we also considered non-parametric kernel estimates of the income distribution for the smaller sample of single parents with exactly two children ($N = 453$). Figure S.7 shows that this does not make much of a difference.

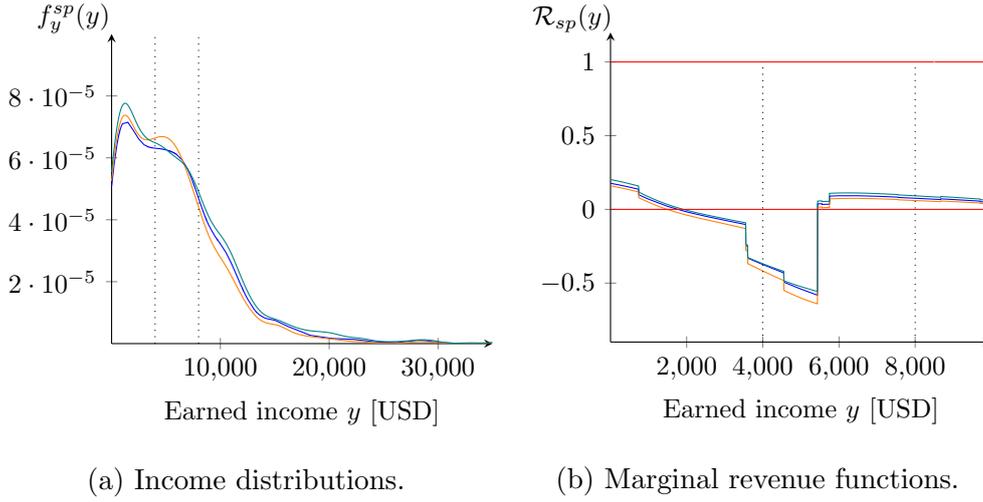


Figure S.8: Revenue function of 1974 US tax-transfer system for single parents, further samples.

Notes: Figure S.8a shows the pdfs of the distributions of earned income based on the samples of single parents aged 25 to 60 with children qualifying for AFDC (blue, benchmark) versus single parents aged 18 to 50 with children qualifying for AFDC (orange line) and single parents aged 25 to 60 with children qualifying for the EITC on the basis of the March 1975 CPS. Figure S.8b gives the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 1974, based on the estimated income distributions for the benchmark sample of single parents (blue) versus single parents aged 18 to 50 (orange) and single parents with children qualifying for the EITC, using benchmark values for the elasticities at both margins. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

We also calculated the revenue functions for the 1974 US tax system with income distributions based on different age restrictions (single parents aged 18 to 50 versus 25 to 60 in the benchmark case) and child age restrictions (children

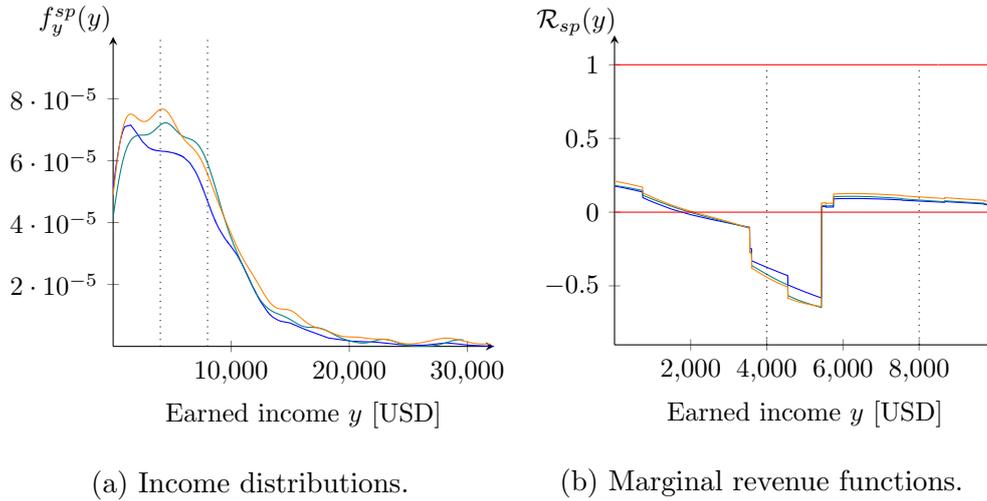


Figure S.9: Revenue function of 1974 US tax-transfer system for single parents, PSID data.

Notes: Figure S.9a shows the pdfs of the distributions of earned income among single parents based on 1975 ASEC-CPS data (blue, benchmark) versus two samples of single parents based on 1975 PSID data (teal and orange lines). Figure S.9b gives the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 1974, using the estimated income distributions based on 1975 ASEC-CPS data (blue, benchmark) versus the same samples based 1975 PSID data (teal and orange line). In both figures, the teal line is for single parents aged 25-60 years without any adults in the same household; the orange line for single parents aged 25-60 without any related adults in the same household. The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

qualifying for the EITC versus children qualifying for AFDC in the benchmark), and with estimated income distributions based on PSID data instead of CPS data as in our benchmark. The alternative samples based on CPS data are somewhat larger than our benchmark sample, whereas the PSID sample is much smaller; see Table S.3 for descriptive statistics. Crucially, Figures S.8 and S.9 show that these variations do not make much of a difference.

Alternative representations of the tax-transfer system. As final robustness checks, we consider three alternative representations of the 1974 US tax-transfer system for single parents. In all three cases, our results are virtually unchanged. First, we take into account the wealth and assets tests for welfare recipients. In 1974, only families with liquid assets below 1,600 USD were eligible for welfare transfers (AFDC, SNAP). While the amount of liquid assets is not reported in CPS data, we can approximate it from the reported capital (investment) income using the procedure suggested by Giannarelli (1992). Specifically, this requires to assume that all assets generate income and have the same rate of return equal to 0.06. Hence, a single parent failed the assets test if she had an annual capital income above 96 USD. Based on this approximation, about 8% of single parents with earned incomes below the relevant labor income thresholds failed the programs' assets test.⁴⁵ Moreover, single parents were only eligible for the EITC if the sum of their earned income and capital income was below 8,000 USD. In the March 1975 CPS data, only less than 1% of the single parents with incomes in the EITC range lost eligibility due to high capital incomes. Hence, we find that about 7% of single parents were eligible for the EITC, but not for welfare programs.

Our benchmark analysis in the main text ignores the assets test. Thereby, we provide an answer to the question whether, in 1974, introducing an EITC with the same assets tests as AFDC and SNAP would have been Pareto-improving. Al-

⁴⁵According to CPS data, the average capital income in 1974 was 224 USD for all single parents, but only 67 USD for single parents with earned incomes below 8,000 USD. About 85% of the latter persons did not have any capital income.

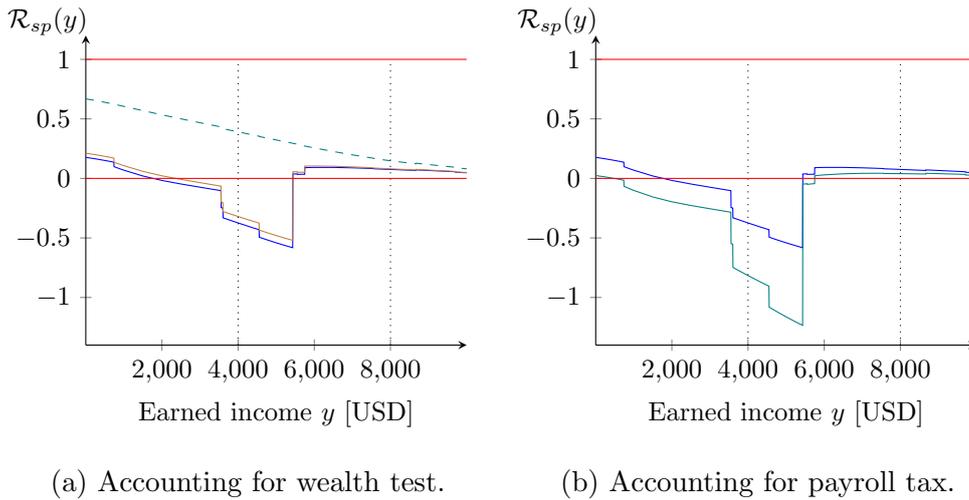


Figure S.10: Revenue function of 1974 US tax-transfer system for single parents, accounting for wealth test or payroll taxes.

Notes: Figure S.10a shows the revenue function for single parents $y \mapsto \mathcal{R}_{sp}(y)$ taking into account the wealth test for AFDC and SNAP eligibility (brown line) versus ignoring it as in the benchmark (blue line). The dashed teal line shows $y \mapsto \mathcal{R}_{sp}(y)$ ignoring AFDC and SNAP, i.e., based on the statutory income tax schedule alone. Figure S.10b plots $y \mapsto \mathcal{R}_{sp}(y)$ taking into account the employee share of social security contributions (teal line) versus ignoring them as in the benchmark (blue line). The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

ternatively, one can ask whether the tax-transfer system would have been Pareto-improved by a tax reform without an assets test, conditional on the assets test for welfare recipients being in place. To answer this question, the brown line in Figure S.10a shows a version of the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ under the assumption that 7% of single parents were not eligible for AFDC and SNAP, while all other single parents were eligible. For comparison, the blue line shows $y \mapsto \mathcal{R}_{sp}(y)$ for the benchmark case where all single parents are eligible for welfare, and the dashed teal line shows $y \mapsto \mathcal{R}_{sp}(y)$ for the counterfactual case where no single parent was eligible for welfare. As Figure S.10a shows, taking into account the wealth test leaves our results qualitatively unchanged: There was scope for Pareto improvements by means of both one-bracket and two-bracket reforms.

Second, we take into account social security contributions (payroll taxes) to compute the effective tax rates for single parents. Specifically, Figure S.10b illustrates the effect of accounting for the employee share of payroll taxes, which was 5.85% at the time. The teal line gives the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for this case, while the blue line shows the benchmark case without payroll taxes (i.e., based in the statutory income tax and welfare transfers only). With payroll taxes, we find even larger inefficiencies in the 1974 US tax-transfer system for single parents, and two-bracket tax cuts for low-income earners become even more attractive. Taking into account both the employee share and the employer share of payroll taxes would reinforce this effect.

Third, we consider a smoothed version of the 1974 tax-transfer schedule for single parents that eliminates kinks. Such smoothed tax schedules are commonly used in the literature on implicit marginal welfare weights. Specifically, Figure S.11a compares the statutory marginal tax $T'_0(y)$ (blue line) with a smooth alternative that is given by the average marginal tax in intervals of length 500 USD, $[T_0(y + 250) - T_0(y - 250)]/500$ (brown line). In Figure S.11b, the brown line shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ based on the smoothed schedule and the corresponding participation taxes, while the blue line is based on the statutory (benchmark) tax schedule. Again, the differences between both curves are small.

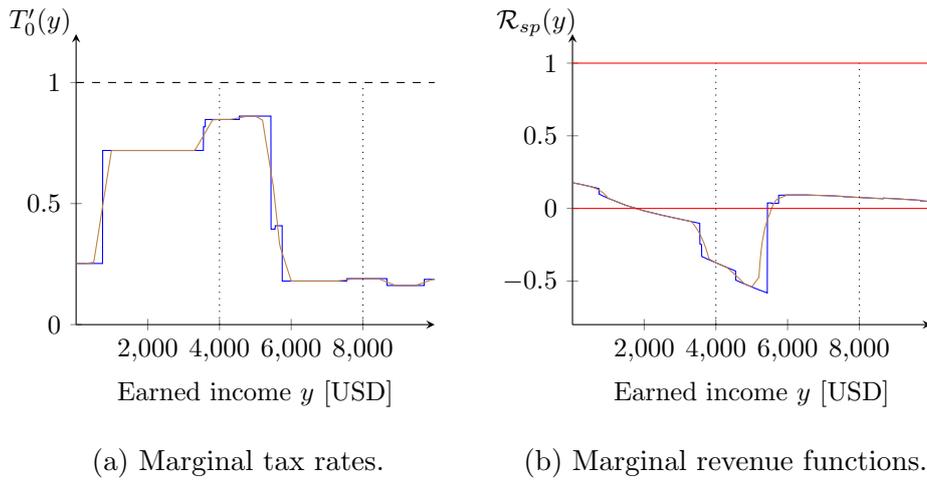


Figure S.11: Revenue function of 1974 US tax-transfer system for single parents, using smoothed tax schedule.

Notes: Figure S.11a shows a smoothed version of the marginal tax rates for single parents (teal line) versus the statutory marginal tax rates (blue line). Figure S.11b shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ based on the smoothed tax schedule (teal line) versus the statutory tax schedule as in the benchmark (blue line). The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

Tax-transfer systems in other US States. Our benchmark scenario uses 1974 marginal tax rates and participation tax rates for single parents based on the federal income tax, SNAP, and AFDC rules from California, the largest state in 1974 as today (see Table S.4). As we pointed out in the main text, AFDC rules varied to some extent across US states. Specifically, the states set different values for two parameters, the payment (need) standards and the maximum benefits, which jointly determined the income thresholds at which AFDC transfers were started to be phased out and at which they were fully phased out. In the following, we report the results of an analysis based on the four largest states next to California: New York, Texas, Pennsylvania, and Illinois. Table S.4 shows descriptive statistics for single parents in these states, based on CPS data. Despite some differences, all these states involve large shares of single parents with earned

incomes below the EITC eligibility threshold of 8,000 USD.

Table S.4: Descriptive statistics for single parents, different US states, 1974.

| Area | Pop. size [M] | Sample size | Avg. age [yrs] | Share zero income [%] | Share EITC range [%] | Earned income [USD] |
|-------------------|---------------|-------------|----------------|-----------------------|----------------------|---------------------|
| USA | 211.4 | 1,494 | 35.7 | 30.9 | 51.3 | 3,968 |
| Largest 10 states | 115.9 | 857 | 35.6 | 35.0 | 46.8 | 3,900 |
| California | 21.2 | 187 | 35.7 | 25.8 | 52.2 | 4,341 |
| New York | 18.0 | 141 | 35.9 | 51.4 | 37.0 | 2,995 |
| Texas | 12.3 | 87 | 36.1 | 27.0 | 48.0 | 3,913 |
| Pennsylvania | 11.9 | 65 | 35.9 | 42.1 | 50.2 | 2,796 |
| Illinois | 11.3 | 65 | 33.4 | 39.7 | 44.1 | 3,734 |

Notes: Table S.4 shows descriptive statistics for single parents in different US states, based on CPS 1975 data. Sample restrictions: Age 25-60, non-married, no adult family member in household, at least one child qualifying for AFDC.

Source: Authors' calculations (see Table S.2 for details).

In a first step, Figure S.12 compares the tax-transfer system across US states. Specifically, the blue lines depict the effective marginal tax rates for single parents in New York (Panel a), Texas (b), Pennsylvania (c) and Illinois (d), while the gray lines in all panels represent the tax rates in California. The main differences are in the income thresholds where AFDC is fully phased out and where, correspondingly, the effective marginal tax rates drop substantially. For New York and Pennsylvania, this drop in tax rates appears at somewhat higher income levels than for California. It appears at a somewhat smaller income level for Illinois (closer to the endpoint of the phase-in range at 4,000 USD), and at a much lower income level for Texas (around 3,150 USD). The latter difference is driven by much less generous maximum transfers in Texas.

In the second step, Figure S.13 provides graphical tests of the Pareto efficiency of the 1974 tax-transfer systems in New York (Panel a), Texas (b), Pennsylvania (c), and Illinois (d). In each panel, the blue line depicts the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents of the corresponding state, while the gray line depicts the revenue function for California, our benchmark state. The patterns for New York,

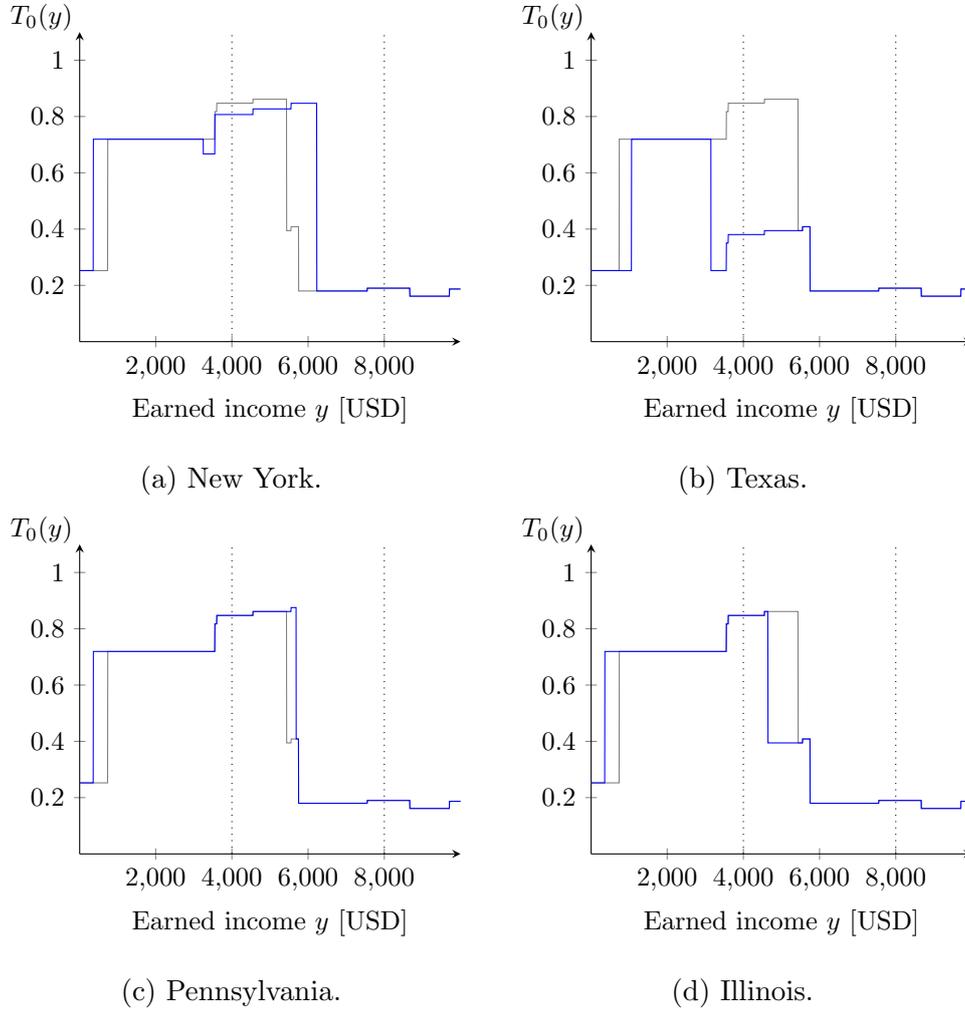


Figure S.12: 1974 tax-transfer systems for single parents, different US states.

Notes: Figure S.12 shows the effective marginal tax $y \mapsto T'_0(y)$ for single parents in different US states. Each panel compares the effective marginal tax for the indicated state (blue line) with the marginal tax for California (gray line). The vertical dotted lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1974 EITC.

Source: Authors' calculations (see Table S.2 for details).

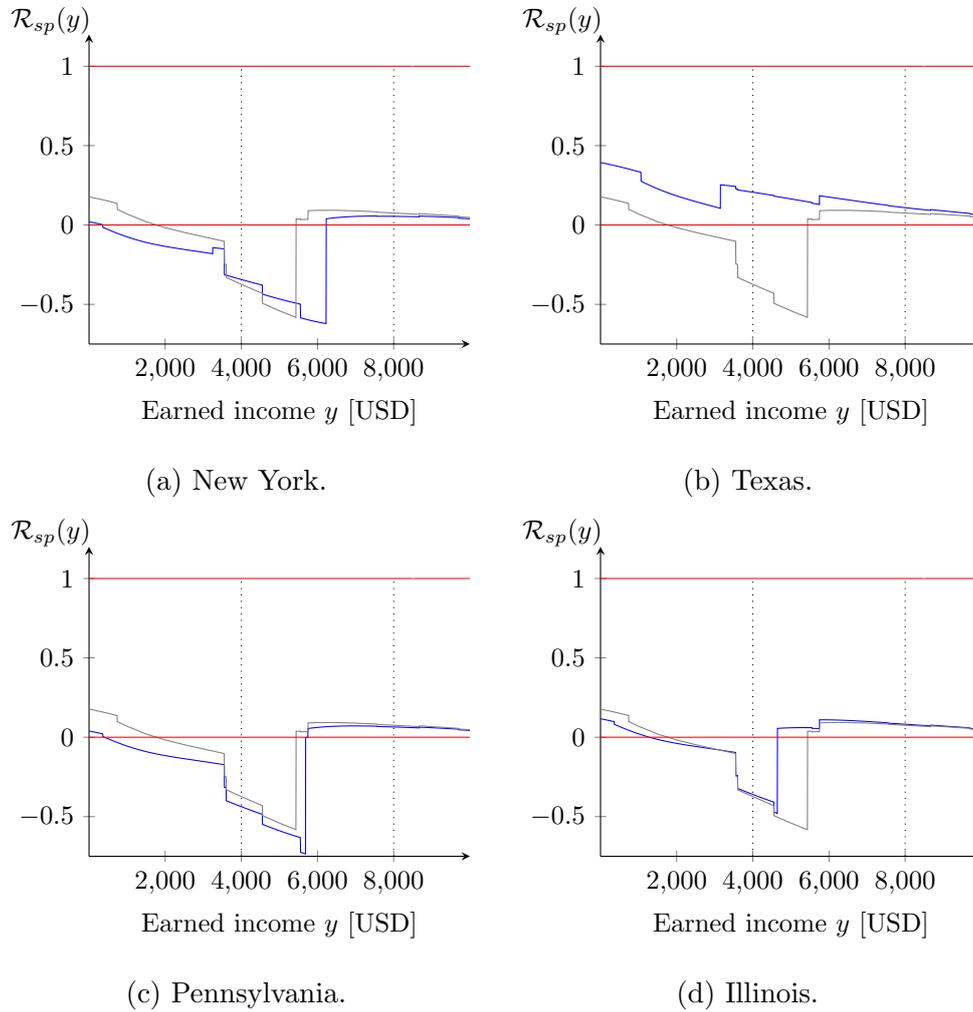


Figure S.13: Revenue functions of 1974 tax-transfer systems for single parents, different US states.

Notes: Figure S.13 shows the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in our benchmark calibration, but based on the 1974 tax-transfer systems in different US states. Each panel compares the revenue function for the indicated state (blue line) with the revenue function for California (gray line). The vertical dotted lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1974 EITC.

Source: Authors' calculations (see Table S.2 for details).

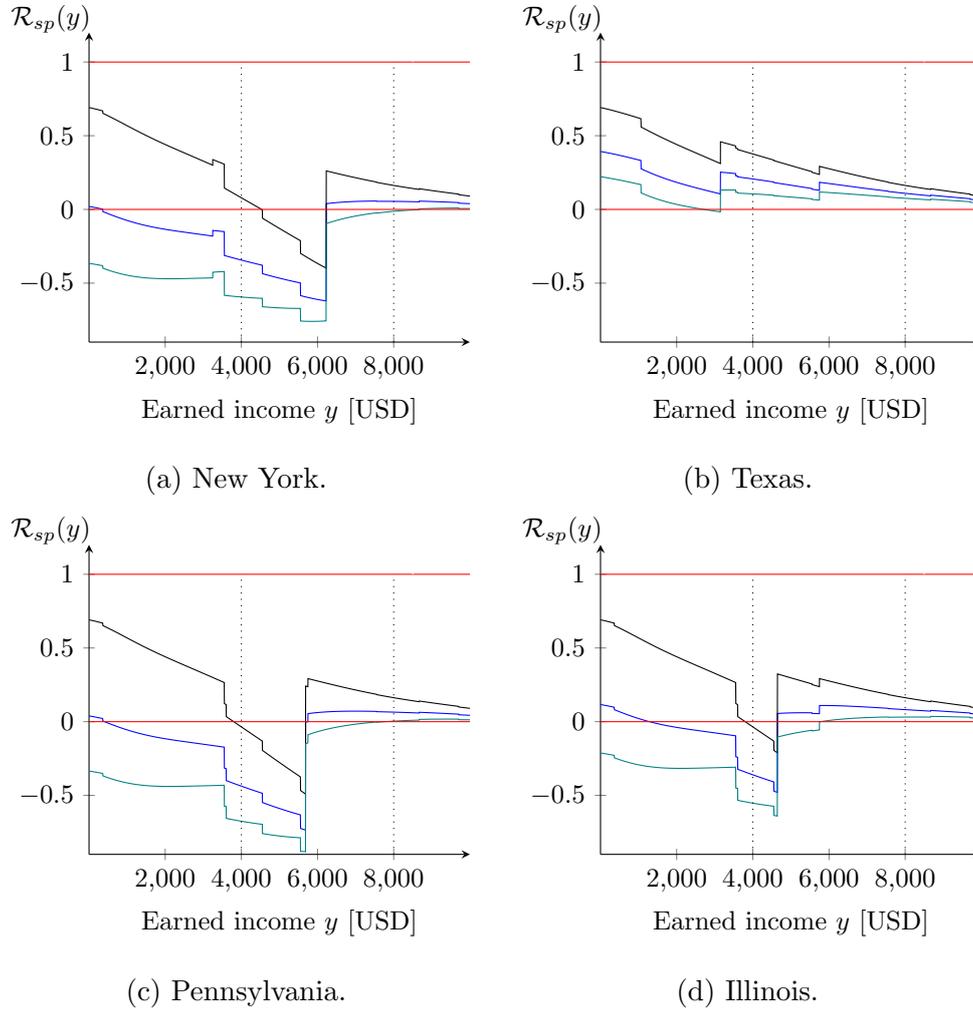


Figure S.14: Revenue functions for single parents, different US states and participation elasticities.

Notes: Figure S.14 shows the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in different US states in 1974, assuming an average participation elasticity of 0.58 (blue line, benchmark) and 0.9 (teal line) and a case without extensive-margin responses (black line). The vertical dotted lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1974 EITC.

Source: Authors' calculations (see Table S.2 for details).

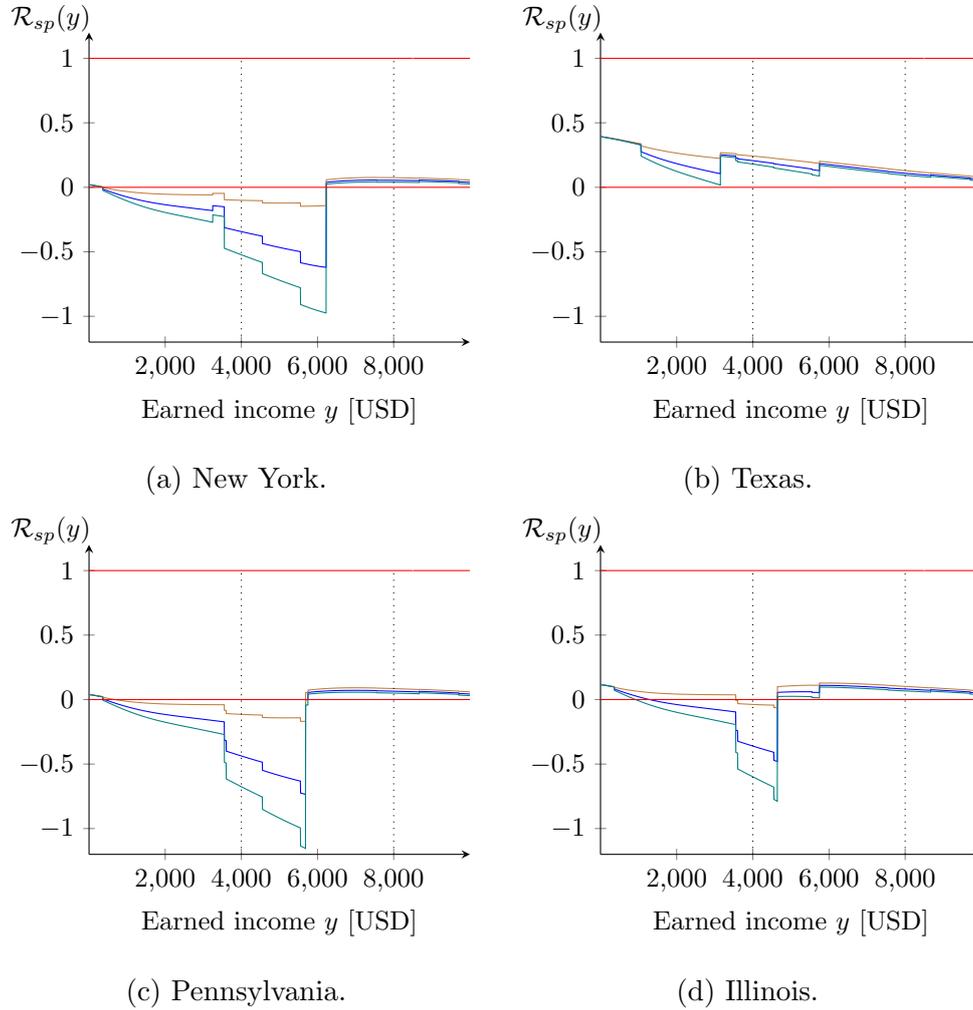


Figure S.15: Revenue functions for single parents, different US states and intensive-margin elasticities.

Notes: Figure S.15 shows the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in different US states in 1974, assuming an intensive-margin elasticity of 0.1 (brown line), 0.33 (blue line, benchmark) and 0.5 (teal line). The vertical dotted lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1974 EITC.

Source: Authors' calculations (see Table S.2 for details).

Pennsylvania, and Illinois are very similar to those for California: The revenue functions violate the lower bound at zero and the monotonicity condition. Hence, there exist Pareto-improving reforms with one bracket and with two brackets. For Texas, the revenue function is also non-monotonic, but below 1 and above 0 at each income level. Hence, a Pareto improvement can only be realized by reforms involving two brackets. These qualitative results are not affected by alternative assumptions on the labor supply elasticities at both margins, see Figures S.14 and S.15.

Table S.5: Revenue effects of 1975 EITC reform and optimal tax reforms, different US states.

| State | 1975 EITC | Optimal tax reform | | | |
|-----------------------|----------------------------|--------------------------|-------------------------|-----------------------|------------------------|
| | Revenue effect [USD pc] | Revenue gain [USD pc] | Phase-in start [USD] | Phase-in end [USD] | Phase-out end [USD] |
| California | -4.03 | 12.64 | 1,248 | 5,748 | 10,248 |
| New York | -4.61 | 18.59 | 0 | 6,225 | 12,450 |
| Texas ^a | -3.33 | 1.12 | 1,900 | 3,148 | 4,396 |
| Pennsylvania | -2.91 | 17.66 | 200 | 5,748 | 11,296 |
| Illinois ^b | 2.26 | 8.58 | 573 | 4,925 | 9,673 |

Notes: Table S.5 considers small reforms for different US states, given the 1974 tax-transfer systems. Column 2 shows the (per-household) revenue effects of a small reform in the direction of the 1975 EITC reform. Column 3 shows the state-specific (per-household) maximum revenue gain from a small tax cut. Columns 4 and 5 reports the income thresholds at which the phase-in range of the state-specific optimal tax reform starts and ends, respectively. Column 6 reports the income thresholds at which the phase-out range of the state-specific optimal tax reform ends. *a*: The optimal tax reform in Texas consists of a two-bracket tax cut between 1,900 and 4,396 USD (revenue gain: 0.93 USD per capita) and another two-bracket tax cut between 5,075 and 5,748 USD (revenue gain: 0.19 USD pc). *b*: The optimal tax reform in Illinois reduces marginal taxes for incomes in [573; 4,925] and in [5,550; 5,748]; it increases marginal taxes for incomes in [4,925; 5,550] and in [5,748; 9,673].

Source: Authors' calculations (see Table S.2 for details).

In the third step, we calculate the (per-household) revenue effects of a small reform in the direction of the *historical* 1975 EITC reform in the same US states.

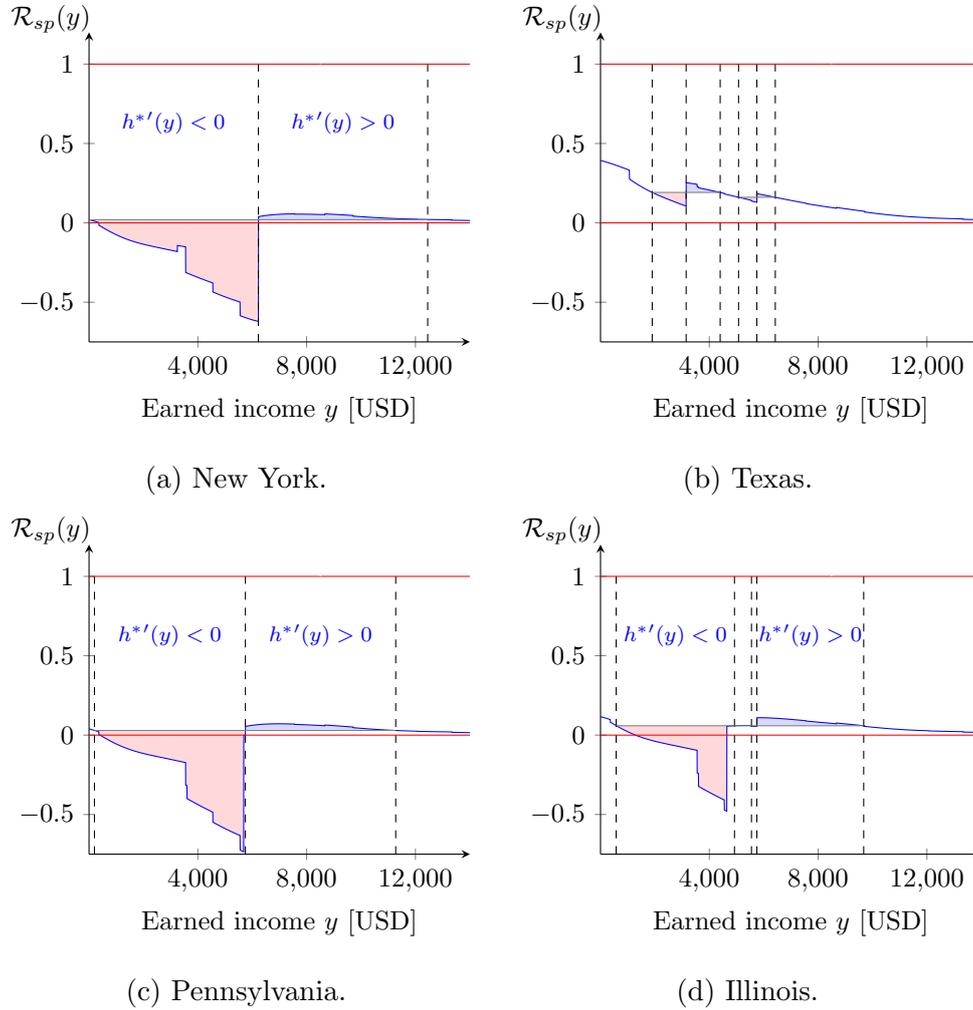


Figure S.16: Optimal tax reforms for single parents, different US states.

Notes: Figure S.16 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents, based on the effective tax rates in different US states. It also illustrates the optimal tax reforms in these states, which would have reduced marginal taxes in red-shaded areas and increased them in blue-shaded areas. The sum of all shaded areas represents the available revenue gain per single-parent family.

Source: Authors' calculations (see Table S.2 for details).

As show in the second column of Table S.5, the reform created revenue losses in California, New York, Taxes, and Pennsylvania. By contrast, it would have been self-financing if restricted to single parents in Pennsylvania only. The remaining columns of Table S.5 provide information about the optimal tax reforms in all states (i.e., the best alternative EITCs). Specifically, column 3 shows that these state-specific reforms would have generated even larger free lunches in New York and in Pennsylvania than in California, and only slightly smaller revenue gains in Illinois. The optimal tax reform in Texas would have created revenue gains of a smaller size. As columns 4 to 6 show, the location of the optimal phase-in and phase-out ranges varies across US states. Figure S.16 illustrates these results by depicting the optimal tax reforms for single parents in New York, Texas, Pennsylvania, and Illinois.

One take-away from this analysis is that, corresponding to the differences in transfer programs, there was also room for the introduction of state-specific EITC schedules. Indeed, since the 1980s, 28 states – including New York and Illinois, but neither Texas nor Pennsylvania – and the District of Columbia have introduced state EITCs, and three more states will introduce them in 2023. It is worth noting, however, than most state EITCs are calculate as a fixed percentage of the federal EITC and do not vary the location of the phase-in and phase-out ranges.

D.3 Heterogeneity with respect to household size

In this section, we take a closer look at the subgroups of single parents with different numbers of children. In 1975, a single EITC schedule was introduced for all taxpayers with children, irrespective of the household size. Decades later, the US authorities introduced additional, more generous EITC schedules for single parents with two or more children (in 1991), and for single parents with three or more children (in 2009). Our analysis provides two insights. First, the 1974 US tax-transfer system was not Pareto-efficient, irrespective of what subgroup of single parents we consider. Second, the size of the inefficiencies differed across these subgroups: Using the inefficiency measure introduced in Section 2.4, the 1974

tax-transfer system was more inefficient for single parents with three children than for those with two children, and the least inefficient for single parents with one child. Correspondingly, the optimal tax reforms differed across subgroups, and the introduction of differentiated EITC schedules would have made sense already in the mid-1970s.

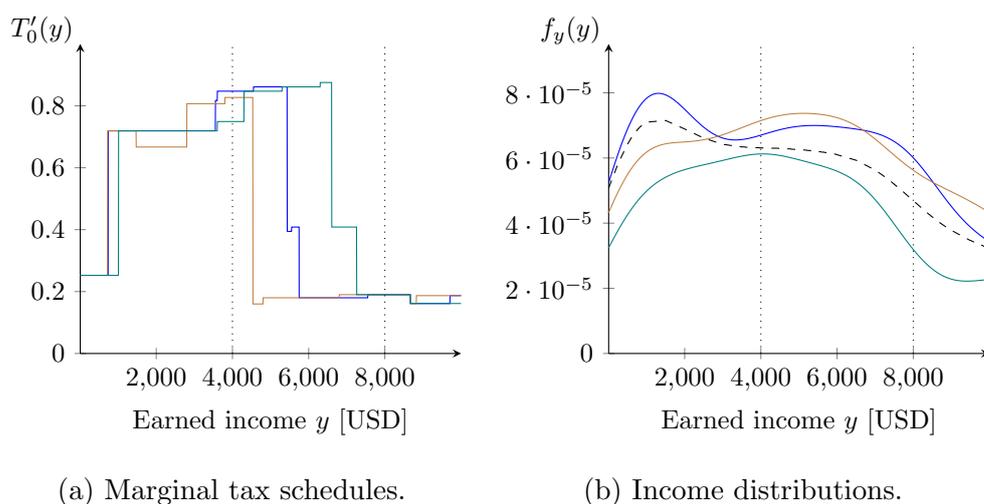


Figure S.17: 1974 US tax-transfer schedules and income distributions for single parents with 1, 2, and 3 children.

Notes: Figure S.17 shows the tax-transfer schedules and the income distributions among single parents with one child (brown lines), two children (blue lines), and three children (teal lines). Specifically, the left panel shows the 1974 effective marginal tax rates $y \mapsto T'_0(y)$. The right panel shows the kernel estimates of the 1974 pdf $y \mapsto f_y(y)$ of earned incomes; the dashed black line plots the estimated pdf among all single parents. In both panels, the dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

Figure S.17 demonstrates pronounced heterogeneity among single parents in two dimensions. First, Figure S.17a shows that the 1974 US tax-transfer schedules differed substantially depending on the number of children: For single parents with more children, effective marginal taxes were higher and remained high up to higher income thresholds. Specifically, marginal taxes above 70% applied for single parents with one child up to annual incomes around 4,500 USD; for those

with two children up to annual incomes around 5,400 USD; for those with three children up to annual incomes around 6,600 USD. The main reason for these differences is that the maximum amount of welfare transfers (AFDC, SNAP) was increasing in family size. As the transfers were phased out at the same rates, the full phasing-out of welfare transfers occurred at higher incomes for families with more children. Second, Figure S.17b shows that there were differences in the distributions of earned income across single parents: Parents with more children had lower average incomes. Specifically, it shows the kernel estimates of the pdfs of earned income based on CPS March 1975 data, using a Gaussian kernel with a bandwidth of 1200 for single parents with one and two children and a bandwidth of 1400 for single parents with three children. For comparison, the dashed black line shows the estimated pdf of the income distribution among all single parents. Among single parents with positive incomes, the average incomes were 5,037 USD for those with one child; 4,165 USD for those with two children; and 3,550 USD for those with three children. The shares of single parents without any earned income were 20.9% (one child); 24.9% (two children); and 43.5% (three children).

To perform separate tests for Pareto efficiency of the 1974 US tax system in these subgroups, Figure S.18 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents with one child in the upper left panel, for those with two children in the upper right panel, and for those with three children in the lower panel. For all cases, we assume an intensive-margin elasticity of 0.33 and participation elasticities that decrease with income from 0.67 to 0.4 as in the benchmark calibration. First, Figure S.18 clarifies that the 1974 US tax-transfer schedules were not Pareto-efficient for any subgroup of single parents. All panels show that there were Pareto-improving tax reforms with both one and two brackets. Second, the relevant income ranges differed: With more children, the inefficiencies arise at higher income levels and cover a wider range of incomes. For single parents with one child, in particular, the inefficiencies are concentrated on income ranges between 2,000 and 4,500 USD, close to the endpoint of the 1975 EITC phase-in range. Third and relatedly, by evaluating Equation (12), we find that the tax reform introducing the 1975 EITC

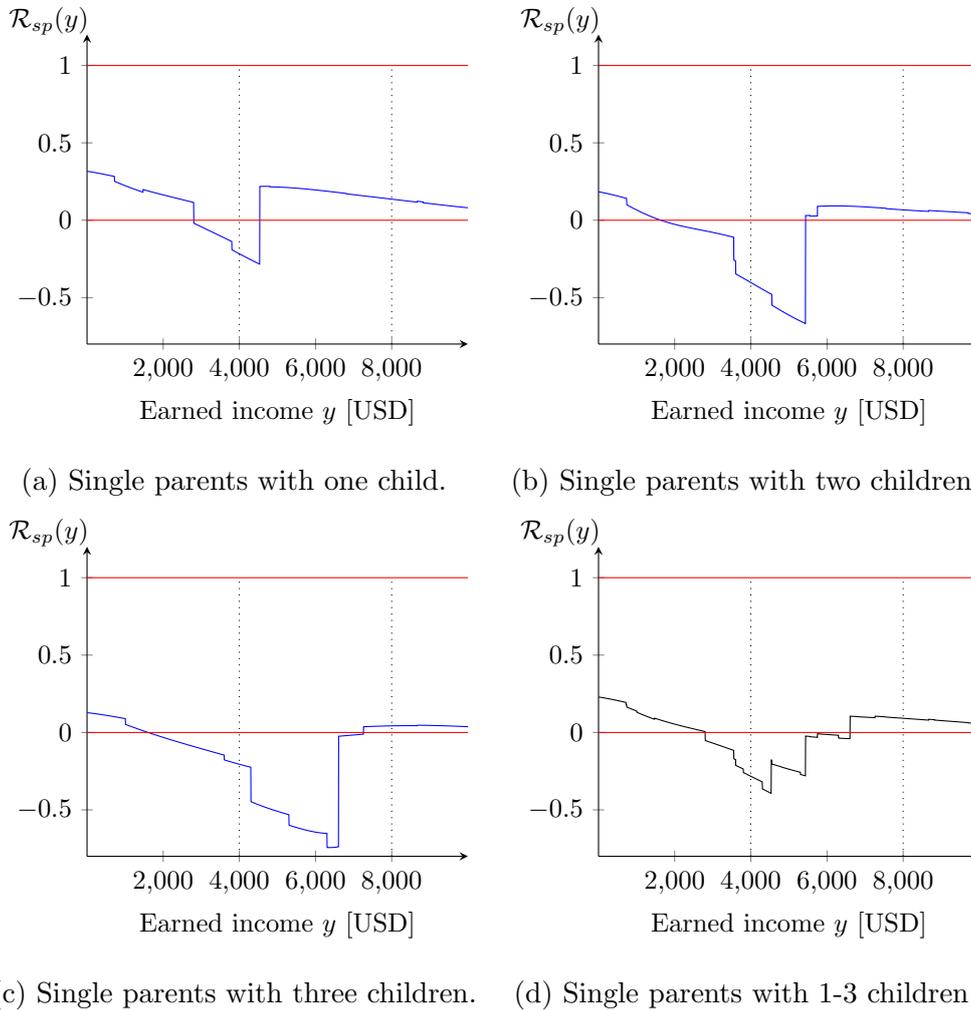


Figure S.18: Revenue functions of 1974 US tax-transfer system, single parents with 1, 2, and 3 children.

Notes: Figure S.18 shows the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ in 1974 for single parents with one child (upper left panel), single parents with two children (upper right panel), single parents with three children (lower left panel), and for all single parents with one to three children (lower right panel). In all panels, the dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

was self-financing and Pareto-improving for single parents with one child, while it was not for those with two or three children, see first column of Table S.6.

The inefficiency measure developed in Section 2.4 confirms the visual impression from Figure S.18. As reported in Table S.6, for single parents with one child, the optimal tax reform was a two-bracket tax cut that reduced marginal taxes in the income range 2,003 to 4,533 USD and increased them in the income range 4,533 to 7,063 USD, accompanied by two further two-bracket tax cuts in the income ranges 1,360 to 1,590 USD and 8,566 to 8,778 USD. This reform allowed to raise tax revenue per household by 6.38 USD per-percentage point change in the marginal tax rates, with more than 99% of that amount coming from the income range 2,003 to 7,063 USD.

For single parents with two children, the optimal tax reform was a two-bracket tax cut that reduced marginal taxes in the income range 1,266 to 5,748 USD and increased marginal taxes in the income range 5,748 to 10,230 USD. This reform allowed to raise the tax revenue per capita by 13.51 USD for each percentage-point change in marginal taxes.

Finally, the optimal tax reform for single parents with three children was a two-bracket tax cut that reduced marginal taxes in the income range 1,305 to 7,257 USD and increased them in the income range 7,257 to 13,219 USD. It allowed to raise revenue per household by 18.51 USD per percentage-point change in marginal taxes.

We note that the optimal tax reforms for all three subgroups of single parents involve a reduction in marginal taxes in the range between approximately 2,000 and 4,500 USD, but they differ substantially above this interval. In particular, at all incomes between 4,533 and 7,063 USD, they require to increase marginal tax rates for some subgroup and to decrease marginal taxes for another one. We conclude that the introduction of a unique EITC schedule for all single parents was not ideal; separate EITC schedules depending on family size (as introduced in the 1990s and 2000s) would have allowed to tackle the observed inefficiencies more precisely.

To quantify this conclusion, we compute the revenue gains that could have been realized by either conducting the optimal tax reforms conditional on family size (reported above), or the optimal unconditional tax reform. We derive the latter based on the revenue function depicted in Figure S.18d, which captures revenue effects of reforms that are adopted identically across all three subgroups. This joint revenue function takes into account the shares of single parents with one child (42.9%), two children (36.8%), and three children (20.3%), the differences in tax-transfer rates they face and the heterogeneity of their income distributions. We find that the best unconditional tax cut would have raised tax revenue by 9.16 USD per percentage point change for each single parent. By contrast, the best conditional tax cuts in all three subgroups would have yielded a total revenue gain that was about 25% larger, i.e., 11.47 USD per percentage point change for each single parent.

Table S.6: Revenue effects of 1975 EITC reform and optimal tax reforms for single parents, different subgroups.

| No. of children | 1975 EITC | Optimal tax reform | | | |
|-----------------|-------------------------|-----------------------|----------------------|--------------------|---------------------|
| | Revenue effect [USD pc] | Revenue gain [USD pc] | Phase-in start [USD] | Phase-in end [USD] | Phase-out end [USD] |
| 1 child | 0.40 | 6.38 | 2,003 | 4,533 | 7,063 |
| 2 children | -4.73 | 13.51 | 1,266 | 5,748 | 10,230 |
| 3 children | -12.93 | 18.51 | 1,305 | 7,257 | 13,219 |
| 1-3 children | -4.20 | 9.15 | 2,225 | 6,602 | 11,000 |

Notes: Table S.6 considers small reforms for different subsets of single parents (rows 2-4) and for the set of all single parents with one to three children (row 5), given the 1974 tax-transfer systems. Column 2 shows the (per-household) revenue effects of a small reform in the direction of the 1975 EITC reform. Column 3 shows the (per-household) maximum revenue gain from a small tax cut, conditional on the number of children. Columns 4 and 5 reports the income thresholds at which the phase-in range of the subgroup-specific optimal tax reform starts and ends, respectively. Column 6 reports the income thresholds at which the phase-out range of the subgroup-specific optimal tax reform ends.

Source: Authors' calculations (see Table S.2 for details).

For completeness, we note that this heterogeneity is mainly driven by the

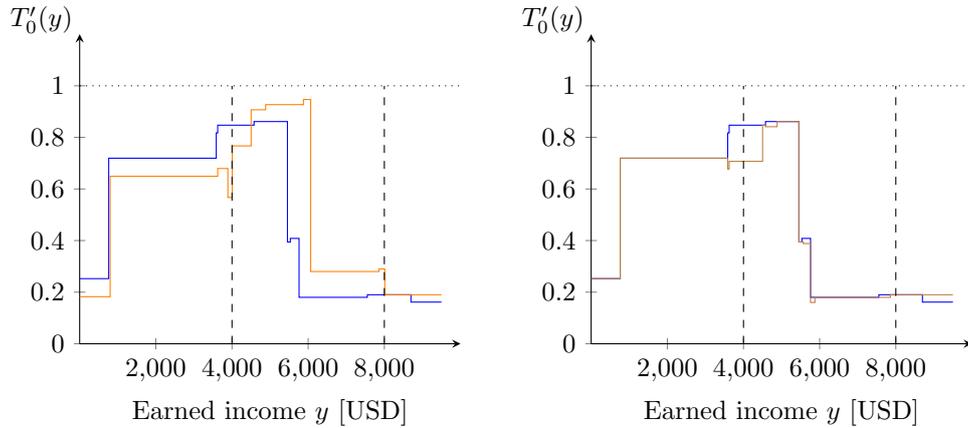
differences in the tax-transfer schedules shown in Figure S.17a. If we compute the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ using the same income distribution for all subgroups of single parents, the inefficiencies differ even more in size than if we use subgroup-specific income distributions as in Figure S.18.

Summing up, we find that the inefficiencies in the 1974 US tax-transfer system were substantially larger for single parents with more children. Moreover, the best available reforms differed across the subgroups: For lone parents with more children, EITC-like reforms on much wider and higher income ranges were Pareto-improving than for those with less children. We conclude that, already in 1974, the introduction of differentiated EITC schedules – as implemented decades later – would have been desirable.

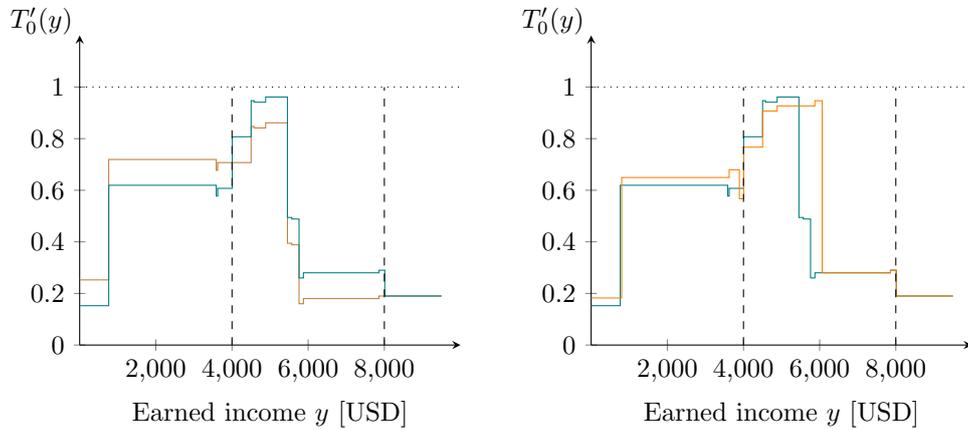
D.4 Results for further post-reform years

We now extend our analysis to the US tax-transfer systems after the 1975 EITC introduction. We start by explaining how the effective marginal tax rates for single parents with two children changed from 1974 to 1975. We also provide a graphical decomposition of the overall change due to adjustments of the statutory income tax, changes in the transfer schedules, and changes due to the introduction of the EITC. Second, we demonstrate that the 1975 tax-transfer schedule was not Pareto-efficient. We then show that these inefficiencies remained and even increased until 1978, partly due to adjustments of the tax schedule and welfare transfers that were independent of the EITC. Third, we investigate the 1979 reform that implemented the first EITC expansion. We find that this reform went into a Pareto-improving direction. In a final step, we show that these results are robust to various alternative assumptions about labor supply responses.

Figure S.19 illustrates the 1975 changes in the tax-transfer schedule for single parents. The upper left panel compares the marginal taxes for low-income earners in 1974 (blue line) and 1975 (orange line). The other three panels associate these changes to the elements of the tax-transfer system: the statutory income tax, the EITC, and the welfare transfers AFDC and SNAP. In the statutory income



(a) Marginal taxes 1974 versus 1975. (b) Changes in statutory income tax.



(c) Changes due to EITC. (d) Changes in AFDC and SNAP.

Figure S.19: Changes in effective marginal taxes for single parents, 1974 to 1975.

Notes: Panel (a) of Figure S.19 compares the effective marginal taxes for single parents with two children in 1974 (blue line) and 1975 (orange line). The remaining panels decompose these differences into three elements of the tax-transfer systems. Specifically, Panel (b) shows how the marginal tax in 1974 was changed by adjustments of the statutory income tax (brown line vs. blue line); Panel (c) shows the effect of the EITC introduction (teal line vs. brown line); Panel (d) accounts for changes in AFDC and SNAP parameters (orange line vs. teal line). The dotted vertical lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC.

Source: Authors' calculations (see Table S.2 for details).

tax, there were minor adjustments in the standard deduction and a disregard for personal exemptions. This implied a reduction of marginal taxes in the lowest tax bracket (see upper right panel). The introduction of the EITC reduced the marginal tax by 10 percentage points between 0 and 4,000 USD, and increased the marginal tax by 10 percentage points between 4,000 and 8,000 USD (see lower left panel). Finally, adjustments in the AFDC and SNAP parameters implied increases in the maximum transfer amounts and in the eligibility thresholds. As a result, both transfers were phased out over wider income ranges, and marginal tax rates increased substantially between 5,400 and 6,100 USD (see lower right panel.) We note that the EITC introduction implied the most substantial tax changes for single parents, but the AFDC and SNAP adjustments also played a role.

In Section 3 in the main text, we study whether the 1975 EITC introduction went into a Pareto-improving direction, holding all other elements of the tax-transfer system fixed. This is the reform we seek to evaluate, and therefore it would not be appropriate to evaluate several changes in the tax-transfer system simultaneously. We find that this reform was not Pareto-improving according to our benchmark calibration, while a similar two-bracket tax cut with different phase-in and phase-out ranges would have been. In particular, the 1975 EITC further increased the excessively high marginal taxes between 4,000 and 5,750 USD, while a reduction would have been Pareto-improving. As can be seen from Figure S.19a, the same pattern emerges if we take account of all changes in the tax-transfer schedule between 1974 and 1975. This suggests that, due to the suboptimal design of the EITC, the 1975 US tax schedule neither was Pareto-efficient.

We confirm this hypothesis applying the conditions for Pareto efficiency in Theorem 1 to the tax-transfer system for single parents prevailing after the introduction of the EITC. Figure S.20 provides a graphical illustration of this test for 1975, right after the EITC introduction, and for 1978, before the first EITC expansion. It shows that the US tax-transfer system remained inefficient through-

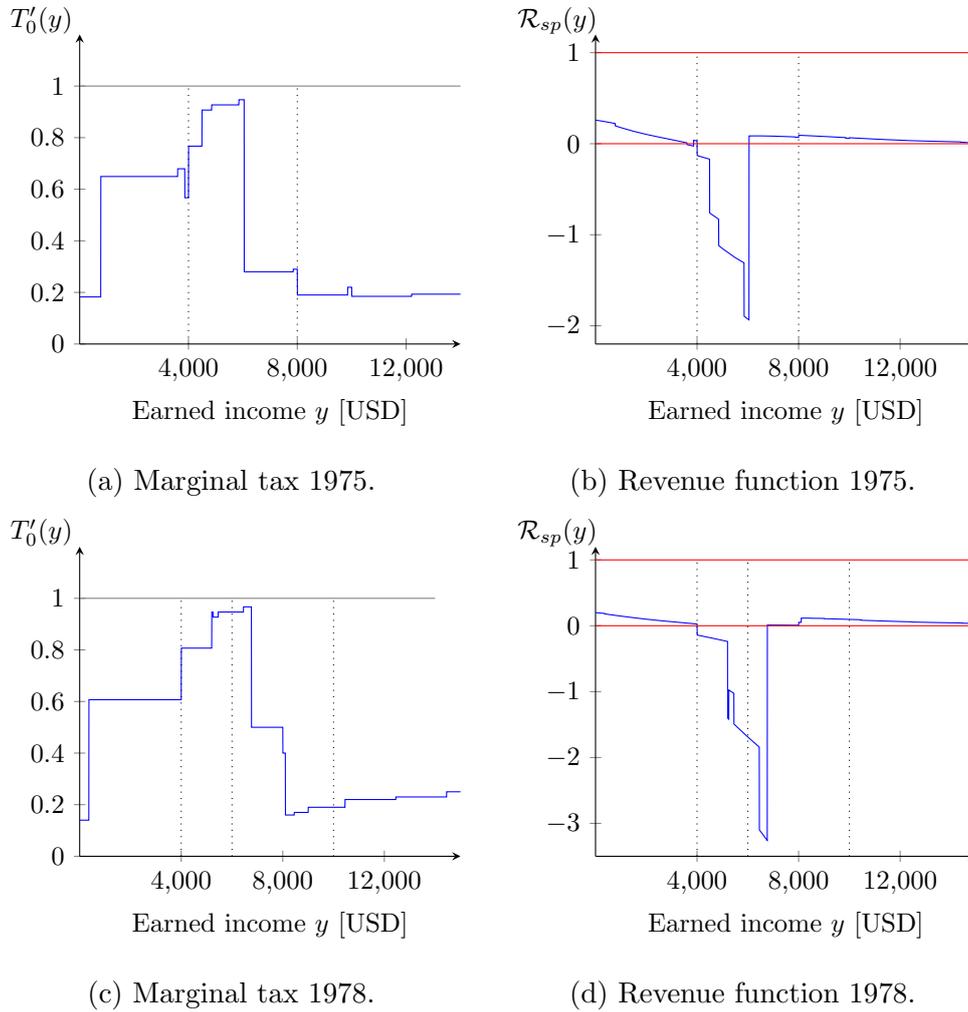


Figure S.20: Revenue functions of US tax-transfer system for single parents, 1975 and 1978.

Notes: Figure S.20 shows the marginal tax rates (left panels) and the revenue functions $y \mapsto \mathcal{R}_{sp}(y)$ (right panels) for single parents in the years 1975 and 1978. In the upper panels, the vertical dotted lines mark the endpoints of the phase-in range at 4,000 USD and the phase-out range at 8,000 USD of the 1975 EITC. In the lower panels, the dashed vertical lines mark the income ranges in which marginal taxes were reduced (4,000 to 6,000 USD) and increased (6,000 to 10,000 USD) by the 1979 EITC reform, respectively.

Source: Authors' calculations (see Table S.2 for details).

out this time. More specifically, the upper right panel shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for the 1975 tax system. On the one hand, it shows that the conditions for Pareto efficiency were now violated on a smaller income range than in 1974. In 1975, the optimal tax reform was a two-bracket reform that reduced marginal taxes between 2,525 and 6,045 USD, and increased them between 6,045 and 9,565 USD (i.e., on a smaller interval than in 1974). On the other hand, function $y \mapsto \mathcal{R}_{sp}(y)$ attains even more negative values for some incomes. As a result, we find that the post-EITC tax system was even more inefficient than the pre-EITC tax system: In 1975, a two-bracket tax cut allowed to raise tax revenue by 21.9 USD per percentage point change for each single parent household. In 1974, only a revenue gain of 12.6 USD per percentage point change was feasible.

From 1975 to 1978, the income range with inefficiencies widened, eventually covering the majority of the phase-out range of the EITC between 4,000 and 8,000 USD (see Figure S.20d). This increase was driven partly by secular changes such as wage growth, and partly by a series of minor adjustments in deductions, income disregards and AFDC/SNAP parameters. Finally, the tax-transfer system in 1978 involved larger inefficiencies than both in 1974 and 1975. It was possible to realize revenue gains per single parent of up to 35.4 USD by a two-bracket tax cut that changed marginal taxes by only one percentage point. Specifically, this optimal tax reform required to reduce marginal taxes between 3,256 and 8,100 USD and to increase them between 8,100 and 12,944 USD.

Interestingly, the 1979 EITC reform was fairly similar to this optimal tax reform. The US government strongly expanded the EITC in the tax year 1979, using a two-bracket tax cut with the following properties: The phase-in range was expanded from 4,000 to 5,000 USD; a plateau range between 5,000 and 6,000 USD was introduced and the phase-out range was adjusted and now went from 6,000 to 10,000 USD. While the phase-in rate was kept at 10%, the phase-out rate was raised from 10% to 12.5%. Altogether, the 1979 EITC reform was a tax

cut in direction

$$\tilde{h}_{79}(y) = \begin{cases} -0.2 & \text{for } y \in [4.000; 5.000) , \\ -0.1 & \text{for } y \in [5.000; 6.000) , \\ 0.025 & \text{for } y \in [6.000; 8.000) , \\ 0.125 & \text{for } y \in [8.000; 10.000) , \\ 0 & \text{for } y \notin [4.000, 10.000) . \end{cases}$$

To check whether this reform went into a Pareto-improving direction, we evaluate condition (7) using the reform direction $y \mapsto \tilde{h}_{79}(y)$ and the 1978 revenue function $y \mapsto \mathcal{R}_{sp}(y)$. Based on our benchmark calibration, we find that the 1979 EITC reform direction was indeed Pareto-improving: It raised tax revenue and allowed to increase the base transfer to all single parents, while directly reducing the tax liabilities for all taxpayers with earned incomes between 4,000 and 10,000 USD.

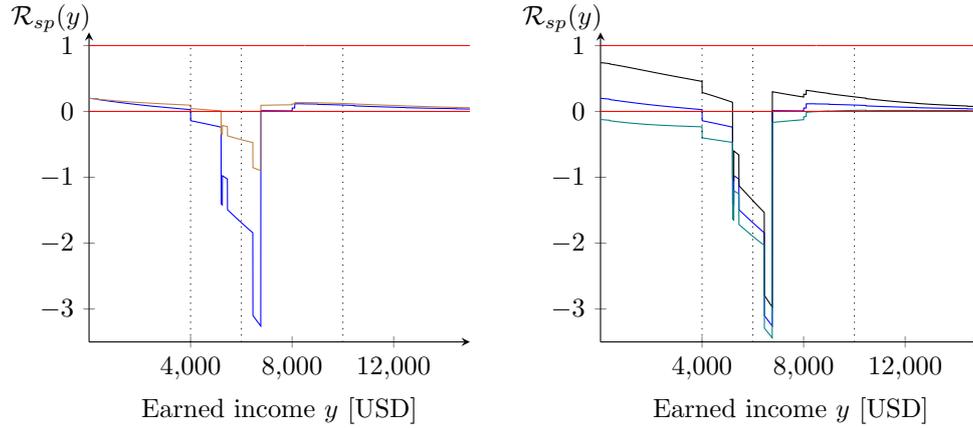
Table S.7: Revenue effects of 1979 EITC reform, different scenarios.

| Scenario | ε | $\bar{\pi}$ | MPE | Revenue gain p.c. [USD] |
|--------------------|---------------|-------------|-------|-------------------------|
| Benchmark | 0.33 | 0.58 | 0 | 1.36 |
| Low ε | 0.1 | 0.58 | 0 | 0.45 |
| High ε | 0.5 | 0.58 | 0 | 2.04 |
| Low $\bar{\pi}$ | 0.33 | 0 | 0 | 0.73 |
| High $\bar{\pi}$ | 0.33 | 0.9 | 0 | 1.74 |
| Medium MPE | 0.33 | 0.58 | -0.12 | 1.22 |
| High MPE | 0.33 | 0.58 | -0.24 | 1.10 |

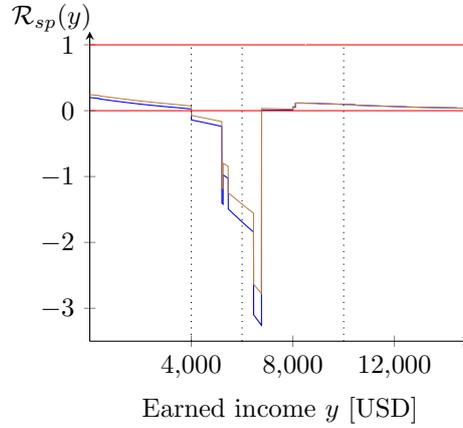
Notes: Table S.7 reports the revenue effects of a small reform in the direction of the 1979 EITC reform that reduces tax liabilities for earned incomes between 4,000 and 10,000 USD for different combinations of the intensive-margin elasticity ε , the average participation elasticity $\bar{\pi}$, and the marginal propensity to earn (MPE) out of unearned income, a measure for the strength of income effects.

Source: Authors' calculations (see Table S.2 for details).

As a final step, we show that our results on the 1978 tax-transfer schedule and the 1979 EITC expansion are remarkably robust with respect to the assumptions on labor supply responses. Figure S.21 depicts the revenue function for the 1978



(a) Different intensive elasticities. (b) Different participation elasticities.



(c) Income effects.

Figure S.21: Revenue function of 1978 US tax-transfer system for single parents, robustness checks.

Notes: Figure S.21 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ (right panels) for single parents in 1978, given different scenarios. Panel (a) compares cases with an intensive-margin elasticity of 0.1 (brown line) and 0.33 (blue line, benchmark). Panel (b) considers cases with a participation elasticity of 0 (black line), 0.58 (blue line, benchmark) and 0.9 (teal line). Panel (c) considers cases with income effects at both margins corresponding to an MPE of -0.24 (brown line) and without income effects (blue line, benchmark). The 1979 EITC reform reduced marginal taxes for incomes between 4,000 and 6,000 USD (first and second dotted lines) and raised them for incomes between 6,000 and 10,000 USD (second and third dotted lines).

Source: Authors' calculations (see Table S.2 for details).

tax-transfer schedule under alternative assumptions on the intensive-margin elasticities (upper left panel), the extensive-margin elasticities (upper right panel), and the magnitude of income effects (lower panel). Across all scenarios, we find that there existed Pareto-improving reforms with one and with two brackets. Table S.7 reports the per capita revenue gains of a small reform in the direction of the 1979 EITC reform for the same scenarios. Again, our finding that the 1979 reform was a self-financing tax cut and, therefore, Pareto-improving holds across all cases.

To sum up, we have demonstrated that the 1974 EITC reform did neither go in a Pareto-improving direction nor did it restore a Pareto-efficient tax system; mainly because the 1974 EITC was too limited in the income range that it covered. By contrast, the 1979 reform expanded the phase-in of the EITC to an income of 6,000 USD, and its phase-out to incomes of up to 10,000 USD, thereby correcting the initial weaknesses and achieving a Pareto improvement.

E Sufficient statistics with behavioral responses at the intensive and the extensive margin

We consider a setting with two-dimensional heterogeneity. Individuals differ both in fixed and variable costs associated with the generation of income. Such a framework has been suggested by the literature that analyzes earnings subsidies from an optimal-tax perspective (see, e.g., Saez (2002), Jacquet et al. (2013), or Hansen (2021)). The analysis in this part of the Supplementary Material yields a characterization of $y \mapsto R_{\tau\ell}^s(0, 0, y)$ that depends on labor supply elasticities at the intensive and the extensive margin.

Henceforth, fixed costs are captured by a parameter γ , variable costs by a parameter ω . Thus, we write $\theta = (\omega, \gamma)$ for an individual's type. Variable cost types and fixed cost types belong, respectively, to subsets of the positive reals that we denote by $\Omega = [\underline{\omega}, \bar{\omega}]$ and $\Gamma = [\underline{\gamma}, \bar{\gamma}]$. The joint distribution is denoted by F . The utility that an individual with type (ω, γ) derives from a (c, y) -pair

that involves positive earnings is denoted by $u(c, y, \omega, \gamma)$. We denote by $u^{no}(c_0)$ the utility of individuals with no earnings. The function u^{no} is assumed to be increasing and weakly concave.

Variable costs. To capture variable costs, we assume that preferences satisfy the Spence-Mirrlees single-crossing property: Consider two individuals with the same fixed cost type γ , and an arbitrary point in the (c, y) -space with $y > 0$. At any such point, an individual with a higher ω -type has a flatter indifference curve. The interpretation is that, due to her lower variable costs, she needs less compensation for a marginal increase of her earnings. Formally, for any given γ , $\omega' > \omega$ implies that

$$-\frac{u_y(c, y, \omega', \gamma)}{u_c(c, y, \omega', \gamma)} < -\frac{u_y(c, y, \omega, \gamma)}{u_c(c, y, \omega, \gamma)}.$$

for any pair (c, y) with $y > 0$.

Let $C : y \mapsto C(y)$ be a non-decreasing function, interpreted as the boundary of a budget set that individuals face. An implication of the Spence-Mirrlees single crossing property is as follows: Consider two individuals who differ only in the variable cost type. If type ω weakly prefers an earnings level y' over an earnings level $y < y'$, then any type $\omega' > \omega$ strictly prefers y' over y . More formally, for any γ , any pair ω', ω with $\omega' > \omega$, and any pair y', y with $y' > y$,

$$u(C(y'), y', \omega, \gamma) \geq u(C(y), y, \omega, \gamma) \quad \text{implies} \quad u(C(y'), y', \omega', \gamma) > u(C(y), y, \omega', \gamma).$$

Fixed costs. Fixed costs affect the compensation that individuals demand for positive earnings. Let $\pi(c, y, \omega, \gamma)$ be such that

$$u(c + \pi(c, y, \omega, \gamma), y, \omega, \gamma) = u^{no}(c_0).$$

We assume that π is a strictly increasing function of γ .⁴⁶

⁴⁶This property holds for various preference specifications that have been explored in the literature. In particular, it holds for separable utility functions of the form $u(c, y, \omega, \gamma) = \tilde{u}(c, y, \omega) - \gamma \mathbf{1}_{y>0}$, where $\mathbf{1}$ is the indicator function. It also holds for specifications with monetary fixed costs $u(c, y, \omega, \gamma) = \tilde{u}(c - \gamma \mathbf{1}_{y>0}, y, \omega)$. The two classes coincide if the function \tilde{u} is quasi-linear in c .

An implication is as follows: Consider two individuals who differ only in their fixed-cost type γ . Given a non-decreasing consumption schedule $C : y \mapsto C(y)$, if type γ weakly prefers an earnings level of 0 over an earnings level of $y > 0$, then any individual with a fixed-cost type $\gamma' > \gamma$ strictly prefers 0 over y ; for any ω , any pair γ', γ with $\gamma' > \gamma$, and any $y > 0$,

$$u^{no}(c_0) \geq u(C(y), y, \omega, \gamma) \quad \text{implies} \quad u^{no}(c_0) > u(C(y), y, \omega, \gamma') .$$

The Spence-Mirrlees single-crossing property preserves monotonicity of choices in variable costs. For a given continuous consumption schedule $C_0 : y \mapsto C_0(y)$, $y_0(\omega, \gamma)$ is the utility-maximizing choice of type (ω, γ) . By the Spence-Mirrlees single-crossing property, $\omega' > \omega$ implies $y_0(\omega', \gamma) \geq y_0(\omega, \gamma)$. In particular, $y_0(\omega', \gamma) = 0$ implies $y_0(\omega, \gamma) = 0$. Thus, for any given γ , there is a cutoff type $\hat{\omega}_0(\gamma)$ so that $\omega < \hat{\omega}_0(\gamma)$ implies $y_0(\omega, \gamma) = 0$, whereas $\omega \geq \hat{\omega}_0(\gamma)$ implies $y_0(\omega, \gamma) > 0$.

The earnings function will generally exhibit an upward jump at $\hat{\omega}(\gamma)$. With C continuous, raising y slightly above 0 comes only with a small gain in consumption utility, but an upward jump of effort costs. Thus, a significant increase of earnings is needed to have a gain in consumption utility that offsets these effort costs. Moreover, by our assumption on fixed costs, $\gamma' > \gamma$ implies that $\hat{\omega}(\gamma') \geq \hat{\omega}(\gamma)$.

The earnings function y_0 is bounded away from zero for all (ω, γ) with $\omega > \hat{\omega}(\gamma)$. Over this domain, we take y_0 to be a non-decreasing function of γ . Thus, $y_0(\omega, \gamma) > 0$, $y_0(\omega, \gamma') > 0$ and $\gamma' > \gamma$ imply that $y_0(\omega, \gamma') \geq y_0(\omega, \gamma)$.⁴⁷

For a given reform direction h , we denote by $y^*(e, \tau, h, \omega, \gamma)$ the solution to the problem

$$\max_{y \geq 0} u(c_0 + e + y - T_0(y) - \tau h(y), y, \omega, \gamma)$$

and the corresponding indirect utility by $v(e, \tau, h, \omega, \gamma)$. The parameter e stands for a source of income that is exogenous from an individuals' perspective. In the

⁴⁷For separable utility functions of the form $u(c, y, \omega, \gamma) = \tilde{u}(c, y, \omega) - \gamma \mathbf{1}_{y>0}$, $y_0(\omega, \gamma) > 0$, $y_0(\omega, \gamma') > 0$ and $\gamma' > \gamma$ imply that $y_0(\omega, \gamma') = y_0(\omega, \gamma)$. For specifications $u(c, y, \omega, \gamma) = \tilde{u}(c - \gamma \mathbf{1}_{y>0}, y, \omega)$ with concave consumption utility, $y_0(\omega, \gamma) > 0$, $y_0(\omega, \gamma') > 0$ and $\gamma' > \gamma$ imply that $y_0(\omega, \gamma') > y_0(\omega, \gamma)$.

subsequent analysis, e will be equal to the change in tax revenues, $e = R(\tau, h)$. If $y^*(e, \tau, h, \omega, \gamma) = 0$, then $v(e, \tau, h, \omega, \gamma) = u^{no}(c_0 + e)$.

The earnings function y^* exhibits a discontinuity at $\hat{\omega}_0(\cdot)$. Earnings are zero for types below $\hat{\omega}(\cdot)$ and bounded away from zero for types above. Individuals with type $\hat{\omega}(\cdot)$ are indifferent between earnings of zero and a strictly positive earnings level. It is convenient to assume that these individuals have positive earnings. Thus, we assume that

$$y^*(R(\tau, h), \tau, h, \hat{\omega}(\cdot), \gamma) > 0 .$$

Intensive-margin responses. For one-bracket reforms, the derivative of the function y^* with respect to τ specifies how earnings respond to small changes in marginal tax rates for incomes that lie in that bracket. These are the behavioral responses at the intensive margin.

Extensive-margin responses. For a given reform direction h , we view the cutoff type $\hat{\omega}$ not only as a function of γ , but also as a function of the size of the reform as measured by τ . Formally, for given γ , the cutoff type $\hat{\omega}(\tau, \gamma)$ is defined as the value of ω that solves

$$u^{no}(c_0 + R(\tau, h)) = v(\tau, \omega, \gamma) .$$

The effect of a small change of the reform intensity τ on the cutoff type $\hat{\omega}$ is obtained by computing a total differential of this equation. This yields, invoking again the envelope theorem,

$$u_c^{no}(\cdot) R_\tau(\tau, h) = u_c(\cdot) (R_\tau(\tau, h) - h(\cdot)) + u_\omega(\cdot) \hat{\omega}_\tau(\tau, \gamma) ,$$

where the functions u_c , u_ω and h are evaluated at $y = y^*(R(\cdot), \tau, \hat{\omega}(\cdot), \gamma)$. Equivalently,

$$\hat{\omega}_\tau(\tau, \gamma) = \frac{u_c^{no}(\cdot)}{u_\omega(\cdot)} R_\tau(\tau, h) - \frac{u_c(\cdot)}{u_\omega(\cdot)} (R_\tau(\tau, h) - h(\cdot)) . \quad (\text{E.1})$$

To interpret these expressions, consider the following thought experiment: a fraction $F_\Omega(\hat{\omega}(\tau, \gamma) \mid \gamma)$ of individuals with fixed cost type γ has zero earnings, where

$F_\Omega(\cdot | \gamma)$ is the distribution of variable cost types ω conditional on the fixed cost type being γ . Consider a small increase of transfers only for the unemployed. The marginal effect on $F_\Omega(\cdot | \gamma)$ is given by

$$P^{no}(\tau, \gamma) := f_\Omega(\hat{\omega}(\tau, \gamma) | \gamma) \frac{u_c^{no}(\cdot)}{u_\omega(\cdot)}, \quad (\text{E.2})$$

where the letter P is chosen to indicate a marginal effect on *Participation*. Alternatively, the effect of a transfer only to those with positive earnings is given by

$$P(\tau, \gamma) = f_\Omega(\hat{\omega}(\tau, \gamma) | \gamma) \frac{u_c(\cdot)}{u_\omega(\cdot)}. \quad (\text{E.3})$$

Thus, assuming that the distribution F and these marginal effects are known is equivalent to assuming that the extensive-margin elasticities and semi-elasticities that are ubiquitous in the related literature are known.

Notation. It is convenient to use a shorthand for endogenous variables at the status quo. For instance, we will occasionally write $\hat{\omega}_0(\gamma) := \hat{\omega}(0, \gamma)$ for the type at the participation margin, among those with fixed cost type γ . We write $y_0(\omega, \gamma) := y^*(0, 0, h, \omega, \gamma)$ for income in the status quo, and similarly for other variables. Given a fixed-cost type γ , we denote by $\omega_0(y, \gamma)$ the variable cost type who chooses earnings of y in the status quo. If we evaluate partial derivatives at the status quo, we occasionally write $\hat{\omega}_{0\tau}(\gamma)$ or $y_{0e}(\omega, \gamma)$ and so on. We, moreover, write $\hat{y}_0(\gamma) := y_0(\hat{\omega}_0(\gamma), \gamma)$ for the status quo income of the cutoff type among those with fixed costs of γ . Finally,

$$\mathcal{K}(\omega_0(y, \gamma) | \gamma) = E_\Omega [T'_0(y_0(s, \gamma)) y_{0e}(s, \gamma) | s \geq \omega_0(y, \gamma), \gamma],$$

is a measure of the size of income effects among those individuals with fixed cost type γ who have earnings exceeding y .

Revenue implications of reforms with one bracket. By our analysis in the previous section, to understand whether a given tax system can be reformed in a Pareto-improving way, we need to check whether the function $y \mapsto R_{\tau\ell}^s(0, 0, y)$

is bounded from below by 0, bounded from above by 1, and non-increasing. The following proposition provides a characterization of this function for the given setup with variable and fixed costs of productive effort. A version of this function in terms of empirically estimable objects can be found in (D.1) in part D of the Supplementary Material.

Proposition E.1 *Suppose that, for any given γ , $\omega_0(y, \gamma)$ is strictly increasing in y , whenever $y > 0$. Also suppose that, for all (ω, γ) , income in the status quo satisfies the first-order conditions of utility-maximization whenever $y_0(\omega, \gamma) > 0$. Then,*

$$R_{\tau\ell}^s(0, 0, y) = (1 - \mathcal{M}_0)^{-1} \left(\mathcal{I}(y) - \mathcal{X}(y) \right),$$

where

$$\mathcal{M}_0 = E_{\Gamma} \left[(1 - F(\hat{\omega}_0(\gamma) \mid \gamma)) \mathcal{K}(\hat{\omega}_0(\gamma) \mid \gamma) - T_0(\hat{y}_0^*(\gamma)) (P_0^{no}(\gamma) - P_0(\gamma)) \right],$$

$$\begin{aligned} \mathcal{I}(y) &= T_0'(y) E_{\Gamma} \left[f(\omega_0(y, \gamma) \mid \gamma) \frac{y_{0\tau}(\omega_0(y, \gamma), \gamma)}{y_{0\omega}(\omega_0(y, \gamma), \gamma)} \right] \\ &\quad + E_{\Gamma} [(1 - \mathcal{K}(\omega_0(y, \gamma) \mid \gamma))(1 - F(\omega_0(y, \gamma) \mid \gamma))], \end{aligned}$$

and

$$\mathcal{X}(y) = \int_{\underline{\gamma}}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) P_0(\gamma) \mathbf{1}(\hat{y}_0(\gamma) \geq y) f_{\Gamma}(\gamma) d\gamma.$$

The proof of Proposition E.1 can be found below. Proposition E.1 shows that the revenue effect of a small single-bracket reform at income level y , $R_{\tau\ell}^s(0, 0, y)$, can be decomposed into a term $\mathcal{I}(y)$ that includes a mechanical effect and an intensive-margin effect, an extensive-margin effect $\mathcal{X}(y)$, and a multiplier \mathcal{M}_0 . The assumptions that $\omega_0(y, \gamma)$ is strictly increasing in y and that individual behavior can be described by first-order conditions are made for ease of exposition. They avoid complications due to bunching.

The term $\mathcal{I}(y)$ consists of a mechanical effect, a behavioral response from an income effect and a behavioral response from a substitution effect. The mechanical effect is that individuals with an income larger than y now pay additional taxes. This yields a revenue gain that is proportional to the mass of these people,

$E_{\Gamma} [1 - F(\omega_0(y, \gamma) | \gamma)]$. With income effects, these people also seek to make up for the fact that the tax reform makes them poorer and they respond with an increase of their earnings. Together the mechanical and the income effect amount to $E_{\Gamma} [(1 - \mathcal{K}(\omega_0(y, \gamma) | \gamma))(1 - F(\omega_0(y, \gamma) | \gamma))]$. The substitution effect is relevant for people with an income of y . They have less of an incentive to exert productive effort when the marginal tax rate for incomes close to y is increased. This is captured by the expression

$$T'_0(y) E_{\Gamma} \left[f(\omega_0(y, \gamma) | \gamma) \frac{y_{0\tau}(\omega_0(y, \gamma), \gamma)}{y_{0\omega}(\omega_0(y, \gamma), \gamma)} \right] .$$

Alternatively, this term can be written as

$$\frac{T'_0(y)}{1 - T'_0(y)} \left(1 + \frac{1}{\varepsilon_0(y)} \right) E_{\Gamma} [f(\omega_0(y, \gamma) | \gamma)] ,$$

where $\varepsilon_0(y)$ is the elasticity of earnings with respect to the net-of-tax rate among people with pre-reform earnings of y .

The extensive-margin effect is shaped by the employment response of those individuals who are close to indifferent between staying out of the labor market and entering. More specifically, $\mathcal{X}(y)$ gives an average for all types who choose earnings of at least y when entering the labor force. A change of the marginal rates in a bracket that begins at y has no effect on individuals who only consider incomes lower than y . For those who consider an income of y or above, there is a negative effect on participation and this tends to lower the revenue that is raised by such a reform. Naturally, the effect of this employment response on tax revenue depends on the tax payment that these individuals pay when entering the labor force and which is lost when they stay out.

The multiplier \mathcal{M}_0 mitigates all previously discussed effects. It reflects income effects at both margins, which appear if utility is non-linear in consumption. In this case, a tax reform that yields a revenue gain and hence increased lump-sum transfers has a negative effect on the earnings of market participants at the intensive margin. This effect is captured by the first component of \mathcal{M}_0 ,

$$E_{\Gamma} \left[(1 - F(\hat{\omega}_0(\gamma) | \gamma)) \mathcal{K}(\hat{\omega}_0(\gamma) | \gamma) \right] .$$

In addition, there are also income effects at the extensive margin. Specifically, the increased lump-sum transfers makes non-participation more attractive with concave consumption utility such that $(P_0^{no}(\gamma) - P_0(\gamma)) > 0$. This effect is captured by the second component in \mathcal{M}_0 . It implies a loss of tax revenue if $T_0(y) > 0$.

E.1 Proof of Proposition E.1

The change in tax revenue associated with a one-bracket reform (τ, ℓ, y_a) can be written as

$$R^s(\tau, \ell, y_a) = \mathbb{E}_\Gamma [\mathcal{R}^s(\tau, \ell, y_a \mid \gamma)] ,$$

where $\mathcal{R}^s(\tau, \ell, y_a \mid \gamma)$ is the change in tax revenue due to individuals with a given fixed cost type γ , and \mathbb{E}_Γ is an expectations operator that indicates the computation of a population average using the marginal distribution F_Γ of fixed costs. Also, note that

$$\begin{aligned} \mathcal{R}^s(\tau, \ell, y_a \mid \gamma) &= \mathbb{E}_\Omega [T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) - T_0(y_0(\omega, \gamma)) \mid \gamma] \\ &= \int_{\underline{\omega}}^{\bar{\omega}} \{T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) - T_0(y_0(\omega, \gamma))\} f_\Omega(\omega \mid \gamma) d\omega , \end{aligned}$$

where $f_\Omega(\cdot \mid \gamma)$ is the density of the conditional distribution of ω for given γ . The change in revenue associated with a marginal change of τ can be written as

$$R_\tau^s(\tau, \ell, y_a) = \mathbb{E}_\Gamma [\mathcal{R}_\tau^s(\tau, \ell, y_a \mid \gamma)] , \tag{E.4}$$

where

$$\begin{aligned} \mathcal{R}_\tau^s(\tau, \ell, y_a \mid \gamma) &= \frac{d}{d\tau} \int_{\underline{\omega}}^{\bar{\omega}} T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) f_\Omega(\omega \mid \gamma) d\omega \\ &= \frac{d}{d\tau} \int_{\hat{\omega}(\tau, \gamma)}^{\bar{\omega}} T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) f_\Omega(\omega \mid \gamma) d\omega . \end{aligned}$$

By standard arguments,

$$\begin{aligned} \mathcal{R}_\tau^s(\tau, \ell, y_a \mid \gamma) &= \frac{d}{d\tau} \int_{\hat{\omega}(\tau, \gamma)}^{\bar{\omega}} T_1(y^*(R^s(\cdot), \tau, h, \omega, \delta)) f_\Omega(\omega \mid \gamma) d\omega \\ &= -T_1(y^*(R^s(\cdot), \tau, h, \hat{\omega}(\cdot), \gamma)) f(\hat{\omega}(\cdot) \mid \gamma) \hat{\omega}_\tau(\cdot) \\ &\quad + \int_{\hat{\omega}(\tau, \gamma)}^{\bar{\omega}} \frac{d}{d\tau} T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) f_\Omega(\omega \mid \gamma) d\omega , \end{aligned} \tag{E.5}$$

i.e., the change of tax revenues can be decomposed into a change that comes from extensive-margin responses and a change that comes from intensive-margin responses to the tax reform.

A small change of the marginal tax rate. Equations (E.4), (E.5) and (E.1) imply that

$$R_\tau^s(0, \ell, y_a) = \mathbb{E}_\Gamma [\mathcal{R}_\tau^s(0, \ell, y_a | \gamma)] , \quad (\text{E.6})$$

where

$$\begin{aligned} \mathcal{R}_\tau^s(0, \ell, y_a | \gamma) &= -T_0(y_0(\hat{\omega}_0(\gamma), \gamma))f(\hat{\omega}_0(\gamma) | \gamma) \hat{\omega}_{0\tau}(\cdot) \\ &\quad + \int_{\hat{\omega}_0(\gamma)}^{\bar{\omega}} \left(\frac{d}{d\tau} T_1(y^*(R^s(\cdot), \tau, h, \omega, \gamma)) \right)_{|\tau=0} f_\Omega(\omega | \gamma) d\omega , \end{aligned} \quad (\text{E.7})$$

and

$$\hat{\omega}_{0\tau} = \frac{u_{0c}^{no}(\cdot)}{u_{0\omega}(\cdot)} R_\tau^s(0, \ell, y_a) - \frac{u_{0c}(\cdot)}{u_{0\omega}(\cdot)} (R_\tau^s(0, \ell, y_a) - h(y_0(\hat{\omega}_0(\gamma), \gamma))) . \quad (\text{E.8})$$

By the arguments in Bierbrauer and Boyer (2018),

$$\begin{aligned} &\int_{\hat{\omega}_0(\gamma)}^{\bar{\omega}} \left(\frac{d}{d\tau} T_1(y^*(R^s(\cdot), \tau, \omega, \delta)) \right)_{|\tau=0} f_\Omega(\omega | \gamma) d\omega \\ &= R_\tau^s(0, \ell, y_a) (1 - F_\Omega(\hat{\omega}_0(\gamma) | \gamma)) \mathcal{K}(\hat{\omega}_0(\gamma) | \gamma) + \mathcal{J}(\ell, y_a | \gamma) \end{aligned} \quad (\text{E.9})$$

where

$$\mathcal{K}(\hat{\omega}_0(\gamma) | \gamma) = E_\Omega [T_0'(y_0(s, \gamma)) y_{0e}(s, \gamma) | s \geq \hat{\omega}_0(\gamma), \gamma] ,$$

as defined in the body of the text, and

$$\begin{aligned} \mathcal{J}(\ell, y_a | \gamma) &= \int_{\omega_0(y_a | \gamma)}^{\omega_0(y_a + \ell | \gamma)} T_0'(y_0(\omega, \gamma)) y_{0\tau}(\omega, \gamma) f(\omega | \gamma) d\omega \\ &\quad + \int_{\omega_0(y_a | \gamma)}^{\omega_0(y_a + \ell | \gamma)} (y_0(\omega, \gamma) - y_a) [1 - T_0'(y_0(\omega, \gamma)) y_{0e}(\omega, \gamma)] f(\omega | \gamma) d\omega \\ &\quad + \ell \left(1 - F(\omega_0(y_a + \ell | \gamma)) \right) \\ &\quad - \int_{\omega_0(y_a + \ell | \gamma)}^{\bar{\omega}} T_0'(y_0(\omega, \gamma)) y_{0e}(\omega, \gamma) f(\omega | \gamma) d\omega \end{aligned}$$

and $\omega_0(y_a | \gamma)$ and $\omega_0(y_a + \ell | \gamma)$ are, respectively, the ω -types who choose income levels of y_a and $y_a + \ell$ in the status quo. Equations (E.2), (E.3), (E.7), (E.8) and (E.9) imply that

$$\begin{aligned} R_\tau^s(0, \ell, y_a | \gamma) &= R_\tau^s(0, \ell, y_a) M_0(\gamma) \\ &\quad - T_0(y_0(\hat{\omega}_0(\gamma), \gamma)) P_0(\gamma) h(y_0(\hat{\omega}_0(\gamma), \gamma)) \\ &\quad + \mathcal{J}(\ell, y_a | \gamma) \end{aligned} \quad (\text{E.10})$$

where the multiplier $M_0(\gamma)$ is given by

$$M_0(\gamma) := (1 - F(\hat{\omega}_0(\gamma)) | \gamma) \mathcal{K}(\hat{\omega}_0(\gamma) | \gamma) - T_0(y_0(\hat{\omega}_0(\gamma), \gamma)) (P_0^{no}(\gamma) - P_0(\gamma)) .$$

We let $\mathcal{M}_0 = E_\Gamma[M_0(\gamma)]$. Then, Equations (E.6) and (E.10) imply that

$$R_\tau^s(0, \ell, y_a) = (1 - \mathcal{M}_0)^{-1} \left(E_\Gamma[\mathcal{J}(\ell, y_a | \gamma)] - E_\Gamma[\mathcal{Z}(\ell, y_a | \gamma)] \right) , \quad (\text{E.11})$$

where

$$\mathcal{Z}(\ell, y_a | \gamma) = T_0(y_0(\hat{\omega}_0(\gamma), \gamma)) P_0(\gamma) h(y_0(\hat{\omega}_0(\gamma), \gamma)) .$$

If y_a is a very high, the tax reform affects only high incomes. Plausibly, y_a is then also above the income level that individuals at the extensive margin would consider, i.e., $y_a > \max_\gamma y_0(\hat{\omega}_0(\gamma), \gamma)$. In this case, $h(y_0(\hat{\omega}_0(\gamma), \gamma)) = 0$ and hence $\mathcal{Z}(\ell, y_a | \gamma) = 0$.

A small change of the marginal tax rate for a narrow bracket. We are interested in determining how much additional tax revenue a reform generates that involves a small change of marginal tax rates for incomes in a narrow bracket. To this end, we provide a characterization of $R_{\tau\ell}^s(0, 0, y_a)$, i.e., of the cross-derivative of tax revenue with respect to the change of the marginal tax rate τ and the length of the bracket ℓ at the status quo. To understand the logic of this exercise, note that our previous derivations imply that $R_\tau^s(0, 0, y_a) = 0$, i.e., if the length of the bracket is zero, the marginal tax rate remains unchanged at each income level. Consequently, there is no effect on tax revenue. If, however, $R_{\tau\ell}^s(0, 0, y_a) > 0$, then moving from $\ell = 0$ to some $\ell = \varepsilon$ for $\varepsilon > 0$ but small, implies that $R_\tau^s(0, \varepsilon, y_a) > 0$, so that a small change of the marginal tax rate then has a positive effect on revenue. More formally, if $R_{\tau\ell}^s(0, 0, y_a) > 0$ there exist $\delta > 0$ and $\varepsilon > 0$ so that $R^s(\delta, \varepsilon, y_a) > 0$.

It follows from Equation (E.11) that

$$R_{\tau\ell}^s(0, 0, y_a) = (1 - \mathcal{M}_0)^{-1} \left(E_\Gamma[\mathcal{J}_\ell(0, y_a | \gamma)] - E_\Gamma[\mathcal{Z}_\ell(0, y_a | \gamma)] \right) , \quad (\text{E.12})$$

where, by the arguments in Bierbrauer and Boyer (2018),

$$\begin{aligned} E_\Gamma[\mathcal{J}_\ell(0, y_a | \gamma)] &= T_0'(y_a) E_\Gamma \left[f(\omega_0(y_a, \gamma) | \gamma) \frac{y_{0\tau}(\omega_0(y_a, \gamma), \gamma)}{y_{0\omega}(\omega_0(y_a, \gamma), \gamma)} \right] \\ &\quad + E_\Gamma [(1 - \mathcal{K}(\omega_0(y_a, \gamma) | \gamma))(1 - F(\omega_0(y_a, \gamma) | \gamma))] \end{aligned} \quad (\text{E.13})$$

We now work towards a characterization of $\mathcal{Z}_\ell(0, y_a) := E_\Gamma[\mathcal{Z}_\ell(0, y_a \mid \gamma)]$. Suppose that $y_a \leq \max_\gamma y_0(\hat{\omega}_0(\gamma), \gamma)$ and denote by $\gamma_0(y_a)$ and $\gamma_0(y_a + \ell)$, respectively, the types at the extensive margin who earn, respectively, incomes of y_a and $y_b = y_a + \ell$, i.e.,

$$y_0(\hat{\omega}_0(\gamma_0(y_a)), \gamma_0(y_a)) = y_a, \quad \text{and} \quad y_0(\hat{\omega}_0(\gamma_0(y_a + \ell)), \gamma_0(y_a + \ell)) = y_a + \ell.$$

Armed with this notation, we can write

$$\begin{aligned} E_\Gamma[\mathcal{Z}_\ell(\ell, y_a \mid \gamma)] &= \int_{\gamma_0(y_a)}^{\gamma_0(y_a + \ell)} T_0(\hat{y}_0(\gamma)) P_0(\gamma) (\hat{y}_0(\gamma) - y_a) f_\Gamma(\gamma) d\gamma \\ &\quad + \ell \int_{\gamma_0(y_a + \ell)}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) P_0(\gamma) f_\Gamma(\gamma) d\gamma, \end{aligned}$$

where $\hat{y}_0(\gamma)$ is a shorthand for $y_0(\hat{\omega}_0(\gamma), \gamma)$. Straightforward computations yield

$$\mathcal{Z}_\ell(0, y_a) = \int_{\gamma_0(y_a)}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) P_0(\gamma) f_\Gamma(\gamma) d\gamma.$$

Now, to accommodate both the case $y_a \leq \max_\gamma \hat{y}_0(\gamma)$ and the case $y_a > \max_\gamma \hat{y}_0(\gamma)$, we will write henceforth

$$\mathcal{Z}_\ell(0, y_a) = \int_{\underline{\gamma}}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) P_0(\gamma) \mathbf{1}(\hat{y}_0(\gamma) \geq y_a) f_\Gamma(\gamma) d\gamma.$$

Proposition E.1 as stated in the text now follows upon adopting the more concise notation

$$\mathcal{I}(y) := \mathcal{J}_\ell(0, y) \quad \text{and} \quad \mathcal{X}(y) := \mathcal{Z}_\ell(0, y),$$

for any earnings level y .

E.2 Implications of Proposition E.1

E.2.1 Diamond's model

With quasi-linear in consumption preferences and without fixed costs of labor market participation, we have

$$\mathcal{M}_0 = 0, \quad \text{and} \quad \mathcal{X}(y) = 0,$$

and

$$\mathcal{I}(y) = T'_0(y) f(\omega_0(y)) \frac{y_{0\tau}(\omega_0(y))}{y_{0\omega}(\omega_0(y))} + 1 - F(\omega_0(y)) . \quad (\text{E.14})$$

This can be rewritten as

$$\mathcal{I}(y) = -\frac{T'_0(y)}{1 - T'_0(y)} f(\omega_0(y)) \left(1 + \frac{1}{\varepsilon_0(y)}\right)^{-1} \omega_0(y) + 1 - F(\omega_0(y)) . \quad (\text{E.15})$$

Getting from Equation (E.14) to Equation (E.15) requires to invoke the first-order condition characterizing $y_0(\omega)$ for the purpose of deriving comparative statics results that yield a characterization of $y_{0\tau}(\omega_0(y))$ and $y_{0\omega}(\omega_0(y))$. Thus,

$$R_{\tau\ell}^s(0, 0, y) = -\frac{T'_0(y)}{1 - T'_0(y)} f(\omega_0(y)) \left(1 + \frac{1}{\varepsilon_0(y)}\right)^{-1} \omega_0(y) + 1 - F(\omega_0(y)) ,$$

which is Equation (10) in the main text using that

$$F_y(y) = F(\omega_0(y)) \quad \text{and} \quad f_y(y) = f(\omega_0(y)) \omega'_0(y) = f(\omega_0(y)) \frac{1}{y_{0\omega}(\omega_0(y))} .$$

and that, for any y in $y_0(\Omega)$, $\varepsilon_0(y) = -\frac{y}{1 - T'_0(y)} y_{0\tau}(\cdot)$.

E.2.2 Fixed costs as an extension of Diamond's model

We now consider an extension of Diamond (1998) that includes fixed costs of labor market participation. Preferences are now given by

$$u(c, y, \omega, \gamma) = c - \frac{1}{1 + \frac{1}{\epsilon}} \left(\frac{y}{\omega}\right)^{1 + \frac{1}{\epsilon}} - \gamma \mathbb{1}_{y > 0} ,$$

where ϵ is a parameter. Again, the absence of income effect implies that $\mathcal{M}_0 = 0$.

We can, once more, rewrite $\mathcal{I}(y)$ using the distribution of incomes F_y so that

$$F_y(y) = E_{\Gamma} [F_{\Omega}(\omega_0(y, \gamma) \mid \gamma)] .$$

This yields

$$\mathcal{I}(y) = -\frac{T'_0(y)}{1 - T'_0(y)} \varepsilon_0(y) y f_y(y) + 1 - F_y(y) .$$

While this expression looks exactly as in Diamond's model, the distribution of incomes F_y is now shaped by the joint distribution of fixed and variable costs. We also rewrite

$$\mathcal{X}(y) = \int_{\underline{\gamma}}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) P_0(\gamma) \mathbf{1}(\hat{y}_0(\gamma) \geq y) f_{\Gamma}(\gamma) d\gamma .$$

in a way that is more handy in the context of our application: We first note that, with quasi-linear in consumption preferences, $P_0(\gamma) = f_\Omega(\hat{\omega}_0(\gamma) \mid \gamma)$, and therefore

$$\mathcal{X}(y) = \int_{\underline{\gamma}}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) \mathbf{1}(\hat{y}_0(\gamma) \geq y) f_\Omega(\hat{\omega}_0(\gamma) \mid \gamma) f_\Gamma(\gamma) d\gamma .$$

We then note that \hat{y}_0 is an increasing function.⁴⁸ Thus, if we denote by $\gamma_0(y)$ be value of γ for which $\hat{y}_0(\gamma) = y$, then we can write $\mathcal{X}(y)$ as

$$\mathcal{X}(y) = \int_{\gamma_0(y)}^{\bar{\gamma}} T_0(\hat{y}_0(\gamma)) f_\Omega(\hat{\omega}_0(\gamma) \mid \gamma) f_\Gamma(\gamma) d\gamma .$$

We seek an interpretation of $\mathcal{X}(y)$ in terms of extensive-margin elasticities. Therefore, consider the following thought experiment: For pre-tax incomes in an interval $[y_0, y]$, after-tax incomes are marginally decreased. Overall employment is then reduced by

$$L_E(y_0, y) = \int_{\gamma_0(y_0)}^{\gamma_0(y)} \frac{f_\Omega(\hat{\omega}_0(\gamma) \mid \gamma)}{u_\omega(\cdot)} f_\Gamma(\gamma) d\gamma .$$

Denote the derivative of $L_E(y_0, y)$ with respect to y by $l_E(y)$ and apply Leibnitz rule to derive

$$l_E(y) = f_\Omega(\hat{\omega}_0(\gamma_0(y)) \mid \gamma_0(y)) u_\omega(\cdot)^{-1} f_\Gamma(\gamma_0(y)) \gamma_0'(y) .$$

Using the fundamental theorem of calculus, we can now also write

$$L_E(y_0, y) = \int_{y_0}^y l_E(y') dy' .$$

Thus, $l_E(y)$ is a marginal effect. One that measures the mass of people flowing out of the pool of people who earn y and into the pool of people who earn zero, when consumption for people earning y is slightly reduced. We will now rewrite $\mathcal{X}(y)$ in a way that highlights the significance of this employment response. We perform a substitution using

$$y' = \hat{y}_0(\gamma) \quad \text{and} \quad dy' = \frac{\partial \hat{y}_0^*(\gamma)}{\partial \gamma} d\gamma .$$

⁴⁸This follows from the following observations. First, recall that, by definition, $\hat{y}_0(\gamma) = y_0(\hat{\omega}_0(\gamma), \gamma)$. Second, with quasi-linear in consumption preferences, the function $y_0 : (\omega, \gamma) \mapsto y_0(\omega, \gamma)$ is increasing in the first argument and constant in the second argument. Third, $\hat{\omega}_0 : \gamma \mapsto \hat{\omega}_0(\gamma)$ is increasing, as higher rents from labor market participation are needed to offset larger fixed costs.

Also, note that $\gamma_0 : y' \mapsto \gamma_0(y')$ is the inverse of the function \hat{y}_0 , so that

$$\gamma_0'(y') dy' = d\gamma$$

Finally, note that $\hat{y}_0(\gamma_0(y')) = y'$. Thus,

$$\begin{aligned} \mathcal{X}(y) &= \int_y^{\bar{y}_0} T_0(y') \frac{f_{\Omega}(\hat{\omega}_0(\gamma_0(y')) | \gamma_0(y'))}{u_{\omega}(\cdot)} f_{\Gamma}(\gamma_0(y')) \gamma_0'(y') dy' \\ &= \int_y^{\bar{y}_0} T_0(y') l_E(y') dy' \\ &= \int_y^{\bar{y}_0} \frac{T_0(y')}{y' - T_0(y')} \frac{l_E(y')}{f_y(y')} (y' - T_0(y')) f_y(y') dy' \\ &= \int_y^{\bar{y}_0} \frac{T_0(y')}{y' - T_0(y')} \pi_0(y') f_y(y') dy' , \end{aligned}$$

where \bar{y}_0 is the highest level of y in the support of F_y . We interpret

$$\pi_0(y') := \frac{l_E(y')}{f_y(y')} (y' - T_0(y'))$$

as an extensive-margin elasticity, it relates a percentage change in the mass of people at pre-tax income y due to extensive-margin responses to a percentage change in their after-tax labor income $y - T_0(y)$. The literature typically refers to $\mathbb{E}_y [\pi_0(y')]$ as the participation elasticity.

Upon collecting terms, and upon assuming an unbounded distribution of income in the status quo, we obtain

$$\begin{aligned} R_{\tau\ell}^s(0, 0, y) &= 1 - F_y(y) - \varepsilon_0(y) y f_y(y) \frac{T_0'(y)}{1 - T_0'(y)} \\ &\quad - \int_y^{\infty} f_y(y') \pi_0(y') \frac{T_0(y')}{y' - T_0(y')} dy' . \end{aligned} \tag{E.16}$$

Note that (E.16) coincides with Equation (11) in the body of the text.

F Analysis of the 2018 US tax-transfer system

There were many reforms of the US tax-transfer system since the mid-1970s. To get a sense of whether there was progress in the US tax policy over time, this section studies the 2018 tax-transfer system for single parents. We find that it was still not Pareto-efficient. But, based on the inefficiency measure introduced in Section 2.4, the inefficiencies in the current tax system are quantitatively an order of magnitude smaller than those in the mid-1970s.

Calibration. Again, our analysis focuses on the group of single parents with two children in California. We parametrize the revenue function given in (11) for the 2018 tax system that applied to this group.

Table S.8: Sources for US tax-transfer system, 2018.

| Information | Sources |
|-------------|--|
| Income tax | Internal Revenue Service, “1040 Instructions, Tax Year 2018”, 2019. Accessible at https://www.irs.gov/pub/irs-dft/i1040gi--dft.pdf . |
| EITC | Tax Policy Center, “Earned Income Tax Credit Parameters, 1975-2018”, 2018. Accessible at https://www.taxpolicycenter.org/file/178859 . |
| AFDC | CCWRO, “CalWORKs (also known as AFDC/TANF)”, 2018. Accessible at https://www.ccwro.org/advocateresources/public-assistance-table . City and County of San Francisco Human Services Agency, “CalWORKs Eligibility Handbook”, 2018. URL: https://www.sfhsa.org/file/6406 . |
| SNAP | New Mexico Department of Human Services, “Income Eligibility Guidelines for SNAP and Financial Assistance”, 2018. URL: https://www.hsd.state.nm.us/uploads/FileLinks/26463f122f47474487faee4922e09ce8/ISD_017_Income_Eligibility_Guidelines_for_SNAP_and_Financial_Assistance_FFY18_1.pdf . |

We take account of the federal income tax and various transfer and welfare programs: the federal Earned Income Tax Credit, the Child Tax Credit (CTC, a partly refundable tax credit introduced in 1998), SNAP and Temporary Assistance for Needy Families (TANF, the successor of AFDC). Table S.8 below presents a summary of the sources that we use for this purpose. As of 2018, the maximum

EITC amounts to 5,716 USD for single taxpayers with two children. It is phased in at a rate of 0.4 for annual incomes below 14,290 USD and phased out at a rate of 0.2106 for incomes between 18,660 and 45,802 USD.

Figure S.1 shows the effective marginal tax, $T'_0(y)$, and the participation tax rate, $T_0(y)/y$, for single parents that result from the interplay of these programs. In 2018, effective marginal tax rates for low incomes are much lower than in the 1970s. Most notably, they are negative due to the EITC phase-in for incomes below 2,700 USD (only), and close to 60% in an income range between 14,300 and 18,500 USD. In contrast to the 1974 tax system, the end of the EITC phase-in range is now closely aligned to the full phasing-out of transfers (TANF, SNAP) around 18,000 USD.

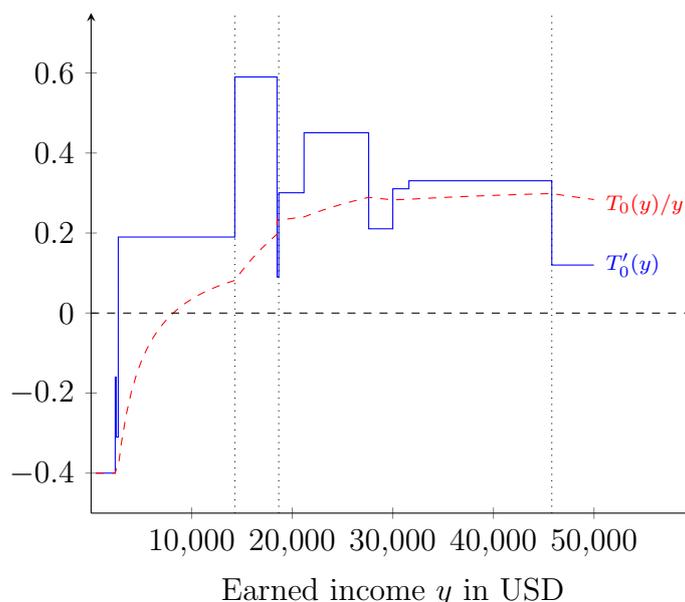


Figure S.1: 2018 US tax-transfer schedule for single parents.

Notes: Figure S.1 shows the 2018 effective marginal tax $T'_0(y)$ (solid blue line) and the participation tax rate $T_0(y)/y$ (dashed red line) for single parents with two children as functions of earned income in 2018 USD. The dotted vertical lines mark the endpoint of the phase-in range at 14,290 USD, the starting point of the phase-out range at 18,660 USD, and the endpoint of the phase-out range at 45,802 USD.

Source: Authors' calculations (see Table S.8 for details).

We estimate the income distribution among single parents based on data from

the March 2019 wave of the CPS. As before, we use a non-parametric kernel density estimation. The share of single parents with no income has gone down considerably since the mid-1970s. In 2018, it amounted to 15.2%, about half of the share in the 1970s. Around 54.1% of single parents had strictly positive incomes below 45,802 USD and were, therefore, eligible for the EITC. In our data, 81.4% of single parents were female.

For the behavioral responses to taxation, we stick to the benchmark case with an intensive-margin elasticity of 0.33 and participation elasticities that are decreasing from 0.67 for very low income levels to 0.4 for incomes above 50,000 USD.

Empirical results. Figure S.2 plots the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for the US tax-transfer system for single parents in 2018. Remember that, for the 1974 US tax system, this revenue function violated two necessary conditions for Pareto efficiency: It was negative at incomes between 2,000 and 5,500 USD, and increasing in the range between 5,000 and 6,000 USD. Hence, there existed Pareto-improving tax reform with one bracket as well as with two brackets. By contrast, the 2018 version of $y \mapsto \mathcal{R}_{sp}(y)$ is bounded from below by 0 and from above by 1. Hence, we no longer find Pareto improvements by means of one-bracket reforms. But there are still several income ranges where the function is increasing. This indicates that, even today, there exist Pareto-improving tax reforms with two brackets.

The visual impression from Figure S.2 is that, despite these non-monotonicities, the inefficiencies in the 2018 US tax system are smaller in magnitude than those in the 1974 tax system. To verify this conjecture, we apply the method introduced in Section 2.4 to find the optimal Pareto-improving tax reform at the time and use the free lunch from this reform as a measure of the current tax system's inefficiency. We find that, in 2018, the optimal tax reform was given by a combination of the following two-bracket tax cuts: (i) one between 15,990 and 20,980 USD that includes the starting point of the EITC phase-out range, (ii) one between 25,200 and 30,000 USD, and (iii) one between 42,359 and 49,245 USD that includes the EITC eligibility threshold. Based on our calibration, this reform would have

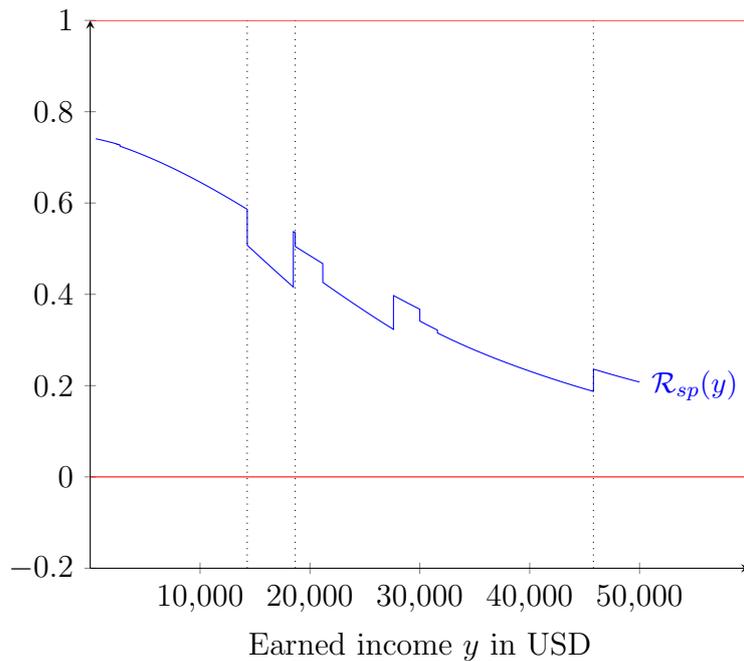


Figure S.2: Revenue function of 2018 US tax-transfer system for single parents.

Notes: Figure S.2 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents in 2018, assuming an intensive-margin elasticity of 0.33 and an participation elasticity that is decreasing from 0.67 at very low incomes to 0.4 at high incomes. The dotted vertical lines mark the endpoint of the phase-in range at 14,290 USD, the starting point of the phase-out range at 18,660 USD, and the endpoint of the phase-out range at 45,802 USD.

Source: Authors' calculations (see Table S.8 for details).

allowed to raise the base transfer to each single-parent household by about three USD (in 2018 units) for each percentage-point change in marginal tax rates. By contrast, the optimal tax reform in 1974 would have allowed to raise the base transfer to single parents by about 12.7 USD in 1974 units, or about 64 USD in 2018 units, per percentage-point change in marginal taxes. We conclude that the set of US tax reforms since the 1970s did not fully eliminate the inefficiencies in the tax-transfer system, but reduced them substantially.

Finally, note that these results for the 2018 US tax system are robust in several dimensions. Most importantly, the recent empirical evidence suggests that labor supply elasticities have decreased in magnitude over time, especially those related

to the extensive margin. We therefore also considered alternative calibrations with smaller labor supply elasticities at both margins. Qualitatively, our results remain unchanged: the marginal revenue function satisfies the lower and upper boundary conditions, but violates the monotonicity condition. Hence, we conclude that the 2018 US tax-transfer system for single parents is still not Pareto-efficient, but these inefficiencies are an order of magnitude smaller than those in the mid-1970s.