Econometrica, Vol. 91, No. 3 (May, 2023), 1077–1103

PARETO-IMPROVING TAX REFORMS AND THE EARNED INCOME TAX CREDIT

FELIX J. BIERBRAUER
CMR, University of Cologne

PIERRE C. BOYER
CREST, École polytechnique, Institut Polytechnique de Paris

EMANUEL HANSEN
CES, Ludwig-Maximilians-Universität Munich and CMR, University of Cologne

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 United States in 1975. A robust finding is that, prior to the EITC, the U.S. 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.

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

Felix J. Bierbrauer: bierbrauer@wiso.uni-koeln.de
Pierre C. Boyer: pierre.boyer@polytechnique.edu
Emanuel Hansen: emanuel.hansen@econ.lmu.de

We thank three anonymous referees for many insightful comments and suggestions. We also 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. Funding Statement: Open Access funding enabled and organized by Projekt DEAL. WOA Institution: Universität zu Köln. Consortia Name: Projekt DEAL.

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 U.S. 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 U.S. 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”; that is, 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”; that is, 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 the additional tax

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

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—that is, 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—that is, 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 United States. Its introduction in 1975 was a substantial policy change for many low-income households; see, for example, 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 4000 USD, and increased them by 10 percentage points in the adjacent phase-out range between 4000 and 8000 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 U.S. 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—that is, the reform yielding the maximal "free lunch"—was qualitatively similar 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 8000 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 U.S. tax-transfer 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, that is, 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) showed 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 by 1. Werning (2007) and Lorenz and Sachs (2016) developed 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 analysed 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

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

[3]There 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).

[4]Lorenz and Sachs (2016) extended 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).

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. 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 U.S. tax system in the mid-1970s, the *introduction* of an EITC can be rationalized as an efficiency-enhancing reform.[9] Note, moreover, that the intro-

[5]For instance, Werning (2007) imposed a specific additively separable utility function and considered 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, for example, 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.

[6]References 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).

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

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

[9]For 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, Tsyvinski, and Werquin (2014), or Jacquet and Lehmann (2021).

duction of the EITC in 1975 did not bring effective tax rates below zero.[10] Our empirical analysis shows that the U.S. tax system 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) estimated 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 focused on a different period and used 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) estimated 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 Appendix A of the Supplemental Material (Bierbrauer, Boyer, and Hansen (2023)). Appendices B and C of the Supplemental Material provide results on the optimal direction for tax reform and an alternative characterization of sufficient conditions for Pareto efficiency, respectively. Appendix D contains a description of the data and the benchmark calibration in our empirical analysis.

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."

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 $u : \mathbb{R}^2_+ \times \Theta \to \mathbb{R}$. 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

[10] Negative effective tax rates were only implemented after subsequent expansions of the federal EITC and the introduction of state EITCs in the 1990s and 2000s.

[11] Further 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) found much smaller ones.

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

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

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

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

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

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

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

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

[13] In the literature, $T_0(y)$ is often referred to as the *participation tax*; see, for example, 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.

[14]There 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.

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

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),

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

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 $h(y_0(\theta))$.

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

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

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 $h \in H$,

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

[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.

[16] For 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.

[17] Our 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

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

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 Appendix A of the Supplemental Material, 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.

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 $(\tau, \ell, \hat{y})$, 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.

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

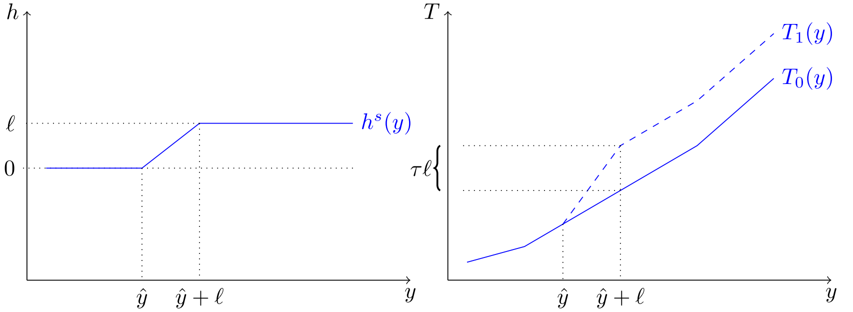

(a) Reform direction h^s. (b) Tax functions T_0 and T_1.

FIGURE 1.—Income tax reforms with one bracket.

[18]Reforms 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.

marginal tax rates equals

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

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^s_τ for the derivative of this function with respect to the first argument, and $R^s_{\tau\ell}$ 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,

$$\tau \ell R^s_{\tau\ell}(0, 0, y)$$

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

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

$$h_2(y) := \tau_1 h^s_1(y) + \tau_2 h^s_2(y), \tag{6}$$

for

$$h^s_1(y) = \begin{cases} 0 & \text{if } y \le y_1, \\ y - y_1 & \text{if } y \in (y_1, y_1 + \ell\ell_1), \\ \ell\ell_1 & \text{if } y \ge y_1 + \ell\ell_1, \end{cases}$$

and

$$h^s_2(y) = \begin{cases} 0 & \text{if } y \le y_2, \\ y - y_2 & \text{if } y \in (y_2, y_2 + \ell\ell_2), \\ \ell\ell_2 & \text{if } y \ge y_2 + \ell\ell_2. \end{cases}$$

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

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

[19]Note that $R^s(\tau, \ell, y) \simeq R^s(0, \ell, y) + \tau R^s_\tau(0, \ell, y) = \tau R^s_\tau(0, \ell, y)$ since $R^s(0, \ell, y) = 0$. Moreover, $\tau R^s_\tau(0, \ell, y) \simeq \tau(R^s_\tau(0, 0, y) + \ell R^s_{\tau\ell}(0, 0, y)) = \tau\ell R^s_{\tau\ell}(0, 0, y)$ since $R^s_\tau(0, 0, y) = 0$.

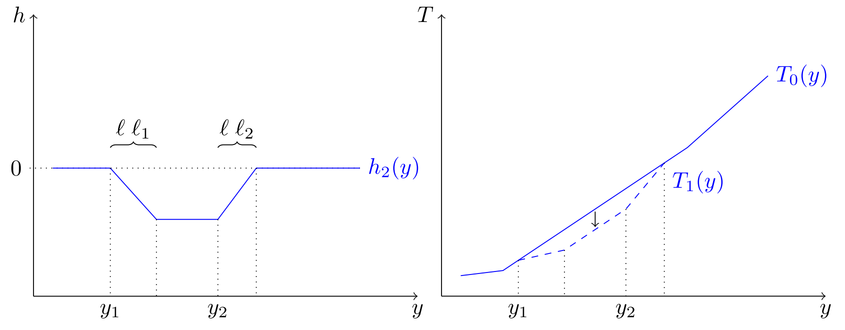

(a) Reform direction h_2. (b) Tax functions T_0 and T_1.

FIGURE 2.—A two-bracket tax cut.

where

$$h_2'(y) = \begin{cases} \tau_1 & \text{for } y \in (y_1, y_1 + \ell\ell_1), \\ \tau_2 & \text{for } y \in (y_2, y_2 + \ell\ell_2), \\ 0 & \text{for } y \leq y_1, y \in [y_1 + \ell\ell_1, y_2], y \geq y_2 + \ell\ell_2. \end{cases}$$

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 $\tau \to 0$ and $\ell \to 0$; 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 such that $h_m(y) :=$

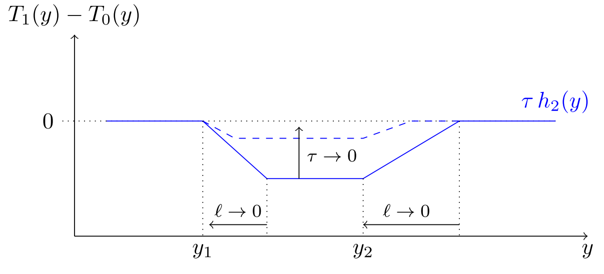

FIGURE 3.—A small two-bracket tax cut.

$\sum_{k=1}^{m} \tau_k h_k^s(y)$, 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, that is, the revenue function is bounded from below by zero and from above by 1. 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 do not 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 2(a). 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—that is, 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 Appendix A of the Supplemental Material 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 have been restricted to the class of tax reforms with finitely many brackets, thereby excluding, for example, 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 $h : \mathcal{Y} \to \mathbb{R}$ 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 $h : \mathcal{Y} \to \mathbb{R}$*.

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 Appendix C of the Supplemental Material, 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 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 Appendix A of the Supplemental Material, 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 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, Jongen,

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

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

[22]This 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 fewer resources that can be used to compensate those adversely affected by the tax reform.

[23]See, for example, Blundell et al. (2009), Bargain et al. (2011), Bourguignon and Spadaro (2012), Lockwood and Weinzierl (2016), Jacobs, Jongen, and Zoutman (2017), or Hendren (2020).

and Zoutman (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 Appendix A. For a 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

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

[24]See 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.

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 $h : \mathcal{Y} \to \mathbb{R}$. 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, $|h'(y)| \le a$ 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 Appendix B of the Supplemental Material characterizes the optimal reform of a tax system that is inefficient because the revenue function $y \mapsto \mathcal{R}(y)$ is increasing 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

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

[26] An 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 was taken by Werning (2007) who proposed to measure the inefficiency of a tax system by the additional tax revenue that can be realized without making anyone worse off. Werning, however, did not specify how this tax revenue would be used.

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

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

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

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 approximation[28]

$$R(\tau, h^*) - \max_{y\in y_0(\Theta)} \tau h^*(y) \simeq \tau a I(T_0, c_0), \tag{9}$$

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.

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 $u(c, y, \theta) = c - \frac{1}{1+\frac{1}{\epsilon}}(\frac{y}{\theta})^{1+\frac{1}{\epsilon}}$, 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

$$\mathcal{R}(y) = 1 - F_y(y) - \varepsilon_0(y) y f_y(y) \frac{T_0'(y)}{1 - T_0'(y)}, \tag{10}$$

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 was first derived by Jacquet, Lehmann, and Van

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

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

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

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

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 Bierbrauer, Boyer, and Hansen (2022), we also derived function $y \mapsto \mathcal{R}(y)$ for a more general framework (see Proposition E.1).[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 U.S. 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). 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 U.S. 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

[30]This 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.

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

[32]While 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*.

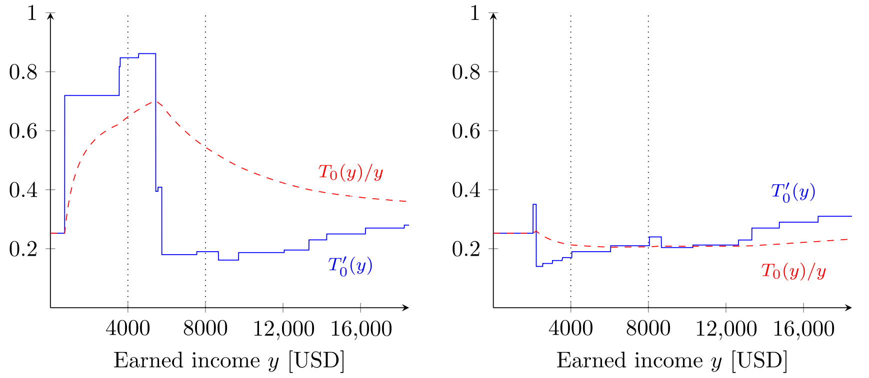

(a) Single parents. (b) Childless singles.

FIGURE 4.—1974 U.S. tax-transfer schedules, single parents and childless singles. *Note*: Figure 4 shows the 1974 effective marginal tax $T'_0(y)$ (solid lines) and participation tax rate $T_0(y)/y$ (dashed 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 4000 USD (first dotted vertical line) and increased them between 4000 and 8000 USD (second dotted vertical line). *Source*: Authors' calculations (see Appendix D of the Supplemental Material for details).

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 4000 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 4000 and 8000 USD. Taxpayers with incomes above 8000 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).

We use data on the most important elements of the U.S. 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 United States

[33]According to CPS data, 51.3% of single parents in the U.S. had earned incomes in the EITC range (i.e., strictly positive and below 8000 USD), and another 30.9% had no earned incomes. Similar figures applied at the state level, for example, 52.2% of single parents with earned incomes in the EITC range and 25.8% with zero incomes in California.

[34]For 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, U.S. authorities also enacted a more modest EITC for childless workers. See Hoynes (2019) for a review.

[35]See Table D.2 in Appendix D of the Supplemental Material for details and sources.

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 4000 USD (first dotted line) and raised them by 10 percentage points in the phase-out range between 4000 and 8000 USD (second dotted line).[39]

We estimate the 1974 income distributions in both subgroups based on data from the March 1975 Current Population Survey (CPS); see Flood, King, Rodgers, Ruggles, Warren, and Westberry (2021). 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 8000 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 Appendix D 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 *s*ingle *p*arents and *c*hildless *s*ingles, respectively. We then use these functions to investigate, first, whether the U.S. tax-transfer system was Pareto-efficient prior to the EITC introduction and, second, whether the 1975 EITC introduction

[36]In Bierbrauer, Boyer, and Hansen (2022), we also present results for the four largest U.S. states next to California: New York, Texas, Pennsylvania, and Illinois.

[37]In our data, the median number of children in single-parent households was two, while the arithmetic mean was 2.2. As we show in Bierbrauer, Boyer, and Hansen (2022), an analysis for single parents with fewer or more children yields similar conclusions.

[38]Recall that we define $T_0(y)$ to capture the participation tax at income y, that is, 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, for example, Kleven (2014).

[39]In Appendix D.4 of Bierbrauer, Boyer, and Hansen (2022), we provide more detailed information on how effective marginal changed from 1974 to 1975.

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," that is, the tax revenue in excess of what was needed to make sure that no one was made worse off.

Was the 1974 U.S. 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 line depicts the revenue function $\mathcal{R}_{sp}(y)$ for single parents, while the dashed line depicts the revenue function $\mathcal{R}_{cs}(y)$ for childless singles.

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 1500 and 5400 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 5000 and 6000 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 dashed line is slightly increasing in the ranges between 2000 and 2500 USD and between 8500 and 9000 USD, respectively. 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 Section 2.4.

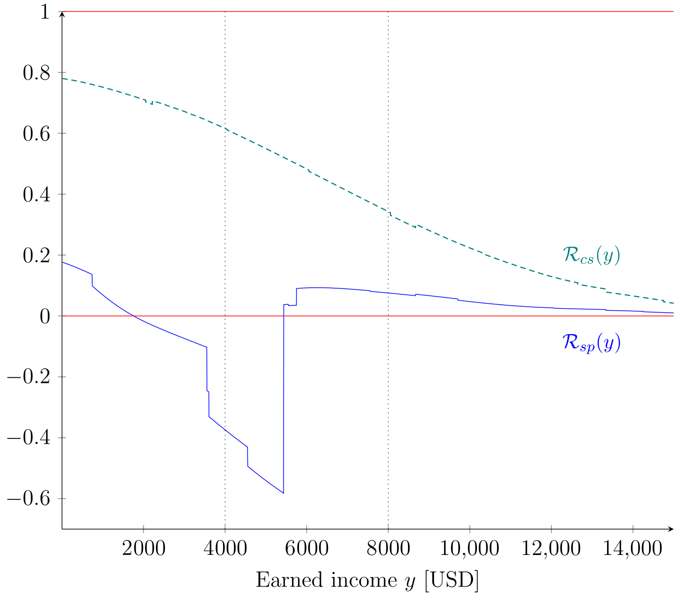

FIGURE 5.—Revenue functions of the 1974 US tax-transfer system. *Note*: Figure 5 shows the revenue functions for single parents $y \mapsto \mathcal{R}_{sp}(y)$ (solid) and childless singles $y \mapsto \mathcal{R}_{cs}(y)$ (dashed) in 1974 as functions of earned income for our benchmark calibration: intensive-margin elasticities of 0.33 for both groups, average participation elasticities of 0.58 for single parents and 0.25 for childless singles. The vertical dotted lines mark the endpoints of the phase-in range at 4000 USD and the phase-out range at 8000 USD of the 1974 EITC. *Source*: Authors' calculations (see Online Appendix D for details).

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 (solid line), for a case with a larger participation elasticity of 0.9 (dashed line), and for the limit case with a participation elasticity of zero (dash-dotted 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, for example, 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 Bierbrauer, Boyer, and Hansen (2022), we provide further robustness 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; see Appendices D.2 and D.3. 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 alter-

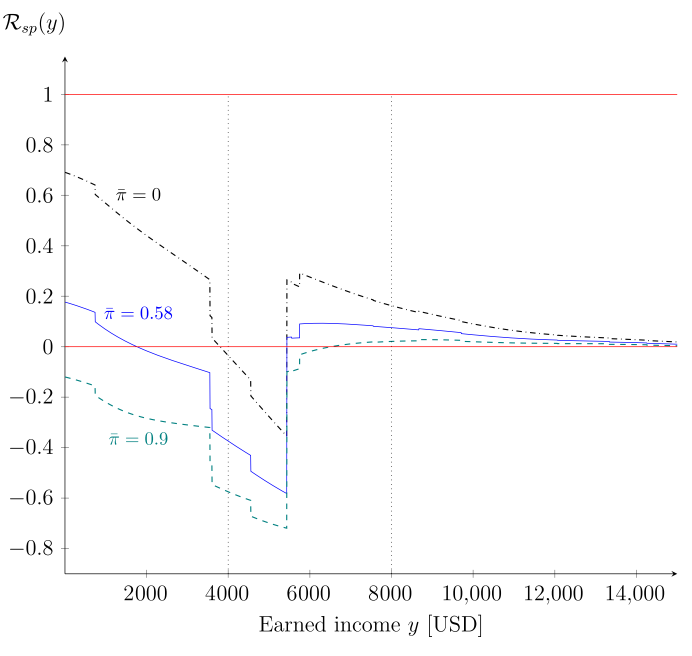

FIGURE 6.—Revenue function of the 1974 U.S. tax-transfer system for single parents, different participation elasticities. *Note*: Figure 6 shows the revenue function $\mathcal{R}_{sp}$ for single parents in 1974, assuming an average participation elasticity of 0.58 (solid line, benchmark), a higher participation elasticity of 0.9 (dashed line) and a case without extensive-margin responses (dash-dotted line). The intensive-margin elasticity is held at the benchmark level of 0.33. The dotted vertical lines mark the endpoints of the phase-in range at 4000 USD and the phase-out range at 8000 USD of the 1975 EITC. *Source*: Authors' calculations (see Appendix D of the Supplemental Material for details).

native representations of the tax-transfer system, for example, including wealth tests or payroll taxes, focusing on single parents with one or three children, or using the transfer schedules from other U.S. 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 U.S. tax-transfer system by a two-bracket tax cut, that is, 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 in 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 4000 USD by 10 percentage points, and increased marginal taxes at all incomes between 4000 and 8000 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

$$R_\tau(0, \tilde{h}_{75}) = -\int_0^{4000} \mathcal{R}_{sp}(y)\, dy + \int_{4000}^{8000} \mathcal{R}_{sp}(y)\, dy > 0 \tag{12}$$

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]

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 1246 and 5748 USD, and increased them between 5748 and 10,250 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 has been 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 U.S. tax policy for people with low incomes? Our approach to study

[40] Some earlier studies indeed estimated participation elasticities in this range; see, for example, Meyer and Rosenbaum (2001).

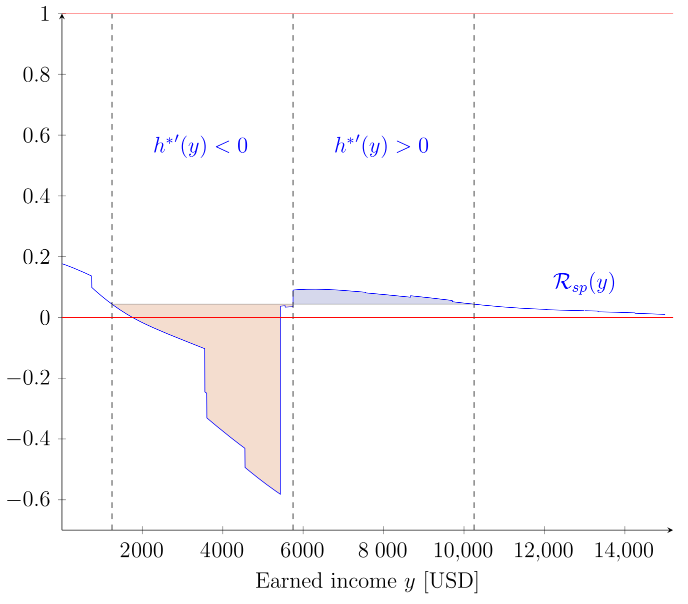

FIGURE 7.—The optimal Pareto-improving reform for single parents. *Note*: Figure 7 shows the revenue function $y \mapsto \mathcal{R}_{sp}(y)$ for single parents given our benchmark calibration. The optimal tax reform would have reduced marginal taxes between 1246 (first dashed line) and 5748 USD (second dashed line) and increased them between 5748 and 10,250 USD (third dashed line). The sum of both shaded areas represents the available revenue gain per single-parent family. *Source*: Authors' calculations (see Appendix D of the Supplemental Material for details).

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 Appendix D.4 of Bierbrauer, Boyer, and Hansen (2022).

First, we study the efficiency of the U.S. 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 1246 and 5748 USD. The 1975 EITC only reduced marginal taxes below 4000 USD, however, and increased them between 4000 and 8000 USD. Thus, the actual reform aggravated inefficiencies at the bottom of the phase-out range. Correspondingly, the 1975 U.S. tax-transfer system did not become Pareto-efficient, and there was further scope for Pareto-improving reforms at higher incomes. In 1979, the U.S. government indeed extended the EITC to higher income levels—by means of a two-bracket tax cut affecting incomes between 4000 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

[41]We 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.

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 realize larger efficiency gains. The 1975 EITC did not condition on the number of children, by contrast. More than a decade later, the U.S. 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, 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 U.S. 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, EHTISHAM, AND NICOLAS STERN (1984): "The Theory of Reform and Indian Indirect Taxes," *Journal of Public Economics*, 25, 259–298. [1081]

AKERLOF, GEORGE A. (1978): "The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning," *American Economic Review*, 68, 8–19. [1090]

BADEL, ALEJANDRO, AND MARK HUGGETT (2017): "The Sufficient Statistic Approach: Predicting the Top of the Laffer Curve," *Journal of Monetary Economics*, 87, 1–12. [1080]

BARGAIN, OLIVIER, MATHIAS DOLLS, DIRK NEUMANN, ANDREAS PEICHL, AND SEBASTIAN SIEGLOCH (2011): "Tax-Benefit Systems in Europe and the US: Between Equity and Efficiency," IZA Discussion Paper No. 5440, https://docs.iza.org/dp5440.pdf. [1081,1090]

BASTANI, SPENCER, SÖREN BLOMQUIST, AND LUCA MICHELETTO (2020): "Pareto Efficient Income Taxation Without Single-Crossing," *Social Choice and Welfare*, 55, 547–594. [1080]

BASTIAN, JACOB (2020): "The Rise of Working Mothers and the 1975 Earned Income Tax Credit," *American Economic Journal: Economic Policy*, 12, 44–75. [1079,1082,1094,1096]

BASTIAN, JACOB E., AND MAGGIE R. JONES (2021): "Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Public Assistance Spending," *Journal of Public Economics*, 196, 104355. [1082]

BIERBRAUER, FELIX J., PIERRE C. BOYER, AND EMANUEL HANSEN (2022): "Pareto-Improving Tax Reforms and the Earned Income Tax Credit," CESifo Working Paper No. 8358, https://www.cesifo.org/DocDL/cesifo1_wp8358.pdf. [1094,1096,1098,1100]

——— (2023): "Supplement to 'Pareto-Improving Tax Reforms and the Earned Income Tax Credit'," *Econometrica Supplemental Material*, 91, https://doi.org/10.3982/ECTA18600. [1082]

BIERBRAUER, FELIX J., PIERRE C. BOYER, AND ANDREAS PEICHL (2021): "Politically Feasible Reforms of Nonlinear Tax Systems," *American Economic Review*, 111, 153–191. [1080]

BLUNDELL, RICHARD (2006): "Earned Income Tax Credit Policies: Impact and Optimality: The Adam Smith Lecture, 2005," *Labour Economics*, 13, 423–443. [1082]

BLUNDELL, RICHARD, AND ANDREW SHEPHARD (2012): "Employment, Hours of Work and the Optimal Taxation of Low-Income Families," *Review of Economic Studies*, 79, 481–510. [1080]

BLUNDELL, RICHARD, MIKE BREWER, PETER HAAN, AND ANDREW SHEPHARD (2009): "Optimal Income Taxation of Lone Mothers: An Empirical Comparison of the UK and Germany," *Economic Journal*, 119, F101–F121. [1081,1090]

BOADWAY, ROBIN, AND MICHAEL KEEN (2000): "Redistribution," in *Handbook of Income Distribution*, Vol. 1, ed. by Anthony B. Atkinson and François Bourguignon. Amsterdam: Elsevier, 677–789. [1080]

BOURGUIGNON, FRANÇOIS, AND AMEDEO SPADARO (2012): "Tax-Benefit Revealed Social Preferences," *Journal of Economic Inequality*, 10, 75–108. [1081,1090,1091]

BRITO, DAGOBERT L., JONATHAN H. HAMILTON, STEVEN M. SLUTSKY, AND JOSEPH E. STIGLITZ (1990): "Pareto Efficient Tax Structures," *Oxford Economic Papers*, 42, 61–77. [1080]

CHETTY, RAJ, ADAM GUREN, DAY MANOLI, AND ANDREA WEBER (2013): "Does Indivisible Labor Explain the Difference Between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," *NBER Macroeconomics Annual*, 27, 1–56. [1096]

CHONÉ, PHILIPPE, AND GUY LAROQUE (2011): "Optimal Taxation in the Extensive Model," *Journal of Economic Theory*, 146, 425–453. [1081]

CHRISTIANSEN, VIDAR, AND EILEV S. JANSEN (1978): "Implicit Social Preferences in the Norwegian System of Indirect Taxation," *Journal of Public Economics*, 10, 217–245. [1081]

DIAMOND, PETER A. (1998): "Optimal Income Taxation: An Example With a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 88, 83–95. [1093]

EISSA, NADA, AND HILARY W. HOYNES (2004): "Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit," *Journal of Public Economics*, 88, 1931–1958. [1082]

EISSA, NADA, AND JEFFREY B. LIEBMAN (1996): "Labor Supply Response to the Earned Income Tax Credit," *Quarterly Journal of Economics*, 111, 605–637. [1082]

FELDSTEIN, MARTIN (1976): "On the Theory of Tax Reform," *Journal of Public Economics*, 6, 77–104. [1081]

FLOOD, SARAH, MIRIAM KING, RENAE RODGERS, STEVEN RUGGLES, J. ROBERT WARREN, AND MICHAEL WESTBERRY (2021): "Integrated Public Use Microdata Series, Current Population Survey: Version 9.0 [Dataset]," IPUMS, Minneapolis, https://doi.org/10.18128/D030.V10.0. [1096]

GOLOSOV, MIKHAIL, ALI SHOURIDEH, MAXIM TROSHKIN, AND ALEH TSYVINSKI (2013): "Optimal Pension Systems With Simple Instruments," *American Economic Review: Papers and Proceedings*, 103, 502–507. [1093]

GOLOSOV, MIKHAIL, ALEH TSYVINSKI, AND NICOLAS WERQUIN (2014): "A Variational Approach to the Analysis of Tax Systems," NBER Working Paper 20780, https://www.nber.org/system/files/working_papers/w20780/w20780.pdf. [1078,1081]

GUESNERIE, ROGER (1995): *A Contribution to the Pure Theory of Taxation*. Cambridge, UK: Cambridge University Press. [1081]

HANSEN, EMANUEL (2021): "Optimal Income Taxation With Labor Supply Responses at Two Margins: When Is an Earned Income Tax Credit Optimal?" *Journal of Public Economics*, 195, 104365. [1081,1098]

HENDREN, NATHANIEL (2020): "Measuring Economic Efficiency Using Inverse Optimum Weights," *Journal of Public Economics*, 187, 104198. [1080,1081,1090]

HOTZ, V. JOSEPH, AND JOHN KARL SCHOLZ (2003): "The Earned Income Tax Credit," in *Means-Tested Transfer Programs in the United States*, ed. by Robert A. Moffitt. Chicago: University of Chicago Press, 141–198. [1082]

HOYNES, HILARY (2019): "The Earned Income Tax Credit," *Annals of the American Academy of Political and Social Science*, 686, 180–203. [1082,1095]

JACOBS, BAS, EGBERT L. W. JONGEN, AND FLORIS T. ZOUTMAN (2017): "Revealed Social Preferences of Dutch Political Parties," *Journal of Public Economics*, 156, 81–100. [1081,1090,1091]

JACQUET, LAURENCE, AND ETIENNE LEHMANN (2021): "Optimal Income Taxation With Composition Effects," *Journal of the European Economic Association*, 19, 1299–1341. [1078,1081]

JACQUET, LAURENCE, ETIENNE LEHMANN, AND BRUNO VAN DER LINDEN (2013): "Optimal Redistributive Taxation With Both Extensive and Intensive Responses," *Journal of Economic Theory*, 148, 1770–1805. [1081, 1093,1094]

JEWITT, IAN (1988): "Justifying the First-Order Approach to Principal-Agent Problems," *Econometrica*, 56, 1177–1190. [1081]

KLEVEN, HENRIK (2021): "The EITC and the Extensive Margin: A Reappraisal," NBER Working Paper 26405, https://www.nber.org/system/files/working_papers/w26405/w26405.pdf. [1082]

KLEVEN, HENRIK JACOBSEN (2014): "How Can Scandinavians Tax so Much?" *Journal of Economic Perspectives*, 28, 77–98. [1083,1096]

KOEHNE, SEBASTIAN, AND DOMINIK SACHS (2022): "Pareto-Improving Reforms of Tax Deductions," *European Economic Review*, 148, 104214. [1080]

LOCKWOOD, BENJAMIN B., AND MATTHEW C. WEINZIERL (2016): "Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions," *Journal of Monetary Economics*, 14 (119), 30–47. [1081,1090]

LORENZ, NORMANN, AND DOMINIK SACHS (2016): "Identifying Laffer Bounds: A Sufficient-Statistics Approach With an Application to Germany," *Scandinavian Journal of Economics*, 118, 646–665. [1080, 1094]

MEYER, BRUCE D., AND DAN T. ROSENBAUM (2001): "Welfare, the Earned Income Tax Credit and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 116, 1063–1114. [1082,1099]

MICHAU, JEAN-BAPTISTE (2014): "Optimal Redistribution: A Life-Cycle Perspective," *Journal of Public Economics*, 111, 1–16. [1093]

MILGROM, PAUL, AND ILYA SEGAL (2002): "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583–601. [1084]

MIRRLEES, JAMES A. (1999): "The Theory of Moral Hazard and Unobservable Behaviour: Part I," *Review of Economic Studies*, 66, 3–21. [1081]

MOFFITT, ROBERT A. (2003): "The Negative Income Tax and the Evolution of U.S. Welfare Policy," *Journal of Economic Perspectives*, 17, 119–140. [1082,1094]

NICHOLS, AUSTIN, AND JESSE ROTHSTEIN (2015): "The Earned Income Tax Credit (EITC)," NBER Working Paper 21211, https://www.nber.org/system/files/working_papers/w21211/w21211.pdf. [1082,1094]

PIKETTY, THOMAS (1997): "La Redistribution Fiscale Face au Chômage," *Revue française d'économie*, 12, 157–201. [1078,1081]

PIKETTY, THOMAS, AND EMMANUEL SAEZ (2013): "Optimal Labor Income Taxation," in *Handbook of Public Economics*, Vol. 5, ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez. Amsterdam: Elsevier, 391–474. [1090]

ROGERSON, WILLIAM P. (1985): "The First-Order Approach to Principal-Agent Problems," *Econometrica*, 53, 1357–1367. [1081]

SAEZ, EMMANUEL (2001): "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 68, 205–229. [1078,1081,1090]

——— (2002): "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *Quarterly Journal of Economics*, 117, 1039–1073. [1081,1098]

SAEZ, EMMANUEL, AND STEFANIE STANTCHEVA (2016): "Generalized Social Marginal Welfare Weights for Optimal Tax Theory," *American Economic Review*, 106, 24–45. [1081]

SCHEUER, FLORIAN (2014): "Entrepreneurial Taxation With Endogenous Entry," *American Economic Journal: Economic Policy*, 6, 126–163. [1080]

SHOURIDEH, ALI, AND MAXIM TROSHKIN (2017): "Incentives and Efficiency of Pension Systems," Working paper, https://www.andrew.cmu.edu/user/ashourid/papers/Troshkindraft56.pdf. [1093]

STANTCHEVA, STEFANIE (2014): "Optimal Income Taxation With Adverse Selection in the Labour Market," *Review of Economic Studies*, 81, 1296–1329. [1080]

STIGLITZ, JOSEPH E. (1982): "Self-Selection and Pareto-Efficient Taxation," *Journal of Public Economics*, 17, 213–240. [1080]

——— (1987): "Pareto Efficient and Optimal Taxation and the New New Welfare Economics," in *Handbook of Public Economics*, Vol. 2, ed. by Alan J. Auerbach and Martin Feldstein. Amsterdam: Elsevier, 991–1042. [1080]

VENTRY, DENNIS J. (2000): "The Collision of Tax and Welfare Politics: The Political History of the Earned Income Tax Credit, 1969–99," *National Tax Journal*, 53, 983–1026. [1094]

WERNING, IVÁN (2007): "Pareto Efficient Income Taxation," Working paper, https://economics.mit.edu/research/publications/pareto-efficient-income-taxation. [1080,1081,1092]

WEYMARK, JOHN A. (1981): "Undominated Directions of Tax Reform," *Journal of Public Economics*, 16, 343–369. [1081]

Co-editor Alessandro Lizzeri handled this manuscript.

Manuscript received 16 June, 2020; final version accepted 9 October, 2022; available online 6 December, 2022.