WELL-BEING, POVERTY, AND LABOR INCOME TAXATION: THEORY AND APPLICATION TO EUROPE AND THE UNITED STATES

François Maniquet, Dirk Neumann

REPRINT | 3200

UCLouvain

CORE | LIDAM
Louvain Institute of Data Analysis and Modeling in economics and statistics

American Economic Journal: Microeconomics 2021, 13(2): 276–310
https://doi.org/10.1257/mic.20180269

Well-Being, Poverty, and Labor Income Taxation: Theory and Application to Europe and the United States[†]

By François Maniquet and Dirk Neumann*

In a model where agents differ in wages and preferences over labor time–consumption bundles, we study labor income tax schemes that alleviate poverty. To avoid conflict with individual well-being, we require redistribution to take place between agents on both sides of the poverty line provided they have the same labor time. This requirement is combined with efficiency and robustness properties. Maximizing the resulting social preferences under incentive compatibility constraints yields the following evaluation criterion: tax schemes should minimize the labor time required to reach the poverty line. We apply this criterion to European countries and the United States. (*JEL* H23, H24, I31, I32, J22)

Eradicating poverty is a social objective that many embrace. In developed societies, this objective often requires that all incomes shall be above some threshold, the so-called poverty line. The United States (US) and all European Union (EU) member states, for instance, publish official poverty lines and official poverty rates, defined as the fraction of the population with incomes below the poverty line. The EU Horizon 2020 objectives include a decrease in the number of people at risk of poverty, that is, having a disposable income below their national poverty line.

With an objective of poverty reduction stated in terms of poverty in income, tax-transfer systems are usually evaluated on the basis of the distribution of incomes they generate, and, in particular, on the distribution of incomes below the poverty line. This creates one main difficulty.

*Maniquet: CORE, Université catholique de Louvain, Voie du Roman Pays 34, 1348 Louvain-la-Neuve, Belgium (email: francois.maniquet@uclouvain.be); Neumann: Federal Ministry for Economic Affairs and Energy, Scharnhorststr. 34-37, 10115 Berlin, Germany (email: Dirk.Neumann@bmwi.bund.de). John Asker was coeditor for this article. We thank four anonymous referees; Roland Benabou, Craig Brett, Koen Decancq, Marc Fleurbaey, Erwin Ooghe, Stefanie Stantcheva, Matthew Weinzierl, John Weymark; discussants and conference participants in Aix-en-Provence (LAGV), Alba-di-Canazei, Italy (Winter School on Inequality and Social Welfare Theory), Antwerp (Inclusive Growth Research Infrastructure Diffusion), Boston (National Tax Association), Louvain-la-Neuve (CORE@50), Lund (Society for Social Choice and Welfare), Mannheim (EEA, ZEW Public Finance), Marseille (CEAFE); as well as seminar participants in Brussels, Leuven, and Louvain-la-Neuve for helpful comments. This paper uses the OECD tax-benefit calculator. We thank the OECD for granting us access to the calculator and Sean Gibson and Daniele Pacifico for detailed assistance on how to use it. Any errors as well as the views presented in this paper are the responsibility of the authors alone. In particular, the views do not represent the official positions of the German Federal Ministry for Economic Affairs and Energy and any other organizations to which the authors are affiliated. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement 269831.

[†]Go to https://doi.org/10.1257/mic.20180269 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

Focusing on income disregards the labor time it takes agents to earn it. However, labor time is also a determinant of well-being, at least if one defines well-being in a consistent way with preference satisfaction as we do in this paper.[1] Thus, when increasing income from below to above the poverty line goes together with an increase in the labor time, antipoverty policies may decrease the well-being of the income-poor, namely if the latter actually prefer to work less and consume less. This tension is reflected in the existing optimal tax literature sharing the objective of income poverty reduction, which we discuss in Section I.

In this paper, we drop the classical objective of income poverty reduction, and we propose an antipoverty requirement that does not conflict with individual well-being. Based on this requirement, we build social preferences, and we derive a criterion to evaluate tax schemes. We apply this criterion to a selected sample of OECD countries.

More precisely, we assume that agents may differ not only in their productivity, like in large parts of the literature mentioned above, but also in their preferences, so that the same opportunity set may lead different agents to different income levels. The key property we require from social preferences is that a transfer from an agent above the poverty line to another agent below the poverty line be a social improvement if both agents have the same labor time. This requirement turns out to be compatible with the property of Pareto efficiency. We prove that together with an auxiliary robustness property, these two properties characterize unique social preferences, which are egalitarian in a specific well-being index representing the preferences of the agents.

Next, we turn to the optimal tax exercise, and we study the consequence of maximizing the social preferences we have characterized under incentive compatibility constraints. These constraints are satisfied when agents choose their labor time given a tax function that determines consumption as a function of earnings. We study both the shape of the optimal tax scheme and the ranking of suboptimal ones. Our main contribution is the construction of a criterion that can be used to evaluate existing income tax schemes. According to that criterion, the pretax earning level (or, equivalently, the labor time needed by minimum wage workers) that is required to reach an after-tax income equal to the poverty line should be minimized.

This criterion allows us to evaluate tax schemes independently of the distribution of incomes they generate. That is, only the opportunities offered to agents, i.e., their budgets, are considered. Our criterion plays the role of a one-number summary of these budgets. The simplicity of the criterion implies that we actually do not need to estimate anything about how individuals react to policy changes to perform welfare analysis. No structural nor reduced-form econometrics is necessary. The sufficient statistics about behavior is entirely contained in the value of our criterion, which can be directly computed from a given tax scheme and poverty line.

[1] Defining well-being consistently with preferences follows from the assumption that choices reveal what is better for the agent (see Hausman 2012 for a long discussion of that assumption). This is the assumption we make, and it is necessary to make sense to the Pareto criterion. The recent behavioral literature, however, has pointed out that circumstances in which poor people make choices may prevent them from revealing their "true" preferences. We discuss the consequences of these observations for our approach further below.

We apply the criterion to evaluate the United States (state of Michigan) and European labor income tax schemes. Those schemes are derived based on simulations using the OECD tax-benefit calculator. We identify the directions of fiscal reforms that would increase social welfare in each of the considered countries. As a by-product of this analysis, we also compare countries and identify those with better tax schemes according to our criterion. It turns out that countries differ widely in the opportunities they offer to poor households, and those performing better are not the countries that transfer the largest benefits to the households whose members do not work at all.

Interestingly, the labor time that is required to reach an after-tax income equal to the poverty line is regularly calculated and published by the OECD as one of their main indicators to evaluate tax-benefit policies.[2] This shows that our criterion can be and is already applied to some extent, to give policy advice to OECD member countries. In that context, our contribution can be understood as providing a theoretical foundation for that criterion and to give it a more central role among a set of different indicators, such as the one provided by the OECD, which also includes indicators for dimensions like "income adequacy" and "work incentives." Indeed, the criterion we advocate for is derived from a complete and transitive notion of social welfare (from which, therefore, income adequacy can also be deduced), and it takes incentive questions into account.

The paper is organized as follows. In Section I, we discuss connections to the existing literature, in particular on optimal taxation, but also including a critical discussion on the notion of preferences and well-being adopted in this paper as well as some other assumptions. In Section II, we present the model. In Section III, we introduce our key property of poverty reduction that does not conflict with individual well-being, and we characterize social preferences. In Section IV, we move to a second-best context and derive an evaluation criterion for tax functions. In Section V, we study the shape of the optimal tax scheme following from our criterion, and we discuss its relationship with tax schemes derived from related social objectives. In Section VI, we use the criterion derived in Section IV to evaluate US and European labor income tax schemes for lone-parent households with two children, under specific assumptions on the poverty line, the relevant social policies, and the minimum wage. In Section VII, we give some concluding comments.

In an online Appendix, we test the robustness of our application with respect to the assumptions we make in the paper over the level of the poverty line and the policies that are simulated. We also identify how each national tax scheme treats different types of households differently. Finally, we identify how our evaluation of tax schemes can be adapted to account for the heterogeneity of unemployment rates among the countries under analysis.

[2] See the website of the OECD tax-benefit calculator, OECD (2019a), and OECD (2019b).

I. Related Literature

The present paper connects to the literature studying optimal taxation in combination with the objective of income poverty reduction. In this literature, as mentioned in the introduction, the tension between fighting against income poverty and increasing well-being (of the poor) is reflected, on which we comment first. Indeed, the tax-transfer systems that are proven optimal to alleviate income poverty distort the opportunities offered to poor (that is, typically, low-productivity) agents so as to incentivize them to earn more, but the loss in opportunities may be accompanied with a loss in well-being.[3]

The following simplified example illustrates this difficulty. Let us assume that two tax-transfer systems are possible. The poverty line is at 11, and the minimum wage is equal to 10. In the first policy, the poorest agent works half-time, thereby earning a pretax income of 5, obtains a transfer of 2, and earns a disposable income of 7. In the second policy, that agent works full time, earning 10, and obtains a transfer of 1 so that her disposable income is equal to 11. Finally, let us assume that the agent strictly prefers to work half-time and get 7 over working full time and getting 11. The first policy is the one that maximizes the well-being of the poor agent, whereas the second policy is the only one that allows the poor agent to reach the poverty line. This illustrates the tension between alleviating income poverty and increasing the well-being of the poor. Also note that in the example, the transfer needed to implement the first policy is higher than the transfer of the second policy, which means that the latter is less costly than the former in terms of well-being of the other, nonpoor, agents. The first policy might therefore not be feasible. The cost of redistribution is a key element in optimal income taxation theory, and it will also be a key element in this paper.

That a decrease in income poverty can be accompanied with a decrease in the well-being of the poor has been noted, for instance, by Kanbur, Keen, and Tuomala (1994a, b); Wane (2001); and, in the case of optimal linear income taxation, Kanbur and Keen (1989) as well as Kanbur et al. (2018). A similar difficulty also arises when poverty is defined in terms of commodity deprivation instead of lack of income (see Pirttilä and Tuomala 2004). Kanbur, Keen, and Tuomala (1994b) suggested a solution to that problem, namely to compare the poverty line not with the actual income of an agent but with her equivalent income. The equivalent income, a concept introduced by Samuelson (1974) and Samuelson and Swamy (1974), is the income level that given some fixed reference wage rate, would leave the agent indifferent to her actual situation. Kanbur, Keen, and Tuomala (1994b) do not give any result, however, and they limit themselves at mentioning the arbitrariness of the choice of the reference wage rate.

[3] Those distortions lead to opportunity sets that typically fail to satisfy the properties derived by optimal tax theory. Kanbur, Keen, and Tuomala (1994a), for instance, conclude that income tax rates should be negative on low incomes when all agents have a positive labor time (and there is no bunching at zero earnings), in contradiction to the classical results (summarized, for instance, in Diamond 1998). Similarly, Pirttilä and Tuomala (2004) conclude that commodity taxation should not be uniform even if preferences are separable in leisure and goods, against the classical theorem of Atkinson and Stiglitz (1976).

Our paper also connects to a recent trend in the optimal taxation literature, which tries to shift the focus from deriving the optimal tax formula to the underlying objectives of optimal taxation. On the empirical side, Weinzierl (2017, 2018) provides evidence that the goals for taxation that the public and policymakers endorse are poorly captured by the classical utilitarian social objective that has dominated the field since Mirrlees's (1971) seminal contribution. On the theoretical side and guided by the same intuition, many authors have studied optimal taxation with a social objective derived from specific redistributive or fairness principles as we do here (see Fleurbaey and Maniquet 2006, 2007; Lockwood and Weinzierl 2015; Saez and Stantcheva 2016, among others, and Fleurbaey and Maniquet 2018, for a survey).

Before we move on in the paper, we would like to also comment on the notion of well-being and preferences we adopt in this paper, and some other assumptions we make, in order to anticipate possible misunderstandings or objections to our approach and contribution that might arise more generally from related areas of research. The reader may skip the rest of this section without any loss of information to study the remaining sections of the paper.

First, as mentioned, the social preferences we derive in this paper are egalitarian in a specific well-being index representing the preferences of the agents. Namely, as we will outline in Section III, the well-being of an agent is computed as a decreasing function of the labor time that associated with a consumption level equal to the poverty line, leaves the agent indifferent to her actual labor time–consumption bundle. It is important to note that this well-being index does not belong to the family of equivalent incomes mentioned above and does not require the arbitrary choice of any reference price. However, it depends on the value of the poverty line, which we take as given. While the resulting objective of income poverty reduction turns out to be compatible with the respect of individual well-being, a current trend in the literature on the measurement of poverty consists of extending the definition of poverty to other dimensions of life than income. For instance, Decancq, Fleurbaey, and Maniquet (2019) propose an approach into the latter direction that is also consistent with the respect of individual preferences.

Second, having mentioned the dependence of our well-being index on the poverty line, one might argue that this is a drawback of our analysis, namely that it keeps silent about the level at which the poverty line should be fixed, pointing, for instance, to the fact that it will be endogenous to the distribution of incomes. In response, here are two reasons why to our view the poverty line should be assumed to be exogenous when one evaluates the ability of tax schemes to alleviate poverty. The first reason is normative. As we prove in the online Appendix, allowing the poverty line to depend on variables of the model, such as the distribution of incomes, would conflict with one of our key requirements, Pareto efficiency. The second reason has to do with the specific (redistributive) policies we are interested in evaluating. We do not believe, indeed, that when policymakers amend a tax system to address poverty, they take account of the influence of their reform on the value of the poverty line. In many countries, the poverty line is fixed at 60 percent of the median income. Poverty alleviation policies do not aim at affecting the median income so as to decrease poverty. This is what we capture in our model. To put it differently, we assess the ability of

tax policies to decrease poverty, given the way governments themselves define what it means to be poor and design the corresponding policies.

Third, some readers may find it questionable to define individual well-being in a way that is consistent with agents' preferences as they are *revealed* by their choices. The current literature on behavioral economics has made it clear that individual choices often fail to follow complete and transitive preferences, creating a gap between "experienced" and "decision" utility.[4] Living in precarious conditions, for instance, exacerbates the otherwise universal bias for the present (see Carvalho, Meier, and Wang 2016 for experimental evidence and Mullainathan and Shafir 2013 for general discussions and references). The immediate cost of the job search may deter poor individuals from actively looking for a job. As a result, one may refrain from developing the kind of welfare analysis that we advocate for here.

We need to distinguish between two radically different versions of this criticism. The first version of the criticism is based on the view that choices do not reveal what the individuals think is good for themselves because choices are constrained. As a result, revealed preferences need to be laundered from optimization constraints, and hence, the welfare analysis should be performed using laundered preferences. This is the approach followed by Chetty (2015) in which preferences are laundered from their present bias.

We need to stress that our analysis remains relevant even under this alternative approach, for two reasons. First, the first part of the paper is devoted to deriving social preferences by aggregating individual preferences in an efficient and fair way. This part assumes that the relevant individual preferences are known. Therefore, whether these individual preferences are revealed or laundered preferences does not matter for this exercise; it is well compatible with both. In other words, even if one considers that the behavior of poor people is biased in many ways and (in the first step) these biases should be removed from their actual preferences before preferences can be viewed as normatively compelling, one (in a second step) still needs a criterion to aggregate such corrected preferences, and our analysis provides such a criterion.

Second, when it comes to the maximization of the social preferences in the presence of individual optimization constraints, our approach does not take these constraints explicitly into account but yet allows us to derive some important conclusions. We need to be precise on this point as several types of optimization constraints have to be considered. If we again think of the present bias, or of other fixed costs of working, then, as we will explain in Section III below, our derivation of the social preferences remains valid. Moreover, when we derive the optimal tax scheme in Section V, we reach similar conclusions to the ones derived from the explicit account that a present bias may lower the efforts to find a job, like it is done in Lockwood (2020). Indeed, the maximization of our social preferences under incentive constraints also leads to imposing negative marginal tax rates on very low

[4] In the psychological and behavioral literature, decision utilities determine choices, while experienced utilities express the satisfaction once these choices are made. Following our assumption above, there is no difference between both types of utilities; i.e., choices are assumed to reveal what people truly prefer.

incomes. Thus, an explicit account of the present bias could only strengthen our results.

If we think of labor market frictions that may prevent individuals from supplying their preferred amount of labor, there are two ways in which our analysis can be applied to this case. First, we may consider that these frictions affect the *wage* of the individuals and not their preferences over labor time. We treat this case in online Appendix E.3. Second, we may consider that the laundered preferences of those who work part-time are more work oriented than their revealed preferences. Even if we do not deliver a full study of this case, we comment on the direction in which it affects our result in Section IV.

If we finally think of credit constraints that might prevent low-income individuals from optimizing, we have to admit that we don't see a way to (even partly) recover the true preferences from the revealed ones in that case. As a consequence, our analysis has to remain silent about it as well as about all other cases in which it is simply impossible to deduce true from observed preferences.

The second version of the criticism against the notion of well-being adopted in this paper is more radical than the one we have just addressed. It claims that even laundered preferences should not be respected in the policy design process. Let us call it the "paternalistic criticism." It is much more difficult to reconcile our work with this view. The typical motivation to fail to respect preferences of the poorest individuals is based on the view that working itself is (future) welfare improving and poor individuals do not take this relationship into account. Increasing the labor force participation and/or labor time of poor individuals is therefore a social improvement, independently of whether it increases their (instantaneous) experienced utility. In spite of the differences in the underlying normative stances, we can partially close the gap between our own starting point and the paternalistic view with the following two remarks. First, the taxation scheme that we end up defending consists in incentivizing poor individuals to participate in the labor force at least to the extent that their consumption reaches the poverty line. Our conclusions are therefore likely to coincide with the paternalistic ones on low incomes. Second, if the choice is between evaluating the ability of a tax system to alleviate poverty either by looking at the lowest incomes or by looking at the maximal labor time it takes to reach the poverty line (as we outline in Section IV), then everyone who is in favor of incentivizing poor individuals to work should also favor our criterion.

Finally, we would like to address the possible objection (from a classical optimal tax perspective) that in this paper, we are just defining one desirable property that the shape of the optimal tax scheme should satisfy (in Section V). Our main contribution, on the contrary, consists in providing (in Section IV) a simple criterion that we use (in Section VI) to compare *suboptimal* tax schemes and identify the one that should be preferred according to our social preferences among any set of available schemes. Following what we discussed in the previous paragraphs about the *optimal* tax scheme and the possible adjustments to our analysis to take behavioral biases into account, however, what we propose remains consistent with the theory if we rank tax schemes in two steps: first, prefer the one that exhibits negative marginal tax rates on low incomes, to provide individuals with the necessary nudge that will induce them to overcome behavioral biases, and second, for the remaining (and

main) part of the tax scheme, apply our criterion. What the present paper fundamentally teaches us is that the ability of a tax scheme to alleviate poverty should not be evaluated by looking at the distance between after-tax incomes of the low-income earners and the poverty line. It should be evaluated by the time it takes them to reach the poverty line.

II. The Model

There are two goods, labor, denoted ℓ, and consumption, denoted c. The population contains n agents. That means that there is a finite number of agents, but we think of that number as a large one so that the population is diverse in the relevant characteristics. A *bundle* for agent $i \in \{1, \ldots, n\}$ is a pair $z_i = (\ell_i, c_i)$, where ℓ_i is agent i's labor and c_i her good consumption. The agents' identical *consumption set* X is defined by the conditions $0 \leq \ell_i \leq 1$ and $c_i \geq 0$. We will refer to ℓ as labor time throughout the paper, but in Section III, devoted to defining social preferences, we could as well see it as effort or any variable for which it is normatively compelling to claim that redistribution should take place between two individuals who have the same ℓ. In Sections IV, V, and VI, however, we will have to restrict ourselves to interpreting ℓ as labor time because it is what will come out of the division of pretax income by hourly wages. In these sections, the level of consumption will be the after-tax income of the agent.

Agents have two characteristics, their preferences over the consumption set and their productivity. For any agent $i \in \{1, \ldots, n\}$, *preferences* are denoted R_i, and $z_i R_i z_i'$ (respectively $z_i P_i z_i'$, $z_i I_i z_i'$) means that bundle z_i is weakly preferred (respectively strictly preferred, indifferent) to bundle z_i'. We assume that individual preferences are continuous, convex, and monotonic.[5] We further assume that consumption is necessary, in the sense that any bundle with a positive good consumption is always strictly preferred to any bundle with a zero consumption. This assumption will play a role during the construction of the social preferences. We let $\mathcal{R}$ denote this set of preferences.

We allow preferences to differ across agents. Moreover, the heterogeneity is arbitrary, and our results hold true whatever the distribution of preferences in the population.

The marginal productivity of labor is assumed to be fixed as in a constant returns to scale technology. Agent i's earning ability is measured by her productivity or *wage rate*, denoted w_i, and is measured in consumption units so that $w_i > 0$ is agent i's production when working $\ell_i = 1$ and, for any ℓ_i, $w_i \ell_i$ is the agent's pretax income (earnings). We assume that $w_i \in [w_{\min}, \infty)$, where $w_{\min} > 0$ stands for the minimum wage rate. We also assume that there are agents in the economy whose wage equals the minimum wage: $w_i = w_m$ for some $i \in \{1, \ldots, n\}$. Our notion of a minimum wage rate can either refer to a legal minimum wage or to the minimal

[5] Preferences are monotonic if $\ell_i \leq \ell_i'$ and $c_i > c_i'$ implies that $(\ell_i, c_i) P_i (\ell_i', c_i')$. Observe that these assumptions assume away preferences being only interested in minimizing labor time. That is a condition under which the property of *Poverty Reduction*, below, is well defined.

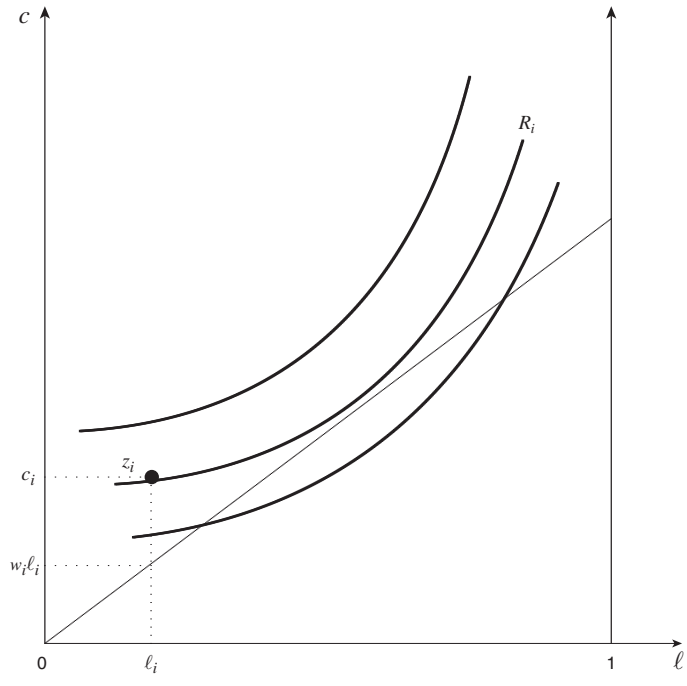

FIGURE 1. ILLUSTRATION OF THE MODEL

statistical wage rate, provided it is observed. In our application in Section VI, we will stick to the legal minimum wage.

The minimum wage being strictly positive is a necessary assumption for all our results. This assumption does not rule out that some people or households in the economy can simply not earn incomes (such as handicapped, homeless, or long-term unemployed people), but the assumption amounts to state that these people can be identified and treated separately by the redistribution system (given that unemployment and minimum income benefits are typically conditional on looking for jobs, we can state that even the long-term unemployed can be identified).

Figure 1 displays the consumption set, with typical indifference curves, and earnings as a function of labor time. As illustrated in the figure, an agent's consumption c_i may differ from her earnings $w_i \ell_i$. This is a typical consequence of redistribution.

An allocation is a list $z = (z_1, \ldots, z_n)$. Social preferences will allow us to compare allocations in terms of fairness principles and efficiency.[6] Social preferences will be formalized as a complete ordering over all allocations in X^n and will be denoted R, with asymmetric and symmetric components P and I, respectively. In other words, zRz' (respectively zPz', zIz') means that z is at least as good as z' (respectively strictly better, equivalent).

[6] We use "fairness" in the classical sense of the theory of fair allocation, according to which economic justice is a matter of resource allocation as opposed to utility level allocation. The poverty-reduction property we define below is consistent with this view.

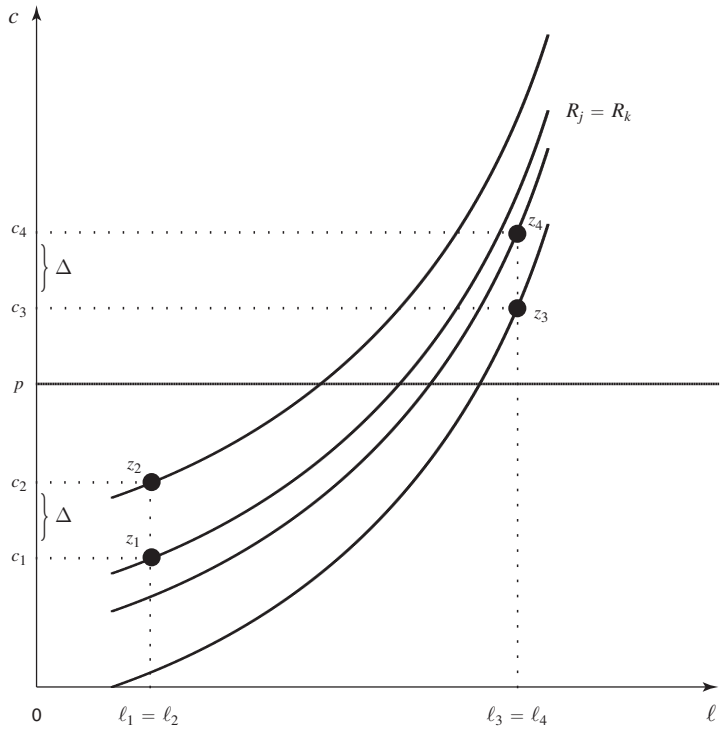

FIGURE 2. INDIVIDUAL WELL-BEING VERSUS THE OBJECTIVE OF CONSUMPTION POVERTY REDUCTION

Social preferences may depend on the population profile of characteristics $(R_1, \ldots, R_n)$ and $(w_1, \ldots, w_n)$. However, we may simplify the analysis and consider that the profile of wage rates is fixed. As a result, social preferences R are formally a *mapping* from the set of population profiles $\mathcal{R}^n$ to the set of complete orderings over allocations.

III. Social Preferences

We assume that there is a poverty line, that is, a consumption level, p, with the property that society considers it unacceptable to let people live with less consumption than p. This poverty line is fixed. As mentioned in the introduction, we study the consequences of allowing the poverty line to depend on the distribution of (after-tax) incomes in the online Appendix.

We now treat the discussion carried out in the introduction more formally. There are two pitfalls that the objective of poverty reduction should avoid. They are illustrated in Figure 2. There are four bundles, z_1, z_2, z_3, and z_4, and two agents, j and k, having identical preferences $R_j = R_k$.

First, let us compare z_2 and z_3. We have $c_2 < p < c_3$ so that we could claim that assigning z_3 to j or k is socially preferable to assigning them z_2. We also have that $z_2 P_j z_3$ (and $z_2 P_k z_3$), though, so a social preference for z_3 over z_2 would immediately conflict with individual well-being and, therefore, the Pareto criterion.

Second, let us consider the allocation at which j consumes z_1 and k consumes z_4. Again, we might consider that j is poor because $c_1 < p$, whereas k is nonpoor because $c_4 > p$. A transfer of income from k to j could then be thought of as a social improvement, seeming in line with the celebrated Pigou-Dalton progressive transfers in the literature on inequality measurement (see Pigou 1912 and Dalton 1920). Contrary to the first pitfall, such a transfer would not conflict with Pareto efficiency as the Pareto criterion is silent about the ranking of allocations (z_1, z_4) and (z_2, z_3).

The transfer is not desirable, though, because it amounts to exacerbate inequality in well-being rather than decrease it. Indeed, in well-being terms, agent k at z_4 is worse off than agent j at z_1, in the sense that z_4 lies on a lower indifference curve than z_1 (this statement only makes sense because the agents both have the same preferences) and well-being inequality is larger at (z_2, z_3) than at (z_1, z_4).

To avoid these two pitfalls, we look for social preferences that satisfy both the Pareto criterion and the following poverty reduction property: a transfer from a rich to a poor is considered to be a social improvement only under the proviso that the labor time of the two agents is the same. Note that this proviso guarantees that both agents agree that the bundle assigned to the agent whose consumption level is above the poverty line is better than the bundle assigned to the agent whose consumption level is below the poverty line. This requirement is illustrated in Figure 3. We state it formally in Property 1.

PROPERTY 1 (Poverty Reduction): *For all economies* $(R_1, \ldots, R_n)$, *for all pairs of allocations* $z = (z_1, \ldots, z_n)$ *and* $z' = (z'_1, \ldots, z'_n)$, *if, for two agents j and k and a positive quantity* Δ,

$$\ell_j = \ell'_j = \ell_k = \ell'_k, \tag{1}$$

$$c'_j = c_j + \Delta \leq p \leq c'_k = c_k - \Delta, \tag{2}$$

whereas $z_i = z'_i$ *for all other agents, then* z' *is socially strictly preferred to* z.[7]

Here are three examples of social preferences that satisfy *Poverty Reduction*. The first example is the social preference relation based on the celebrated income poverty measurement introduced by Foster, Greer, and Thorbecke (1984). We denote that social preference by R^{FGT}. It is formally defined as follows: for all economies $(R_1, \ldots, R_n)$, for all pairs of allocations $z = (z_1, \ldots, z_n)$ and $z' = (z'_1, \ldots, z'_n)$, z is socially weakly preferred to z' if and only if there is less poverty at z than at z', that is,

$$\frac{1}{n}\sum_{i=1}^{n}\left(\frac{\max\{0, p - c_i\}}{p}\right)^{\alpha} \leq \frac{1}{n}\sum_{i=1}^{n}\left(\frac{\max\{0, p - c'_i\}}{p}\right)^{\alpha}.$$

[7] In equation (2), agent j, who benefits from the transfer, does not jump over the poverty line as a consequence of the transfer. A stronger version of the property could allow for such a jump by letting $c'_j \leq c'_k$, $p \leq c'_k$ but not $c'_j \leq p$. All the results below would remain true.

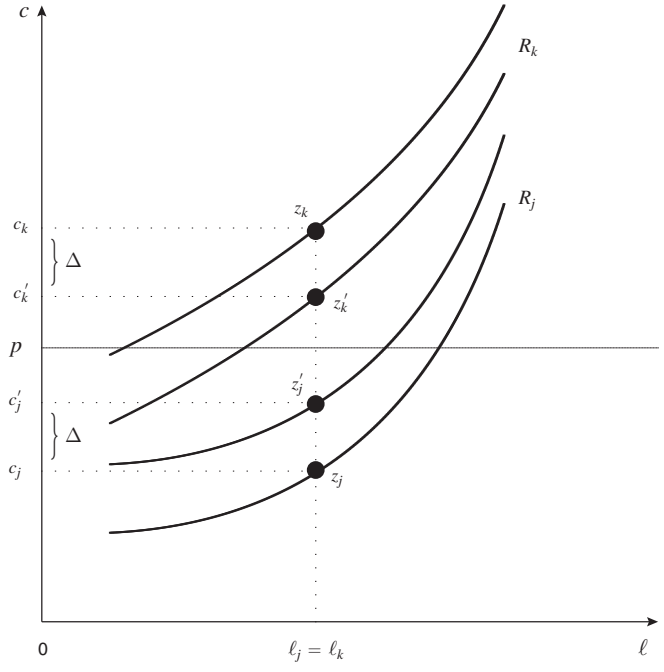

FIGURE 3. *POVERTY REDUCTION*: (z_j', z_k') IS SOCIALLY PREFERRED TO (z_j, z_k)

Social poverty is defined as the sum of individual poverty. Any agent consuming more than the poverty line has a poverty level equal to zero. An agent consuming nothing has a poverty level of one. The α coefficient stands for the degree of inequality aversion among the poor, which amounts to the priority that is given to people at the very bottom of the consumption spectrum. If $\alpha = 0$, all poor agents contribute the same level to global poverty. It amounts to give priority to agents very close to the poverty line; i.e., giving \$1 to an agent so close to the poverty line that this agent quits poverty as a consequence of the transfer decreases poverty more than giving \$1 to a very poor agent. This is the so-called headcount ratio. If $\alpha = 1$, transferring \$1 to a poor agent decreases poverty the same way independently of the consumption level of this agent. This is the so-called poverty-gap ratio, measuring the average share of p that needs to be transferred to poor agents to completely alleviate poverty. If $\alpha > 1$, transferring \$1 to a poor agent decreases poverty more the poorer this agent is. As soon as $\alpha > 0$, all social preferences R^{FGT} satisfy *Poverty Reduction*.

Another example of social preferences that satisfy this property is the generalized utilitarian social welfare function that is often used in optimal taxation theory under the assumption that preferences are quasi-linear (see, for instance, Diamond 1998). We denote that social preference by R^U. It is formally defined as follows: for all economies $(R_1, \ldots, R_n)$ in which all agents have (possibly different) quasi-linear preferences represented by quasi-linear utility functions $(u_1, \ldots, u_n)$, for all pairs of

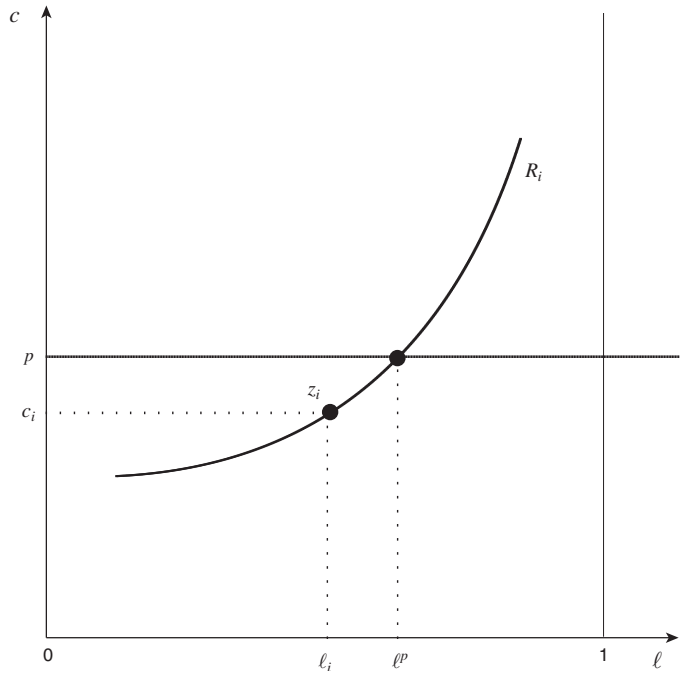

FIGURE 4. ILLUSTRATION OF THE WELL-BEING MEASURE W^p: $W^p(z_i, R_i) = -\ell^p$

allocations $z = (z_1, \ldots, z_n)$ and $z' = (z'_1, \ldots, z'_n)$, z is socially weakly preferred to z' if and only if

$$\sum_{i=1}^{n} g\big(u_i(z_i)\big) \geq \sum_{i=1}^{n} g\big(u_i(z'_i)\big),$$

where g is a strictly concave and strictly increasing real-valued function representing the inequality aversion of the planner. Let us observe that when $\ell_j = \ell'_j = \ell_k = \ell'_k$ and $c'_j = c_j + \Delta < c'_k = c_k - \Delta$, quasi-linearity implies that $u_j(z'_j) - u_j(z_j) = \Delta = u_k(z_k) - u_k(z'_k)$. The fact that R^U satisfies *Poverty Reduction* comes then from the strict concavity of g.

A last example is a new social preference relation that we introduce in this paper. It works by applying the leximin aggregator to some particular well-being representation of individual preferences. We denote that social preference by R^{lex}. The leximin aggregator works by lexicographically applying the maximin aggregator: first maximize the well-being of the worst off, in case of a tie, maximize the well-being of the second worst off, and so on. The new representation of the preferences works as follows. It is a decreasing function of the labor time that leaves an agent indifferent between her current bundle and consuming the poverty line p at that labor time. We denote that labor time ℓ^p. Formally, the well-being of an agent having preferences R_i and consuming bundle $z_i = (\ell_i, c_i)$, denoted $W^p(z_i, R_i)$, is equal to $-\ell^p$ if this agent is indifferent between z_i and (ℓ^p, p). It is illustrated in Figure 4. Note that given this construction, $\ell^p \in [0, 1]$, but the only relevant characteristic of $W^p(z_i, R_i)$

is to be decreasing in labor time. How decreasing does not matter because we apply the leximin aggregator, which only uses the ordinal information on well-being levels.

There are two cases in which no ℓ^p satisfying the indifference condition above exists. The first case is when $z_i P_i (0,p)$, that is, the agent strictly prefers her bundle over not working at all and consuming the poverty line. In this case, there is a consumption level c such that $z_i I_i (0,c)$. We fix the well-being at such a bundle equal to $c-p$.

The second case is when $(1,p) P_i z_i$, that is, the current bundle of this agent is so bad that she would prefer to work full time and consume exactly the poverty line. In this case, there is a consumption level c such that $z_i I_i (1,c)$. We fix the well-being at such a bundle equal to $c-p-1$.

We now check which social preferences satisfy the other properties we are interested in. Our first property is the classical Pareto property, which we now formally define. It guarantees that decreasing the preference satisfaction of an agent will never be a social improvement, even if this decrease goes together with an increase in income above the poverty line.

PROPERTY 2 (Pareto): *For all economies* $(R_1,\dots,R_n)$, *for all pairs of allocations* $z=(z_1,\dots,z_n)$ *and* $z'=(z'_1,\dots,z'_n)$, *if all agents* i *weakly prefer* z_i *to* z'_i, *then* z *is socially weakly preferred to* z'. *If, moreover,* z_j *is strictly preferred to* z'_j *for one agent* j, *then* z *is socially strictly preferred to* z'.

Note, R^{FGT} does not satisfy *Pareto*. It should not be a surprise because R^{FGT} only aggregates consumption levels and remains insensitive to increases in labor times.

On the other hand, R^U satisfies *Pareto*, which follows from the fact that these social preferences are directly defined as a function of the utility levels of the agents. These preferences R^U are defined for economies in which preferences are quasi-linear. We would like to be more general than that and be able to define social preferences even when there are income effects. We need to generalize R^U to all preferences. Yet, we would like social preferences in non-quasi-linear economies to be consistent with that in quasi-linear ones. We capture this requirement with the following property. It requires that social preferences be independent to changes in preferences that do not affect the indifference curves through the bundles we are contemplating. This is a cross-economy robustness property, and it explains why we wanted to define the domain of economies as a function of all possible preference profiles in the economy. This is a way to make our conclusions independent of the precise profile of the preferences in the economy, for instance of the fact that all agents have quasi-linear preferences.

There are many ways in which such independence can be justified. It is a weakening of Arrow's independence property and a weakening that makes it compatible with fairness properties as Samuelson (1977) and Pazner (1979) already mentioned. This independence property is also related to incentive compatibility. We postpone that discussion to the next section.

To define that property formally, we need the following terminology. For some preference relation $R_i \in \mathcal{R}$ and some bundle z_i, we let $I(z_i,R_i)$ denote the indifference curve at z_i, that is, the set of all bundles to which this agent is indifferent.

PROPERTY 3 (Independence): *For all pairs of economies* $(R_1, \ldots, R_n)$ *and* $(R'_1, \ldots, R'_n)$ *and all pairs of allocations* $z = (z_1, \ldots, z_n)$ *and* $z' = (z'_1, \ldots, z'_n)$, *if all agents have the same indifference curves at* z_i *and* z'_i *with preferences* R_i *and* R'_i, *that is, if* $I(z_i, R_i) = I(z_i, R'_i)$ *and* $I(z'_i, R_i) = I(z'_i, R'_i)$ *for all* $i \in \{1, \ldots, n\}$, *then the social preference toward* z *and* z' *is the same in both economies.*

Note, R^{FGT} satisfies this property for all values of α. This immediately comes from the fact that R^{FGT} is defined without reference to preferences. It is therefore independent to all changes in preferences, including the ones that leave indifference curves unaffected. Unfortunately, no social preference can generalize R^U to non-quasi-linear economies so as to satisfy that property. This will come as a corollary of our first result. The social preference relation that we have introduced above, R^{lex}, satisfies all three properties. Moreover, any social preference that satisfies all three properties needs to maximin the well-being measure W^p. That almost amounts to say that R^{lex} is the only social preference that satisfies the three properties.

PROPOSITION 1: *Social preferences* R^{lex} *satisfy Poverty Reduction*, *Pareto*, *and Independence. Conversely*, *if social preferences* R *satisfy Poverty Reduction*, *Pareto, and Independence*, *then for all economies* $(R_1, \ldots, R_n)$, *for all pairs of allocations* $z = (z_1, \ldots, z_n)$ *and* $z' = (z'_1, \ldots, z'_n)$ *such that*

$$\min_i W^p(z_i, R_i) \in [0, 1] \quad \textit{and} \quad \min_i W^p(z'_i, R_i) \in [0, 1],$$

if

$$\min_i W^p(z_i, R_i) > \min_i W^p(z'_i, R_i),$$

allocation z *is socially strictly preferred to* z'.

This proposition only provides us with a partial social ranking of the allocations because it requires that the indifference curve of the worst-off individual crosses the poverty line (the well-being level according to W^p is between zero and one). However, even if it is not stated in the proposition, the three axioms imply more than that. For instance, any allocation that Pareto dominates another one is also socially preferred. As a result, any allocation in which all individuals have their entire indifference curves above the poverty line is socially preferred to any allocation in which even only one agent's indifference curve crosses the poverty line while those of all other agents lie above.

The formal proof of the proposition is similar to the proof of the main result in Maniquet and Sprumont (2004) in a public-good model.[8] We provide a complete proof in the online Appendix.

Two aspects of Proposition 1 are unexpected. First, it tells us that we need to maximize the minimal well-being level; that is, we need to apply an infinite degree

[8] The main differences between the two models are that labor time is bounded above and consumption is bounded below in the model of this paper.

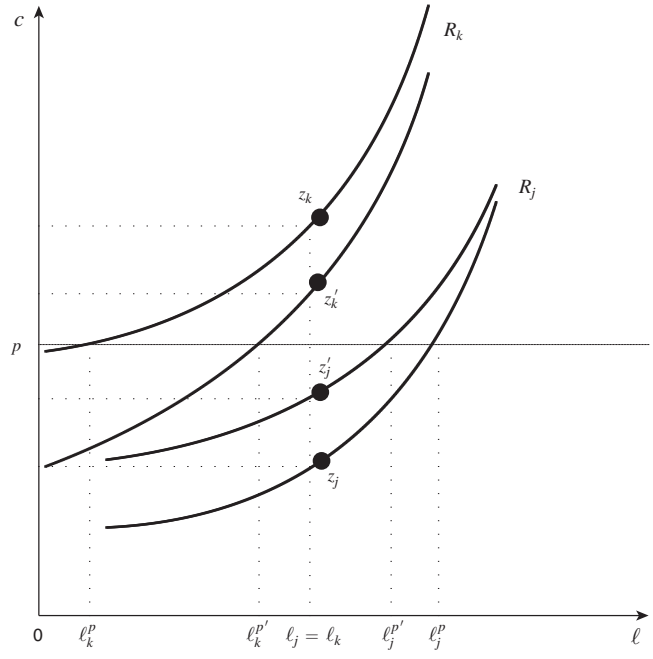

FIGURE 5. ILLUSTRATION OF WHY SOCIAL PREFERENCES NEED TO BE MAXIMIN

of inequality aversion, whereas the property of *Poverty Reduction* is nothing more than a transfer principle, which is compatible with any degree of inequality aversion. The intuition of this result is not difficult to grasp, though. Let us look at Figure 5, a variant of Figure 3. The transfer from k to j needs to be declared a social improvement. In terms of well-being, however (that is, in terms of the labor time it takes to reach the poverty line along those indifference curves), we see that receiving $z'_j - z_j$ increases agent j's well-being by $\ell_j^p - \ell_j^{p'}$, whereas losing $z_k - z'_k$ decreases agent k's well-being by $\ell_k^{p'} - \ell_k^p$, a much larger amount. One can easily imagine that the well-being increase for agent j is arbitrarily small, whereas the corresponding decrease for agent k is arbitrarily large. An arbitrarily small increase of the poorer agent has to compensate an arbitrarily large decrease for the relatively richer one. This can only be achieved with a maximin objective.[9]

Second, Proposition 1 tells us to measure individual well-being according to W^p, which is expressed in labor time, whereas *Poverty Reduction* is a property of money transfer. Again, the intuition can be deduced from Figure 5. In order to satisfy *Poverty Reduction* with egalitarian social preferences, one needs to be sure that agent j is considered poorer (that is, at a lower well-being level) at z_j than k at z_k, whatever their preferences R_j and R_k. As suggested in the picture, the indifference curves of these two agents may cross each other. Moreover, they may cross at the

[9] Maximin results are frequent in the literature on social preferences based on fairness properties. See Fleurbaey and Maniquet (2011) for a detailed presentation.

vertical of any point either left of $\ell_j = \ell_k$ or right. Consequently, the only certainty one can have about the indifference curves of agent j at z_j and of agent k at z_k is that the former crosses the poverty line at a larger labor time than $\ell_j = \ell_k$, whereas the latter crosses it at a smaller labor time. This is why well-being is measured in labor time along the poverty line, and agent k is claimed to have a larger well-being at z_k than agent j at z_j.

We close this section with a more technical remark. Contrary to most of the literature on optimal taxation, we assume in this paper that agents' preferences may differ. This assumption does not make our result more difficult to derive. Quite the contrary, it helps. Indeed, the proof of the theorem requires that the domain of preferences be sufficiently rich. As a result, on the one hand, the theorem would not hold true under the assumption that all agents have the same preferences. It would even not hold true under the assumption that all agents have quasi-linear preferences, or any other restriction on preferences. On the other hand, the theorem would remain valid even if we extend the domain of admissible preferences. One such extension is particularly relevant given what we discussed in the introduction. Indeed, a reduced-form model of present bias, or all other elements that amount to fixed costs of working, may assume that preferences are concave over bundles containing very low labor time. The proof of the theorem remains valid if preferences are allowed to be concave, and so would the second-best result of the next section. The only result that would need to be revised is the first-best result of Section V. Given that we consider it less important, we will not explore how it would generalize to nonconvex preferences. The main message we are left with is then that taking fixed costs of working into account will not affect the criterion that we derive in the next section.

IV. Second Best

In the previous section, the objective was to construct a complete ranking of allocations, independently of the information that was needed to rank them. In this section, we introduce the information constraints facing the redistribution designer; that is, we switch to the classical second-best context, whose formalism dates back to Mirrlees (1971). We assume that only earned income $y_i = w_i \ell_i$ is observed so that redistribution is made via a tax function $\tau(y_i)$.[10] This tax is a subsidy when $\tau(y_i) < 0$. In this context, meeting the incentive compatibility constraints is equivalent to letting agents choose their labor time in the budget set modified by the tax schedule. Each agent i chooses ℓ_i knowing that it yields a consumption level $w_i \ell_i - \tau(w_i \ell_i)$.

A tax function τ is feasible only if it meets the incentive compatibility constraints and it balances the budget, that is,

$$\sum_{i \in N} \tau(y_i) \geq B,$$

[10] See, e.g., Stiglitz (1987, 1002–04) or Boadway and Keen (2000, 737–38) for simple presentations of this second-best setting.

where B stands for the public expenditures that are not related to the redistribution system. In this section, we simply assume that the tax functions we have to rank are feasible. Of course, in practice, identifying which tax functions are feasible is key, but it is precisely the question on which the literature on optimal taxation has been more prolific, and many formulas exist to check whether tax functions are feasible (see, for instance, Boadway 2012 for an overview on the numerous insights on this question). This paper does not contribute to the literature on identifying which tax functions are feasible, but we rather build on it. Note that whether the selection among feasible tax functions follows some fairness, welfarist, or other criteria does of course not matter at the stage of identifying what is feasible but only once feasibility is given. This paper precisely contributes to this second stage.

If feasible tax functions could guarantee that all agents are lifted out of poverty, that is if $y - \tau(y) \geq p$ for all y, it is clear that the optimal tax functions should be found among them. In most economies, though, such tax functions are not feasible, for budgetary or sometimes political reasons. This is why a criterion needs to be developed to evaluate tax functions that do not succeed in bringing everybody above the poverty line. We derive this criterion now.

Given the information available to the tax function designer, the relevant space becomes the earnings-consumption space. Individual preferences in that space are denoted R_i^*, and they are derived from ordinary preferences over labor time–consumption bundles by

$$(y,c)\, R_i^*\, (y',c') \;\Leftrightarrow\; \left(\frac{y}{w_i},c\right) R_i \left(\frac{y'}{w_i},c'\right).$$

It is a notably difficult task to characterize the optimal tax function when agents differ in wage rates and in preferences (see Jacquet and Lehmann 2017 for a recent solution to that problem). We escape the need of such heavy derivations here thanks to the leximin nature of our objective. Indeed, it is sufficient for our purpose to deduce which agent in the population has the lowest well-being index W^p, or, more precisely, it is sufficient to identify the earning level chosen by the agent with the lowest index.[11]

We derive our main result under the following assumption, which we discuss and justify at the end of the section. It requires that whatever the tax function we are contemplating, there is at least one agent with the minimum wage rate who chooses the pretax income level that is just sufficient to reach a consumption level equal to the poverty line, provided such a pretax income level exists. If no such pretax income level exists, because after-tax incomes are too low, then there is at least one minimum wage agent working full time.

ASSUMPTION 1: *For each τ, if there exists $y^p \leq w_m$ such that $y^p - \tau(y^p) = p$, then there exists $i \in \{1,\ldots,n\}$ such that $w_i = w_m$ and $y_i = y^p$. If $y - \tau(y) < p$ for all $y \leq w_m$, then there exists $i \in \{1,\ldots,n\}$ such that $y_i = w_m$.*

[11] A similar simplification due to the leximin nature of the social objective is used in Fleurbaey and Maniquet (2006).

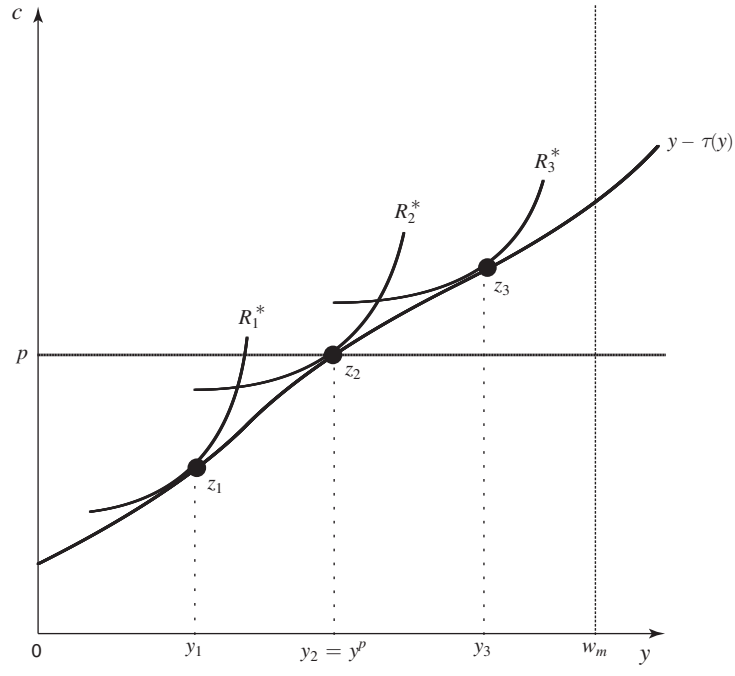

FIGURE 6. SECOND BEST: WHO IS THE WORST OFF?

The argument leading to the result is illustrated in Figure 6. Let us assume that τ is such that there exists $y^p \leq w_m$ such that

$$y^p - \tau(y^p) = p. \tag{3}$$

First, let us restrict our attention to minimum wage agents, that is, agents with $w_i = w_m$. In the figure, if $w_2 = w_m$, then $W^p(z_2, R_2) = -y^p/w_m$. This type of agent always exists, by Assumption 1. The picture describes the entire indifference curve through bundle (y^p, p), but we need not know the entire shape of the curve. Knowing that this agent has chosen y^p is enough to compute $W^p(z_2, R_2)$. Also, if $w_1 = w_m$, $W^p(z_1, R_1) > -y^p/w_m$. Indeed, choosing an earning level, say y_1, below y^p reveals a preference for $(y_1, y_1 - \tau(y_1))$ over (y^p, p). If $w_3 = w_m$, $W^p(z_3, R_3) > -y^p/w_m$ as well by the same revealed preference argument. That proves that any minimum wage agent choosing an earning level different from y^p has a higher $W^p(z_i, R_i)$ than agent 2. This agent has the lowest W^p-index among the minimum wage agents.

Second, let us consider agents with $w_i > w_m$. If $w_2 > w_m$, for instance, then agent 2 would need to work less to earn y^p and consume p than with a wage equal to w_m. More specifically, $W^p(z_2, R_2) = -y^p/w_2 > -y^p/w_m$. Again, applying the same revealed preference argument as above, if $w_1 > w_m$, $W^p(z_1, R_1) > -y^p/w_1 > -y^p/w_m$, and if $w_3 > w_m$, $W^p(z_3, R_3) > -y^p/w_3 > -y^p/w_m$. That proves that the lowest $W^p(z_i, R_i)$ is the one of the minimum wage agent earning y^p. We denote this well-being index $W^p_{\min}$.

This reasoning also illustrates that our social preferences do not favor "lazy" agents, which one might expect at first glance. Indeed, any agent like agent 1 in the figure is more averse to work than agent 2, but they have a higher level of well-being so that agent 2 is given priority by the leximin preferences. The social preferences do not favor hardworking agents, either. Indeed, any agent like agent 3 in the figure is less averse to work than agent 2, but they have also a higher level of well-being. The only individual preferences that can be claimed to be favored by our social preferences are those of an agent who would always choose the labor time that yields a consumption level equal to the poverty line.

Let us assume, now, that no y^p satisfying equation (3) exists. Then, either $-\tau(0) > p$ or $w_m - \tau(w_m) < p$. In the former case, the well-being index of an agent who does not work is $-\tau(0) - p > 0$. Agents having a strictly positive earning level reveal strict preference of their bundle over $(0, -\tau(0))$ so that their well-being index satisfies $W^p(z_i, R_i) > -\tau(0) - p$. This means that the lowest well-being index has a strictly positive value so that this allocation is socially preferred to any allocation in which a y^p satisfying equation (3) exists. In the latter case, because Assumption 1 guarantees the existence of some agent i such that $w_i = w_m$ and $y_i = w_m$, we have $W^p(z_i, R_i) = w_m - \tau(w_m) - p - 1 < -1$; that is, the lowest well-being index has a strictly lower value than -1 so that this allocation is socially worse than any allocation in which a y^p satisfying equation (3) exists.

We can summarize the result in the following proposition.

PROPOSITION 2: *Let τ and τ' be two tax functions. Under Assumption 1, the allocation generated by τ is socially better than that generated by τ' if one of the following conditions holds:*

(*i*) *all agents have an after-tax income above the poverty line in the allocation generated by τ but not by τ',*

(*ii*) *the earning level that is just necessary to obtain an after-tax income equal to the poverty line is lower than the minimum wage and is lower in the allocation generated by τ than by τ',*

(*iii*) *the earning level that is just necessary to obtain an after-tax income equal to the poverty line is lower than the minimum wage in the allocation generated by τ but larger in the one generated by τ'.*

Proposition 2 is illustrated in Figure 7. Four stylized budget curves are drawn in the earnings-consumption space. Tax function τ^1 yields an allocation that is socially better than all the others, according to condition (i) in the proposition. Condition (ii) amounts to claiming that the allocation generated by τ^2 is socially better than the one generated by τ^3. Condition (iii) amounts to claiming that τ^4 is the worst tax function among the four represented in the figure.

We have proven that the evaluation criterion for tax functions is the pretax income that is just necessary to reach a consumption level equal to the poverty line because it reveals the labor time it takes for the minimum wage agents to reach that level.

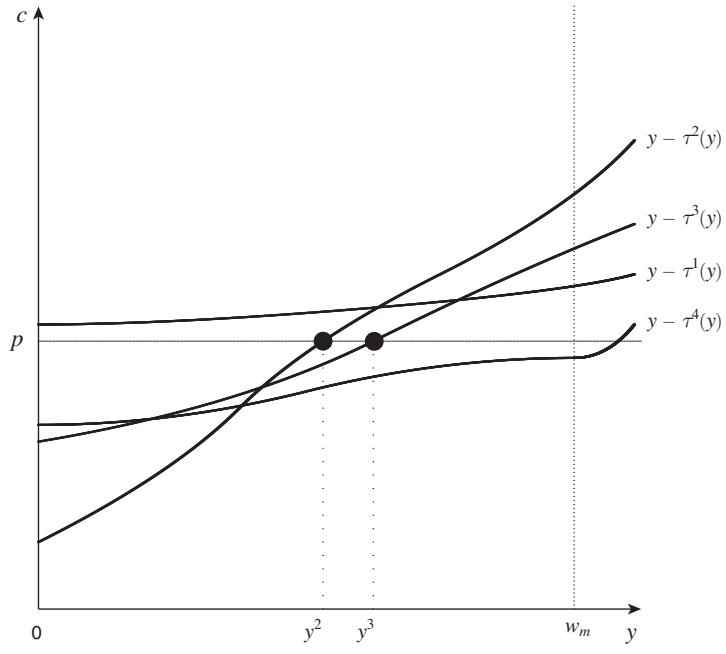

FIGURE 7. ILLUSTRATION OF PROPOSITION 2

Consequently, a reform of a tax function aiming at increasing social welfare should decrease the amount of taxes (or typically, increase the amount of subsidies) at this precise pretax income level. In order to apply the criterion to the evaluation of existing tax schemes, the only statistic that we need to measure is the pretax level that is necessary to reach a consumption level equal to the poverty line. That means that nothing needs to be estimated in terms of individual responses to changes in the tax scheme. This simplicity is quite striking when compared with the typical estimations of behavioral parameters that are needed to apply optimal tax formulas derived from the maximization of general social preferences (see, for instance, Chetty 2009, Saez and Stantcheva 2016, Jacquet and Lehmann 2017).

We would like to underline two properties of the result above that might not be transparent at first glance. The first property is that the result is derived based on the assumption that agents can choose any bundle in their budget set. In terms of the typical algebraic arguments of optimal behavior, this means that agents can adjust their labor supply both at the intensive margin, by changing their labor time marginally, or at the extensive margin, by jumping between zero and strictly positive levels of labor supply. That is, we do not need to restrict the set of available choices in any way.

The second property is that the reasoning remains valid in presence of tagging. A tag is an observable characteristic that may enter the tax function so that the amount of tax paid by the agents may depend on the value of their tag (see Akerlof 1978 for a seminal treatment of the use of tagging in optimal tax theory). Let us assume that a tag allows us to divide the population in two subpopulations so that a tax scheme is composed of two tax functions. Now, to take the example that we have to

compare two tax schemes, one composed of tax functions τ^1 and τ^2, one applied to each subpopulation, and the other one composed of tax functions τ^1 and τ^3. The ranking between these two tax schemes follows from exactly the same reasoning as before. The worst-off agents with the first tax scheme will be the minimum wage agents in the second subpopulation choosing to earn y^2, whereas the worst-off agents with the second tax scheme will be the minimum wage agents in the second subpopulation choosing to earn y^3. The former tax scheme is, therefore, preferable.

We conclude the second-best analysis of this section with three comments, one on our third property, *Independence*, one on Assumption 1, and one on the possibility to take market frictions into account. First, the well-being index that Proposition 1 teaches us to use gives us a numerical representation of preferences that only depends on the indifference curve through the bundle we are contemplating. This immediately follows from *Independence*, which prevents social judgments about bundle z_i to depend on preference information outside that indifference curve. This is the reason why we succeed in deriving Proposition 2. Indeed, the only fact that one agent, agent 2 in the reasoning above, chooses to earn y^p reveals enough of her well-being index to help us conclude that she is the worst off. With social preferences violating *Independence*, we might have had first to check for preferences of agent 2 over bundles not contained in her indifference curve through her chosen bundle, and such a check would have typically been impossible to do because we cannot collect more information than agent 2's best bundle in her budget set. It would not have been possible to derive a second-best criterion corresponding to social preferences requiring more information than indifference curves through the contemplated bundles. The surprise, actually, comes from the fact that *it is* possible to derive a criterion from R^{lex}. Indeed, agent 2's choice does not even reveal her entire indifference curve. Fortunately, the local information it reveals is sufficient to make the criterion work.

Second, Proposition 2 is derived under an unusual assumption that bears on both the types of the agents and the shape of the tax function. We discuss and justify it now. Let us assume that contrary to the assumption, no minimum wage agent chooses earning level y^p. There can be two cases. Either no agent chooses y^p or only higher-wage agents choose it.

In the first case, the part of the budget curve around y^p is irrelevant. That means that τ can be adjusted, decreased, until it becomes relevant, that is, until some agents find it interesting to earn pretax incomes around y^p. This exercise can actually be done until all points of τ are relevant, that is, until it coincides with the lower envelope of agents' indifference curves through their chosen bundles. After τ is adjusted, either the assumption is satisfied or we have reached the second case.

That is, let us assume that all agents earning y^p have a higher wage than w_m. In theory, it can be justified to increase the tax at y^p to collect more money and redistribute it toward worse-off agents. In practice, though, this is irrelevant because it amounts to claim that the tax rate at some $y < w_m$ should be computed by the fiscal authority on the basis that this authority is certain that all agents earning that amount are high-skill agents. It does not sound like a plausible justification because the typical information a fiscal authority has about the distribution of wages and earnings does not allow it to exclude the possibility that agents with a given wage work a given fraction of their time. We view Assumption 1, actually, as reflecting the

imprecision of the available information, with the consequence that the tax designer cannot exclude the existence of an agent having the characteristics mentioned in the formal statement above.

Third, we commented in the introduction on the possibility that revealed preferences do not perfectly reflect what individuals find good for themselves because their choices can be constrained. This affects Proposition 2 in the following way. When individuals work less than what they would like to work, which is mostly the case among individuals working part-time, then their true preferences favor working more than their revealed preferences. In Figure 6 above, if individual 1 is constrained, her true indifference curve through z_1 may actually cross the poverty line right of y_2 so that individual 1 may actually be the worst off given our criterion. This cannot be the case of individual 3. We do not fully study this possibility here as it would require to have a precise theory of the relationship between revealed and true preferences, but what we can say is that the new criterion would require that we contemplate the tax scheme not only at the point at which it crosses the poverty line *but also below*. How far below the poverty line we should extend the domain of interest would depend on how different true preferences can be from revealed preferences.

V. Optimal Income Taxation

Before we apply the criterion defined in Proposition 2, we briefly study in this section the tax scheme that is optimal given our social objective. Finding the formula of the optimal tax has been the main focus of optimal taxation theory (see the surveys in Boadway 2012 or Kaplow 2008). It is well known that when agents are heterogeneous in both wages and preferences, it is impossible to give the general formula of the tax scheme that maximizes social preferences without making strong assumptions on the distribution of types. In this section, we do not look for the optimal tax formula, but we derive a property of it that does not depend on any strong assumption on the distribution of types.[12] This property is that marginal tax rates need to be negative on average for earnings yielding consumption levels below the poverty line. It will allow us to contrast the shape of our optimal tax scheme with the shape of two optimal tax schemes recently derived by Saez and Stantcheva (2016) from maximizing criteria embedding poverty alleviation also being compatible with Pareto efficiency. We will also compare our result to other rationales for negative marginal tax rates on low incomes.

Let τ^* denote the tax scheme that maximizes our social preferences. If τ^* looks like τ^1 or τ^4 in Figure 7 above, that is, if either $-\tau^*(0) \geq p$, or $w_m - \tau^*(w_m) \leq p$, then we (unsurprisingly) do not have much to say about its shape. In the (more relevant) case in which $y^p - \tau^*(y^p) = p$ for some $y^p \leq w_m$, we are able to derive the following property of the optimal tax scheme. The income subsidy obtained by the

[12] Note that without the formula of the optimal tax, we cannot derive classic comparative statics results of, e.g., a change in the distribution of wages or a change in the elasticity of labor supply. Deriving the optimal tax from the statistical knowledge of the distribution of wages and labor supplies in a specific economy, on the other hand, can be easily done using simulation methods as, e.g., the one proposed and illustrated in Fleurbaey and Maniquet (2018).

agents earning exactly y^p is the largest subsidy: for all $y \geq 0, \tau^*(y) \geq \tau^*(y^p)$. Let us prove this claim graphically in Figure 8.

Let us assume the claim is not true. Then, the optimal tax scheme τ^* is such that for some $\underline{y}$, $\tau^*(\underline{y}) < \tau^*(y^p)$. In the figure, $\underline{y} < y^p$, but the argument holds in the case $\underline{y} > y^p$ as well. The figure shows the curve $y - \tau^*(y)$ as well as a 45° line through $\left(y^p, y^p - \tau^*(y^p)\right)$. This line represents the function $y - \tau(y)$ defined by for all y, $\tau(y) = \tau^*(y^p)$. Note that τ amounts to transfer a strictly positive amount of money to all agents, which is obviously unfeasible.

Let us consider a new tax scheme, τ', defined by

$$\tau'(y) = \max\left\{\tau(y), \tau^*(y)\right\}.$$

This is equivalent to stating that $y - \tau'(y) = \min\left\{y - \tau(y), y - \tau^*(y)\right\}$. At each earning level, the new tax amount is either unaffected or it is larger. Let us note that $\tau'(y^p) = \tau^*(y^p)$ so that the value of the social criterion is the same under τ' as under τ^*. This also means that all agents earning y^p under τ^* still choose to earn y^p under τ'.[13] More generally, all agents choosing y under τ^* such that $\tau'(y) = \tau^*(y)$ still choose y under τ'. The other agents are likely to change their earning level, but, given the way τ' is defined, they will pay more tax under τ' than under τ^*. In summary, τ' leads to the same level of the social criterion but collects a budget surplus, which can be redistributed to agents earning around y^p, increasing the social criterion. That proves that τ^* is not optimal, a contradiction.

As a result, the graph of the optimal tax scheme τ^* needs to lie everywhere on or below the 45-degree line through (y^p, p). If the graph is strictly below the line at $y = 0$, then marginal tax rates are on average strictly negative and there is a negative participation tax rate (a participation subsidy). We can summarize the argument above in the following proposition.[14]

PROPOSITION 3: *Under Assumption 1, the optimal tax scheme τ^* satisfies the following property: either $-\tau^*(0) \geq p$, or $w_m - \tau^*(w_m) \leq p$, or for y^p such that $y^p - \tau^*(y^p) = p$,*

(*i*) $\lim_{y \to y^p} d\tau^*(y)/dy \leq 0$ *and marginal tax rates are on average nonpositive below y^p, and*

(*ii*) $\lim_{y^p \leftarrow y} d\tau^*(y)/dy \geq 0$ *and marginal tax rates are on average positive above y^p.*

[13] In this reasoning, chosen earning levels are those that maximize utility globally, which means that we take account of behavioral responses both at the extensive and the intensive margins.

[14] It is possible to be more precise about the shape of the optimal tax under additional assumptions on the distribution of types. An immediate example can be derived from Brett and Weymark (2017). They study the shape of the tax scheme that any specified agent would find optimal for herself given her preferences over her own bundles, under the assumption that wage and preference heterogeneities can both be captured by a single parameter and the assumption of a continuum of agents. Let us assume that the optimal tax scheme τ^* according to our criterion is also the preferred tax scheme of some of the agents choosing (y^p, p). This is more likely the more convex preferences of those agents are around bundle (y^p, p). Brett and Weymark (2017) prove that marginal tax rates are negative below y^p, except at the minimal earning level if it is positive; there is a kink at y^p; and marginal tax rates are positive above y^p, where they are designed in such a way as to maximize the collected tax.

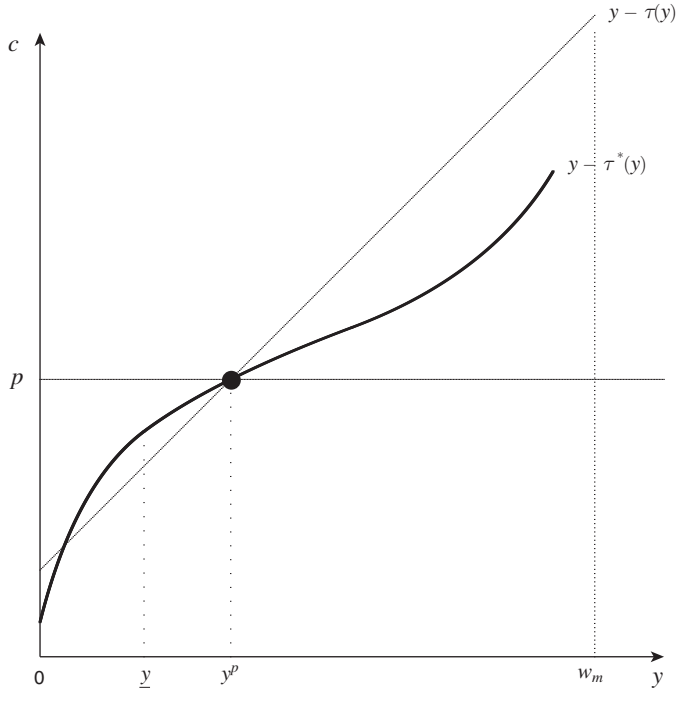

Figure 8. Illustration of Proposition 3

This result should be contrasted with related results recently presented by Saez and Stantcheva (2016). They propose a generalized social weight approach to the optimal tax formula. There are two main features of their approach. First, they focus on local optimality rather than global optimality. A tax scheme is locally optimal if no budget-neutral perturbation of the tax scheme increases the value of the social criterion. Second, this focus on local perturbations allows them to express the effect on the value of the social criterion as a weighted sum of the marginal changes in incomes.

Based on this, the authors study criteria of poverty alleviation that are different from our criterion but that are also compatible with Pareto efficiency. Their first criterion is the minimization of the income poverty-gap ratio. That criterion minimizes the sum of income gaps, that is, the difference between actual incomes and the poverty line. In a sharp contrast to what we obtain, they prove that the optimal tax scheme involves positive marginal tax rates below the poverty line.[15]

Their second criterion is the head count ratio (see Saez and Stantcheva 2016, proposition B2, online Appendix, 11). They find that the optimal tax scheme involves a negative marginal tax rate on the lowest incomes but not on all incomes below y^p. The optimal tax scheme that minimizes the head count ratio would, therefore, be closer to our optimal tax scheme than the one that minimizes the poverty-gap ratio.

[15] This result is more in line with the optimal tax schemes derived from weighted utilitarian social welfare functions when the weights represent a concern for poverty alleviation, such as in Kanbur and Tuomala (2011).

This may come as a surprise. Indeed, remember that our criterion is derived from the normative property that a disposable income transfer from an agent above the poverty line to an agent below it should be a social improvement. We could have thought that this property would bring us close to the poverty-gap ratio, but it does not.

Recent developments in the literature on optimal labor income taxation have reached the conclusion that optimal tax rates on low incomes could be negative. The justifications for such negative tax rates are twofold. First, negative tax rates turn out to be optimal when there are good reasons to incentivize agents to participate in the labor force and accept jobs. This may come from the existence of fixed cost of labor force participation (such as in Saez 2002; see also Diamond 1980 and the empirical application of Blundell and Shephard 2012) or a present bias (such as in Lockwood 2020). These results are complementary to ours. It is striking that we obtain the same property of the optimal tax scheme without assuming any cost to participate in the labor force. That means that our results would simply be reinforced if we added such costs to the model. This relates to the discussion that we held at the end of Section I.

Second, negative tax rates (or, to put it differently, increasing subsidies on low incomes) can emerge from the maximization of normative objectives (remember that marginal tax rates on low incomes are positive in the Mirrlees 1971 solution). Here again we need to distinguish between two streams of the literature.

The first stream sticks to the traditional utilitarian objective. Then, negative tax rates are obtained when individuals have different preferences and social weights are a function of these preferences, like in Boadway et al. (2002).

The second stream studies social objectives that embed fairness principles as done in this paper. In Fleurbaey and Maniquet (2006), the normative objective includes the principle that no redistribution should take place in a hypothetical society in which all individuals have the same wage. As a result, the tax system should try to replicate such a hypothetical society for low incomes, with the common wage as high as possible and, typically, larger than the actual minimum wage. This is why individuals working full time at the minimum wage should receive a larger subsidy than those working less. In Fleurbaey and Maniquet (2007) and Henry de Frahan and Maniquet (forthcoming), the normative objective is an aggregation of individual well-being indices, and an individual's well-being level does not directly depend on her labor time but on the hypothetical first-best budget to which she is indifferent. As a result, the tax system should try to replicate a first-best budget on low incomes, and this is achieved through a zero marginal tax, sometimes complemented with negative rates just below the earning level of low-income earners working full time.

A maybe surprising feature of Proposition 3 is that the properties of the optimal tax scheme that are listed do not depend on the distribution of types in the population. This can be explained by the fact that these properties have been deduced from the fact that our social preferences are egalitarian in well-being measures so that only the minimal well-being level is relevant. The distribution of types would matter a lot, however, if we were to identify the precise value of y^p for which the properties listed in the proposition hold and the amounts of taxes $\tau(y)$ for all y. For instance, for a fixed distribution of wages and a fixed poverty line, the usual result that the

tax system will be more generous toward the low-income individuals (higher $-\tau(y)$ for $y \leq y^p$) the lower the behavioral responses to taxation (i.e., the lower the labor supply elasticity) still holds. Let us note that Proposition 2 was proven by taking both the intensive and extensive margin behavioral responses into account.

VI. Application

We now turn to the application of the criterion developed in Section IV, i.e., its use to evaluate existing tax schemes, which are thus assumed to be *suboptimal*, i.e., not maximizing our social preferences. For this exercise, we need to draw the actual budgets that agents face. This requires that we take account of the fact that taxes and subsidies depend on the household composition, through, for instance, family allowances or in-work benefits that depend on the number of children. Our theory, however, does not teach us anything about comparing well-being or poverty across households of different size. As a consequence, we refrain ourselves from engaging in this kind of comparisons. Rather, we partition the population into household types and study each household type in turn.[16] Budgets are drawn for these household types using the OECD tax-benefit calculator, which takes account of all relevant regulatory aspects that transform pretax incomes into after-tax incomes in OECD countries.[17]

It is important to note that we draw the budgets under the assumption that information constraints are the only departure from a first-best world; that is, we assume away all other sources of frictions on the labor market. Nonetheless, our methodology lends itself particularly well to the main frictions that result in distortions of the budget sets facing agents. For instance, the lack of childcare can be represented as a truncation of the budget sets above some maximal feasible labor time. Long-term unemployment can be represented by a budget that only contains the point at which labor time is zero. In these cases, the application of our criterion follows immediately. Given that the criterion is constructed as a maximin aggregator, absolute priority should be given to enhancing the well-being of the worst off, who are the constrained agents. The conclusion would then be that the well-being of these agents should be enhanced, either by increasing the transfers to these agents, by decreasing unemployment, or by providing universal childcare, whichever is less costly. This is of course a very important conclusion to reach, but it is fair to recognize that we don't need the machinery of our criterion to reach such an obvious conclusion. Moreover, the nature of our social preferences is that once the well-being of these constrained agents has been maximized, absolute priority should then be given to the worst off among the working agents. This is when our criterion exhibits its full power, and this is the exercise we develop in this section.[18]

[16] This means that in case we analyze the well-being of households with several decision-makers, we implicitly assume a unitary model of such households.

[17] See the OECD website "Benefits and wages" (OECD 2019a), which gives access to the OECD tax-benefit calculator as well as country-specific information on the policies modeled.

[18] In online Appendix E.3, we additionally propose a simple approximation to take into account the fact that the probability to find a job might be lower than one.

We evaluate the ability of tax-transfer policies to alleviate poverty given official poverty lines, the official notion of disposable income, and the legal minimum wage. As a result, we restrict our sample to the EU15 countries that have a legal minimum wage[19] and the United States (US), taking the tax-transfer system of the state of Michigan as an example,[20] and calculate budgets for tax-transfer rules of 2013. Precisely, the tax-transfer policies that are taken into account are income support and social assistance (SA),[21] family and child benefits (FB), housing benefits (HB), in-work benefits (IW), labor income taxes (IT), and (employee) social insurance contributions (SC). We assume that policies implemented to fight against nontake-up and fiscal evasion are distinct from the definition of the tax-transfer system itself, and we thus abstract from the latter phenomena here. We come back to this in Section VII.

Disposable income is derived on the basis of the labor income that is earned when increasing hours worked from zero to full time. The wage earned is the legal monthly minimum wage in each country in 2013 as reported by EUROSTAT (2015a); see fifth column in Table 1.[22]

In the Europe 2020 strategy (European Commission 2010), the European member states have agreed to use 60 percent of national median equivalized disposable income as the at-risk-of-poverty indicator. We thus use this poverty threshold for the European countries in our analysis as reported by EUROSTAT (2015b). For the United States, we rely on the Supplemental Poverty Measure reported by the US Census Bureau.[23] Using country-specific poverty lines means that we do not try to compare poverty across countries, which would have required to use a common poverty line. Despite the ambiguous question of how to define such a common poverty line, we wish to compare how tax systems do succeed in alleviating poverty in the selected countries, given the way governments themselves define the poverty line.

Results are presented for single-parent households with two children (aged 10 and 12). The reason is that those households are known to be at high risk of poverty (EUROSTAT 2020).

The details about the methodology are presented in the online Appendix. There, we also identify how each national tax scheme treats different households differently and test the robustness of our application to changes in the definition of the poverty line, changes in the policies that are simulated, and changes due to the

[19] Belgium (BE), France (FR), Germany (GE), Greece (GR), Ireland (IR), Luxembourg (LU), the Netherlands (NL), Portugal (PT), Spain (SP), the United Kingdom (UK).

[20] Tax-transfer systems in the United States are largely state specific. The OECD calculator uses the tax-transfer system of Michigan, a typical manufacturing region, to represent the United States. Michigan's Temporary Assistance for Needy Families (TANF) and unemployment insurance benefits (not modeled in the present paper) are somewhat above the average for all states. In the following, for simplicity and consistency, we continue referring to the United States to denote Michigan's tax-transfer system.

[21] Unemployment *insurance* benefits are not taken into account as they are typically conditional on past labor force participation and social contributions. As a result, young or long-term unemployed people typically do not benefit from it. Unemployment *assistance* benefits, which are not based on previous contributions, are considered to be part of SA.

[22] Germany introduced a legal minimum wage in 2015. The wage for Germany is this minimum wage deflated to 2013.

[23] Detailed information about this measure and its derivation for 2013 can be found in Short (2014). A historical comparison of official poverty measures used in the United States up to the Supplemental Poverty Measure can be found in Meyer and Sullivan (2012).

Table 1—Social Welfare and Its Decomposition for Single Households with Two Children

	$W^p_{\min}$	$W^p_{\min} w_m$	b	t	w_m	p	$\frac{w_m}{p}$	$\frac{b}{p}$
US	51.13	646.72	1,018.28	7.02	1,264.96	1,739.71	72.71	58.53
BE	97.95	1,953.65	2,128.09	84.42	1,994.57	2,284.44	87.31	93.16
FR	97.00	1,842.41	1,571.74	62.58	1,899.48	2,225.79	85.34	70.61
GE	0.00	0.00	2,116.80	68.25	1,938.84	2,076.62	93.37	101.93
GR	107.09	972.47	182.87	12.58	908.10	890.04	102.03	20.55
IR	0.00	0.00	2,353.30	75.37	1,941.48	2,025.62	95.85	116.18
LU	15.95	397.07	3,195.76	84.76	2,489.11	3,538.16	70.35	90.32
NL	0.00	0.00	2,321.35	74.69	1,951.51	2,213.57	88.16	104.87
PT	112.99	849.07	495.40	61.36	751.48	868.79	86.50	57.02
SP	146.97	1,469.54	771.64	70.32	999.86	1,436.79	69.59	53.71
UK	0.00	0.00	2,510.71	77.68	1,659.93	1,986.20	83.57	126.41

Notes: $W^p_{\min}$ in percent of full time. w_m/p and b/p in percent. All other values in monthly US$. t denotes effective marginal tax rates, b the disposable income earned when not working, p denotes official national poverty lines, w_m statutory minimum wages. Germany introduced a legal minimum wage in 2015. w_m for Germany is this minimum wage deflated to 2013.

heterogeneity of unemployment rates among the countries under analysis. Finally, we systematically provide the full set of results of our application, that is, for all selected countries and six different household types.

A. *Measuring and Decomposing Social Welfare*

All elements shown in Table 1 are directly deduced from the budget curves constructed with the OECD calculator or calculated making additional use of the country-specific minimum wages and poverty lines. The first column shows $W^p_{\min}$, the measure of social welfare according to R^{lex}, as derived in Section IV, in percentage of a full-time job, for single-parent households with two children. For instance, the worst-off households of that type in the United States are those who work 51.13 percent of their time (where 100 percent of one's time means a full-time job). The percentage of labor time needed to reach the poverty line varies from 0 percent in Germany, Ireland, the Netherlands, and the United Kingdom, where even those who do not work have the opportunity to get out of poverty, to countries in which it is impossible for low-skill households to get out of poverty even by working full time, namely Greece, Portugal, and Spain.

The second column shows the pretax income corresponding to this measure, $W^p_{\min} w_m$. For instance, the worst off in the United States, according to R^{lex}, are those who work at the minimum wage and earn US$646.72. This is the main result delivered by our approach: if a policymaker in the United States is interested in the normative property of *Poverty Reduction* defined above, a reform of the tax-transfer system should go in the direction of increasing—in the limit of what is feasible—the disposable income of those earning US$646.72.

The next columns of the table decompose $W^p_{\min}$ into the policy parameters that determine it. For the sake of simplicity, we approximate the tax function $\tau(y)$ on low incomes by $\tau(y) = -b + ty$; that is, we do as if the marginal tax rate were constant over low incomes. The first policy parameter is what can be called "basic"

income, b, that is, the disposable income of those who do not earn anything. It gives us the *level* of the opportunity set available to those who do not work. The second one is the rate t at which low incomes are effectively taxed. It shows, given any additional dollar earned, how much of it is taken away by the tax-transfer system, on average, below the minimum wage. This tool gives us a summary of the *shape* of the opportunity set of low-skill households, that is, of how their labor is rewarded. The third one is the minimum wage itself, w_m, but its effect on $W^p_{\min}$ is best seen when it is expressed as a percentage of the poverty line, w_m/p, which is done in the column next to it. It shows that except in Greece, no household of the type we are looking at could reach the poverty line without benefits.

If the tax function τ were indeed linear on low incomes, then our measure of social welfare would satisfy the equation $p = b + W^p_{\min} w_m(1-t)$, which gives us

$$W^p_{\min} = \frac{p-b}{w_m(1-t)},$$

illustrating how the combination of the three policy parameters, b, t, and w_m, determines social welfare and how they can be used to increase it. Social assistance, family benefits, and housing benefits typically determine b. How these benefits fade out when gross income increases, in-work transfers, income tax, and social security payments together determine t. Finally, w_m is a direct policy instrument.

A more detailed look at the table reveals two facts. First, there is less heterogeneity in the minimal income policy (w_m/p) than in the basic income (looking at b/p in the last column, for better comparison) and tax rates. Second, the generosity of the tax systems toward the zero-income earners requires to tax low incomes more, consistently with classical labor income taxation theory (the coefficient of correlation between b/p and t is equal to 71.6 percent).

Besides these simple observations, the analysis of the relationship between policies and the ability of tax systems to alleviate poverty points toward two major lessons. First, at the risk of stressing the obvious, the only way for a country to completely alleviate poverty (under the assumption of full take-up of all benefits) is to have a basic income larger than the poverty line $(b \geq p)$. This is the case for Germany, Ireland, the Netherlands, and the United Kingdom. Note that out of these four countries, only Ireland eliminates poverty of all types of households as can be seen from the tables in online Appendix F.2.[24]

Second, among the other countries, there is no clear policy pattern that guarantees to have a high social welfare. Greece, for instance, has a larger minimum wage (w_m/p) and taxes less (a lower t) than Belgium, France, and Luxembourg but has a lower social welfare because the basic income is also (much) lower than in these countries. Belgium is more generous toward zero-income earners and has a larger minimum wage than the United States or France but reaches a lower social welfare because of a larger tax rate. From a similar table computed for couples without

[24] Out of the benefits that we take into account, housing benefits are the ones with the lowest take-up rate. If we remove these benefits, then no country succeeds in lifting all households out of poverty. This is discussed in online Appendix E.2.

children (see online Appendix F.2), we can see that Portugal is less generous than Spain in basic income and taxes incomes more, but it reaches a larger social welfare because of a significantly larger minimum wage.

An important consequence of this second observation is that the countries which seem to have a tax system closer to the one that the theory proves optimal (see Proposition 3 in Section V), that is, the countries with the lowest income taxation rates, are not the ones reaching the largest social welfare. This suggests that all these countries with a positive $W^{p}_{\min}$ are far from maximizing social welfare as it is defined in this paper; that is, they are far from reducing poverty as much as they can.

B. *Extended Cross-Country Comparison*

In Figure 9, we draw the entire budget curves for single-parent households with two children for all countries over the relevant income span. Country-specific budget curves are made comparable by rescaling the axes such that all minimum wages (respectively poverty lines) correspond to coordinate 1 along the horizontal (respectively vertical) axis.

The figure clearly shows the large variety of policies across countries. The budget curves in Belgium, Luxembourg, the Netherlands, and Spain are characterized by a 100 percent tax rate on the lowest incomes, sometimes after a small interval of lower tax rates. This 100 percent tax rate is produced by the \$1 decrease in social assistance following any \$1 increase in gross income over that interval. The budget curves in Germany, France, and Portugal are strictly increasing but at a slow rate, illustrating the way social assistance and, in the case of France and Germany, housing benefits fade out as gross income increases. The budget curves in Ireland, the United Kingdom, and the United States are steeper, illustrating the in-work benefits existing in these three countries, even if its implementation gives a much smoother budget curve in the United States than in the other two countries. The budget curve in Greece is characterized by the nonexistence of basically any income support program, except modest family benefits, coupled with very low income tax rates.

VII. Concluding Remarks

In this paper, we have proposed a normative property to make social preferences sensitive to poverty in a way that is compatible with the Pareto criterion. From that property, we have built social preferences, from which we have derived a simple criterion to evaluate tax schemes: tax policymakers adopting the proposed social preferences should minimize the labor time needed to get out of poverty. This criterion has then been applied to analyze tax schemes in Europe and the United States. One of the main achievements of this paper is indeed to have been able to go in a consistent way from an abstract property of social preferences to propositions of reforms for *existing* income tax systems. As an additional result, we have shown that the *optimal* tax scheme following from our criterion implies negative marginal tax rates at very low incomes.

Our application has been developed by considering households that benefit from all the policies we have contemplated. Given that some policies that are part of

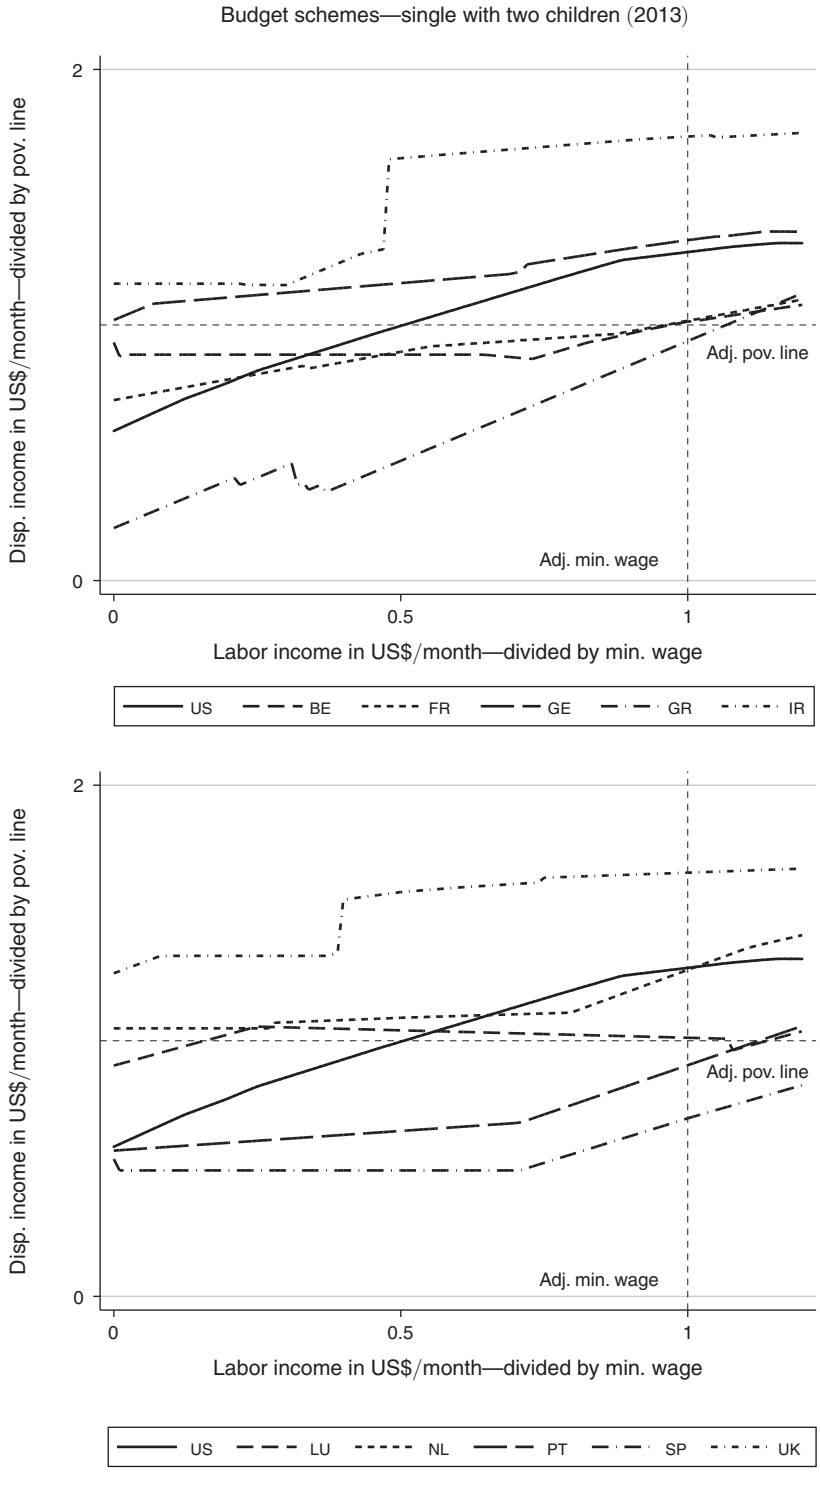

FIGURE 9. CROSS-COUNTRY COMPARISON OF ENTIRE BUDGET CURVES FOR SINGLE HOUSEHOLDS WITH TWO CHILDREN

the fiscal system do not have a full take-up rate, other households do actually live below the budget curves that we have drawn. We consider that our approach is the correct one, though. Indeed, we should distinguish between the evaluation of a tax system dedicated to be applied to all households and the observation that there is room for increasing the take-up rate. Understanding the causes for non-take-up is a very important, though different, research topic (see, for instance, Currie 2004 for a survey of the research on this topic).

A given tax system could yield two different income distributions if applied to two different economies. The social preferences we define in this paper, however, would be said to be equally satisfied in these two economies. Indeed, under the assumption that agents choose their best possible bundle, our criterion only depends on the opportunities they are given by the tax system, not on the specific choices they make when facing these opportunities (provided, of course, Assumption 1 in Section IV is satisfied). This is in sharp contrast to evaluating the ability of tax-transfer systems to alleviate poverty on the basis of the statistical distribution of incomes. We consider this an advantage of our approach, which clearly distinguishes between the design of the policy and the way agents react to it. Two systems offering the same opportunities to agents should receive the same praise and blame, independently of agents' choices.

The criterion we have defined evaluates an entire tax scheme on the basis of a single point, the intersection between the graph of the tax scheme function and the graph of the poverty line. That implies that among suboptimal allocations, groups of socially indifferent allocations are easily identified. This suggests using complementary criteria to discriminate among these groups. In case two allocations are deemed socially indifferent, it seems natural to us to look at the pretax income necessary to reach a consumption level just below the poverty line, and so on in case of further indifference. Among the allocations analyzed in this paper (and the online Appendix), we don't find any exact indifference but, instead, a clear ranking of tax systems. However, in some cases, the labor time required to reach the poverty line comes very close for pairs of countries, and one might be interested in exploring these cases further, arguing that they could be considered as almost indifferent allocations. In Figure 9 above, a near tie appears for France and Belgium (see also $W^{p}_{\min}$ in Table 1). Taking our suggestion to look at marginally lower poverty lines in that case, we then conclude that France has a better tax system than Belgium because it requires to work less than in Belgium to reach the "new" poverty line.

We have presented above the derivation of a simple evaluation criterion for tax schemes from social preferences. It is important to note that this exercise did not require the detour through the derivation of a formula describing the optimal tax-transfer system, the kind of formula on which the literature in the field has devoted its largest effort so far. The normative property we have studied in this paper is natural but certainly not the only one worth investigating, even if one concentrates on the goal of poverty alleviation.

REFERENCES

Akerlof, George A. 1978. "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning." *American Economic Review* 68 (1): 8–19.

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

Blundell, Richard, and Andrew Shephard. 2012. "Employment, Hours of Work and the Optimal Taxation of Low-Income Families." *Review of Economic Studies* 79 (2): 481–510.

Boadway, Robin. 2012. *From Optimal Tax Theory to Tax Policy*. Munich Lectures in Economics. Cambridge, MA: MIT Press.

Boadway, Robin, and Michael Keen. 2000. "Redistribution." In *Handbook of Income Distribution*, Vol. 1, edited by Anthony B. Atkinson and François Bourguignon, 677–789. Amsterdam: Elsevier.

Boadway, Robin, Maurice Marchand, Pierre Pestieau, and María Del Mar Racionero. 2002. "Optimal Redistribution with Heterogeneous Preferences for Leisure." *Journal of Public Economic Theory* 4 (4): 475–98.

Brett, Craig, and John A. Weymark. 2017. "Voting over Selfishly Optimal Nonlinear Income Tax Schedules." *Games and Economic Behavior* 101: 172–88.

Carvalho, Leandro S., Stephan Meier, and Stephanie W. Wang. 2016. "Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday." *American Economic Review* 106 (2): 260–84.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1: 451–88.

Chetty, Raj. 2015. "Behavioral Economics and Public Policy: A Pragmatic Perspective." *American Economic Review* 105 (5): 1–33.

Currie, Janet. 2004. "The Take up of Social Benefits." NBER Working Paper 10488.

Dalton, Hugh. 1920. "The Measurement of the Inequality of Incomes." *Economic Journal* 30 (119): 348–61.

Decancq, Koen, Marc Fleurbaey, and François Maniquet. 2019. "Multidimensional Poverty Measurement with Individual Preferences." *Journal of Economic Inequality* 17 (1): 29–49.

Diamond, P. 1980. "Income Taxation with Fixed Hours of Work." *Journal of Public Economics* 13 (1): 101–10.

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

European Commission. 2010. *EUROPE 2020—A European Strategy for Smart, Sustainable and Inclusive Growth*. Brussels: European Commission.

EUROSTAT. 2015a. "Minimum Wages." European Commission. https://ec.europa.eu/eurostat/web/products-datasets/-/tps00155 (accessed August 3, 2015).

EUROSTAT. 2015b. "Mean and Median Income by Household Type." European Commission. https://ec.europa.eu/eurostat/web/products-datasets/-/ilc_di04 (accessed August 7, 2015).

EUROSTAT. 2020. "People at Risk of Poverty or Social Exclusion by Income Quintile and Household Type." European Commission. https://ec.europa.eu/eurostat/web/products-datasets/-/ilc_peps03 (accessed May 5, 2020).

Fleurbaey, Marc, and François Maniquet. 2006. "Fair Income Tax." *Review of Economic Studies* 73 (1): 55–83.

Fleurbaey, Marc, and François Maniquet. 2007. "Help the Low Skilled or Let Hardworking Thrive? A Study of Fairness in Optimal Income Taxation." *Journal of Public Economic Theory* 9 (3): 467–500.

Fleurbaey, Marc, and François Maniquet. 2011. *A Theory of Fairness and Social Welfare*. Econometric Society Monographs. Cambridge, UK: Cambridge University Press.

Fleurbaey, Marc, and François Maniquet. 2018. "Optimal Income Taxation Theory and Principles of Fairness." *Journal of Economic Literature* 56 (3): 1029–79.

Foster, James, Joel Greer, and Erik Thorbecke. 1984. "A Class of Decomposable Poverty Measures." *Econometrica* 52 (3): 761–66.

Hausmann, Daniel M. 2012. *Preference, Value, Choice, and Welfare*. Cambridge, UK: Cambridge University Press.

Henry de Frahan, Lancelot, and François Maniquet. Forthcoming. "Preference Responsibility vs Poverty Reduction in the Taxation of Labor Incomes." *Journal of Public Economics*.

Jacquet, Laurence, and Etienne Lehmann. Forthcoming. "Optimal Income Taxation with Composition Effects." *Journal of the European Economic Association*.

Kanbur, Ravi, and Michael Keen. 1989. "Poverty, Incentives and Linear Income Taxation." In *The Economics of Social Security*, edited by Andrew W. Dilnot and Ian Walker, 99–115. Oxford: Clarendon Press.

Kanbur, Ravi, Michael Keen, and Matti Tuomala. 1994a. "Optimal Non-linear Income Taxation for the Alleviation of Income-Poverty." *European Economic Review* 38 (8): 1613–32.

Kanbur, Ravi, Michael Keen, and Matti Tuomala. 1994b. "Labor Supply and Targeting in Poverty Alleviation Programs." *World Bank Economic Review* 8 (2): 191–211.
Kanbur, Ravi, Tuuli Paukkeri, Jukka Pirttilä, and Matti Tuomala. 2018. "Optimal Taxation and Public Provision for Poverty Reduction." *International Tax and Public Finance* 25 (1): 64–98.
Kanbur, Ravi, and Matti Tuomala. 2011. "Charitable Conservatism, Poverty Radicalism and Inequality Aversion." *Journal of Economic Inequality* 9: 417–31.
Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.
Lockwood, Benjamin B. 2020. "Optimal Income Taxation with Present Bias." *American Economic Journal: Economic Policy* 12 (4): 298–327.
Lockwood, Benjamin B., and Matthew Weinzierl. 2015. "De Gustibus non est Taxandum: Heterogeneity in Preferences and Optimal Redistribution." *Journal of Public Economics* 124: 74–80.
Maniquet, François, and Dirk Neumann. 2021. "Replication data for: Well-Being, Poverty, and Labor Income Taxation: Theory and Application to Europe and the United States." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E118721V1.
Maniquet, François, and Yves Sprumont. 2004. "Fair Production and Allocation of an Excludable Nonrival Good." *Econometrica* 72 (2): 627–40.
Meyer, Bruce D., and James X. Sullivan. 2012. "Identifying the Disadvantaged: Official Poverty, Consumption Poverty, and the New Supplemental Poverty Measure." *Journal of Economic Perspectives* 26 (3): 111–36.
Mirrlees, J.A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (2): 175–208.
Mullainathan, Sendhil, and Eldar Shafir. 2013. *Scarcity—Why Having Too Little Means so Much*. New York: Times Books.
OECD. 2019a. "Benefits and Wages." OECD. https://www.oecd.org/social/benefits-and-wages/ (accessed May 5, 2020).
OECD. 2019b. "Hours of Work Needed to Escape Poverty for Workless Families." OECD. https://stats.oecd.org/Index.aspx?DataSetCode=HOURSPOV (accessed May 5, 2020).
Pazner, Elisha A. 1979. "Equity, Nonfeasible Alternatives and Social Choice: A Reconsideration of the Concept of Social Welfare." In *Aggregation and Revelation of Preferences*, edited by Jean-Jacques Laffont. Amsterdam: Elsevier.
Pigou, Arthur C. 1912. *Wealth and Welfare*. London: Macmillan.
Pirttilä, Jukka, and Matti Tuomala. 2004. "Poverty Alleviation and Tax Policy." *European Economic Review* 48 (5): 1075–90.
Saez, Emmanuel. 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics* 117 (3): 1039–73.
Saez, Emmanuel, and Stefanie Stantcheva. 2016. "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." *American Economic Review* 106 (1): 24–45.
Samuelson, Paul A. 1974. "Complementarity: An Essay on the 40th Anniversary of the Hicks-Allen Revolution in Demand Theory." *Journal of Economic Literature* 12 (4): 1255–89.
Samuelson, P.A. 1977. "Reaffirming the Existence of 'Reasonable' Bergson-Samuelson Social Welfare Functions." *Economica* 44: 81–88.
Samuelson P.A., and S. Swamy. 1974. "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis." *American Economic Review* 64 (4): 566–93.
Short, Kathleen. 2014. *The Supplemental Poverty Measure: 2013*. Washington, DC: United States Census Bureau.
Stiglitz, Joseph E. 1987. "Chapter 15—Pareto Efficient and Optimal Taxation and the New New Welfare Economics." In *Handbook of Public Economics*, Vol. 2, edited by Alan J. Auerbach and Martin Feldstein, 991–1042. Amsterdam: Elsevier.
Wane, Waly. 2001. "The Optimal Income Tax when Poverty Is a Public Bad." *Journal of Public Economics* 82 (2): 271–99.
Weinzierl, Matthew. 2017. "Popular Acceptance of Inequality Due to Innate Brute Luck and Support for Classical Benefit-based Taxation." *Journal of Public Economics* 155: 54–63.
Weinzierl, Matthew. 2018. "Revisiting the Classical View of Benefit-Based Taxation." *Economic Journal* 128 (612): F37–64.