Why we do not Tag to Tax: Pareto Optimality versus Horizontal Equity

Vinzenz Ziesemer*

Erasmus University Rotterdam

November 2020

Abstract

Currently accepted Welfarist theories of optimal taxation suggest making taxes dependent on height, race, and gender. This paper establishes the precise condition under which such tags are Pareto improving, and thus *must* be used by any Welfarist policy maker: tags are Pareto improving if and only if they identify Laffer effects. In practice however, many such tags are intuitively rejected on the basis of vague notions of Horizontal Equity, i.e. the equal treatment of equals. The paper formalizes that notion by minimally constraining the classic Welfarist approach. In doing so, it achieves a close correspondence between the tags prescribed by theory and those used in practice.

1 Introduction

Arguably, the biggest difference between the theory and practice of taxation is the use of tags (cf. Mankiw, Weinzierl, and Yagan, 2009). Theory, asking how a government can maximize the total welfare of its people, suggests that we should use any information available to optimize our tax system. As a result, tax theory suggests we should tax the tall more than the short, differentiate taxes by gender and ethnicity, and use genetic information on people's ability to earn money to determine how much taxes they should pay.

*Email: vinzenz.ziesemer@eui.eu. I am grateful to Árpád Ábrahám, David Koll, Nils Grevenbrock, Albert-Jan Hummel and Thomas Ziesemer for extensive comments and discussions.

Using a tag such as height for taxation may seem ridiculous or discomforting. Yet, there is one very strong argument in favor of doing so: the Pareto Principle. Indeed, it may be true that height-based taxation could make everyone better off in a Welfarist sense.[1] As Mankiw and Weinzierl (2010), who discuss height-based taxation, put it: *"Nevertheless, if a nontrivial Pareto-improving height tax were possible, and if people both understood and were convinced of that possibility, it is our sense that most people would be comfortable with such a policy"* (p. 173).

This paper contains two contributions to the literature. First, it answers the following question: Under which conditions is it precisely that tags facilitate a Pareto improvement? It considers the standard Mirrleesian environment, the most commonly used theoretical framework in the study of optimal taxation. Tags are shown to be Pareto improving *if and only if* they identify a subgroup of the population for which, compared to the optimal tax schedule without these tags, taxes can be reduced without loss of revenue. In other words, tags are Pareto improving if and only if they identify Laffer effects. It immediately follows that if such a tag exists, then a Welfarist policy maker must use it. This implies a clear test for Pareto improvements by tagging. And as more data become available, more tags are likely to pass this test.

Yet in practice tax systems use almost no such information on personal characteristics. Why is it that we do not tag? The intuitive discomfort that comes with proposals such as height-based taxation is often related to concerns for Horizontal Equity, often defined as the principle that *equals should be treated equally* (cf. Diamond and Saez, 2011). This leaves much to be desired of course: what do we mean by equal? And by equal treatment? And what kind of tags does this principle then rule out?

As a second contribution, this paper develops a constrained Welfarist objective criterion in which there is a minimal concern for Horizontal Equity. The constraint works as follows: when an individual exerts the same effort and achieves the same result as another, he or she should be made at least as well off. Notions of Horizontal Equity have been much criticized for their reliance on the 'status quo' (the set of policies that happen to be in place at any given time), or on a 'natural state' that is hard to envision in reality (see Kaplow, 1989, 2000).[2] In addition, Horizontal Equity based objectives are often far from the more common Welfarist approach, and therefore lead to wildly different conclusions. This is undesirable, because other than the issue of tagging, the Welfarist objective does a good job at explaining why governments tax as they do. The constrained Welfarist criterion that this paper develops

[1]Welfarism, to be defined below, describes a criterion of welfare in which only the totality of (weighted) individual outcomes is of importance, but their relation to each other is not.

[2]The concept of a 'natural state' in this context typically describes a state in which there is no government.

relies neither on a 'natural state', nor does it diverge far from the original Welfarist criterion. Instead, it introduces a counterfactual comparison across individuals that constrains an otherwise Welfarist planner.

The constrained criterion proposed in this paper results in sharp predictions: it allows for some tags, but not for others, with little ambiguity. In particular, it prescribes the *equal treatment of equals* in a precisely defined sense, and entirely rules out the use of a class of tags that we will refer to as *diffuse*. The result is a close match with tags that are observed in reality, as well as with those that are not. Interestingly, the criterion also rules out a range of recent academic proposals on differentiated taxation, for example by gender.

In addition, the paper discusses extensions of its criterion to dynamic settings. A recent literature on dynamic optimal taxation suggests making taxes dependent on the history of past incomes. The paper discusses the extent to which these hold up to concern for Horizontal Equity. Under some extensions of the principle, history-dependent policies can be ruled out by the same rationale this paper uses to rule out tags such as those based on height. This potentially has implications for the normative dynamic optimal taxation literature.

In what follows, section 2 discusses previous literature. Section 3 introduces a Mirrleesian framework and analyzes the relation between tags, Pareto optimality, and the Welfarist criterion. It then introduces an constrained Welfarist criterion with concern for Horizontal Equity, and demonstrates some implications for tagging. Next, it discusses how the constrained criterion can be carried into a framework of dynamic taxation, and what implications this has. Finally, section 4 shows how the prescriptions of the constrained criterion match the practice of taxation, and compares this to the prescriptions of a standard Welfarist framework. Section 5 concludes.

2 Literature

On the first contribution of this paper, the main reference is Werning (2007). Werning considers an environment in which preferences are separable and homogeneous, and finds sufficient and necessary conditions under which a tax system is Pareto efficient. The conditions amount to a test for Laffer effects. He then generalizes these conditions to the case of tagged groups, finding that the same conditions for Pareto efficiency now apply at the level of the group. This paper concerns a more general environment and thus provides a more general result. In addition, the exposition focuses on the role of tags rather than general conditions for Pareto efficiency.

The second contribution of this paper, on the positive role of Horizontal Equity in taxation, relates to a larger literature. Weinzierl (2014) introduces the idea that governments might in reality follow a *mixed objective*. While they care for standard Welfarist objectives, they are also concerned with horizontal equity, which Weinzierl formalizes as a principle of *Equal Sacrifice*. This approach differs in several ways from the one pursued here. First, this formalization requires a 'natural state'. Second, the resulting framework, rather than just reducing the role of some tags, also alters policy prescriptions on other issues.

The main advantage of the approach of Weinzierl is that it allows to uphold the Pareto principle. As a result, it prescribes a reduced use of tags such as height, but as this paper shows the use of such tags must remain a requirement. Similarly, Saez and Stantcheva (2016) introduce the concept of generalized social marginal welfare weights. They show how the principle of Horizontal Equity can be embedded, combining social preferences for equity with an adherence to the Pareto principle. As this paper shows, such adherence to Pareto optimality cannot do without certain forms of discrimination.

Empirical work by Sausgruber and Tyran (2014) investigates the role of mixed objectives in the acceptance of taxes. They find that discriminatory taxes remain unpopular even when they deliver clearly dominant outcomes.

Previous literature on tagging has focused on solving optimal taxation problems with tags under standard Welfarist criteria. Akerlof (1978) provides the first extensive treatment of the tagging problem. Immonen, Kanbur, Keen, and Tuomala (1998) explain how to solve a problem with tags in an otherwise standard Mirrleesian framework. Viard (2001a), Viard (2001b), and Boadway and Pestieau (2006) report results for similar problems. Cremer, Gahvari, and Lozachmeur (2010) also calibrate their model, and report substantial welfare gains from the use of tags.

Also related to this paper is any literature that derives welfare criteria from basic principles (including Horizontal Equity, e.g. Musgrave, 1990), seeks to understand how social preferences and political decisions arise from individual preferences, considers welfare criteria further removed from the classic Welfarist criterion (including other formalizations of Horizontal Equity, such as the one by Auerbach and Hassett, 2002), or considers environments further removed from the classic Mirrleesian economy. I do not review these literatures here, because the focus of this paper is on tags in the context of positive optimal taxation.

3 Optimal Taxation, Tagging, and Horizontal Equity

I start by introducing a standard Mirrlees (1971) environment - which has one unobserved choice, one observed outcome, and one type of unobserved heterogeneity - with two added features: observed heterogeneity in preferences and one extra dimension of observed heterogeneity that can be correlated with everything else. I then analyze the link between tags, Pareto optimality, and Welfarism. Next, I introduce a constrained Welfarist criterion that respects Horizontal Equity, and discuss its implications for tagging. Finally, I devote some space to generalizing the principle of Horizontal Equity to a dynamic setting, as well as to how one might approach solving optimal taxation problems under Horizontal Equity.

3.1 Tags in a Mirrlees Environment

Denote individuals by index $i \in I$, which is a finite set. Each individual has an ability $\theta \in \Theta$ (a finite set of positive reals), which in combination with work time n results in production $y^i = \theta^i n^i$. Each individual has preferences over consumption $c \geq 0$ and work time $0 \leq n \leq 1$, which are ranked by a real-valued function $u^i(c^i, n^i)$ that is defined over the domain of the inputs. All u^i are assumed to be strictly increasing and concave in c, and strictly decreasing and convex in n.

As is standard in the Mirrlees setup, a social planner maximizes some criterion by setting tax policies. The planner observes neither θ^i nor n^i. Instead, he observes y^i, and some other individual characteristic, which is represented by a real number $\gamma^i \in \Gamma$. Hence, the planner can set a function for taxes and transfers $T(y^i, \gamma^i)$. In doing so, the planner is subject to a budget constraint:[3]

$$\mathbb{E} \sum_{i \in I} T(y^i_*, \gamma^i) \geq R. \tag{1}$$

Here, R represents some net revenue requirement of the government, which is a real number. The planner has knowledge of the functions u^i, and of the conditional distributions (cumulative density functions) of ability $F(\theta|\gamma, u^i)$.

Consumption is finally determined by the function $c^i = y^i - T(y^i, \gamma^i)$. From here on out, n^i_* denotes an optimal choice of agent i given the environment, and y^i_* and c^i_* the corresponding output and consumption. The problem of the social planner is then completed by a welfare criterion which the planner maximizes. For the remainder of this paper, we study the following generalized Welfarist objective function:

Definition 1 (Welfarist Criterion)**.**

$$\begin{aligned} U(\{u^i(c^i_*, n^i_*)\}_{i \in I}) =& \mathbb{E} \sum_{i \in I} w^i u^i(c^i_*, n^i_*), \\ & \forall i : w^i > 0. \end{aligned}$$

The utility functions must be interpreted cardinally and must be known to the planner, in order to allow for a trade-off between the welfare of different individuals. Positive welfare

[3] The use of the expectations operator reflects a technical issue: the planner faces a finite set of individuals, and so is uncertain about the revenues that any given tax schedule delivers.

weights must be attributed to all.[4] That social preferences are a function of individual utilities only, without further objectives or significant restrictions, is the defining feature of the Welfarist approach.

Allowing for heterogeneity in individual utilities is conceptually difficult in the Welfarist framework.[5] Formally, I simply assume the existence of cardinal u^i that are known to the planner. One interpretation of this framework that is equivalent for the purposes of this paper is as follows: we simply see utility functions as part of the social welfare function of the planner, so that their cardinality is a normative view the planner takes. At the same time, the planner knows the behavioral responses of the agents (at least at the group level). This allows us to formalize the problem as if the u^i are known to the planner.[6]

How does a Welfarist social planner proceed? In this type of environment, the planner would ideally (in first-best) set taxes based on individual ability and individual characteristics. This is ruled out, however, because information on individual abilities is not available to the planner. Instead, he resorts to providing incentives. As Mirrlees (1971) shows, the optimal tax function then depends (amongst other things) on the distribution of ability in the economy.

Tags provide additional information on the distribution of abilities, and therefore get the planner closer to his first-best. Because the tag splits the population into groups, which on average may have differing abilities, the planner can simply solve the optimal taxation problem per group, given some revenue requirement from the group. Next, the group revenue requirements are adjusted to minimize the social cost of meeting the total revenue requirement. Combining both steps yields a solution to the problem of the social planner where groups may be treated differentially. The reader is referred to Immonen, Kanbur, Keen, and Tuomala (1998) for a more extensive exposition.

3.2 Pareto Implications

When does the planner make the tax schedule dependent on tags? When tags are entirely uninformative of ability, they are not used. As Weinzierl (2014) points out, the opposite

[4]The case where w^i is equal for all individuals i is commonly referred to as the Utilitarian Criterion.

[5]First, consider this example: Two individuals provide the same n and receive the same c. Yet, they derive differential levels of welfare from this. Does a Welfarist treat them differentially? The answer seems to be outside of the scope of common definitions of Welfarism. (The same dilemma applies when welfare is differentially sensible to changes in n and c.) Second, suppose that in addition these cardinal u_i are not observed. Now how would the social planner elicit them? There is no obvious instrument for doing so. Cf. Piacquadio, 2017.

[6]For the interested reader, the optimal Welfarist taxation implications of having heterogeneity in both u and ability are studied in Jacquet and Lehmann (2015).

does not need to hold: welfare weights can be correlated with tags too, so that the net effect might be zero. However, as Mankiw and Weinzierl (2010) point out, these would be knife-edge cases: one would have to construct welfare weights so that tags are excluded, which is hard to align with the individualistic principle underlying Welfarism.

More importantly, I show below that a specific set of tags *must* be used by a Welfarist social planner. I do this by comparison to a counterfactual situation in which the planner does not observe tags, and sets a tax schedule $T(y)$. If tags identify Laffer effects, they are Pareto improving. And if they are Pareto improving, a Welfarist must use them. Thus, it is not true that the Welfarist criterion can always be *made to fit* basic notions of Horizontal Equity by adjusting welfare weights.

Much less trivially, I also show that the set of Pareto improving tags *only* consists of those that identify Laffer effects. The intuition is simple: the planner cannot raise effective taxes, so he must identify a subset of the population for which a tax reduction somewhere does not lead to a loss of tax revenue. Since the original situation was Pareto optimal, it must be the tag that allows the identification of such a subset. Thus, while tags may improve welfare by the Welfarist Criterion in many ways, they only lead to Pareto improvements if they identify Laffer effects. This characterization then implies a clear test for whether or not a tag is Pareto improving. Proposition 1 establishes these links between tags and Pareto optimality. I begin by defining some objects.

Definition 2 (Pareto improvement). *An allocation a, defined as $\{(c^i_a, n^i_a)\}_{i\in I}$, is said to be a Pareto improvement over an allocation b, defined as $\{(c^i_b, n^i_b)\}_{i\in I}$, if all agents (i) are at least as well off, and at least one agent is strictly better off, under a than under b:*

$$\forall i : u^i(c^i_a, n^i_a) \geq u^i(c^i_b, n^i_b), \textit{and}$$
$$\exists i : u^i(c^i_a, n^i_a) > u^i(c^i_b, n^i_b) > 0.$$

Definition 3 ($T^*(y)$). *Let $T^*(y)$ denote the tax schedule that maximizes the Welfarist Criterion if the planner would not observe the individual characteristics γ.*

Definition 4 ($Y_T, Y_T^{\gamma_k}, Y_{T/T'}^{\gamma_k}$). *Let y^i_T denote the income levels that are part of agent i's optimal choice set with respect to tax schedule T. For generic tax schedules T and T', define:*

- *Y_T as the set of incomes that is part of the optimum for some agent under tax schedule T: $\{y : \exists i \ (y \in y^i_T)\}$;*
- *$Y_T^{\gamma_k}$ as the set of incomes that is part of the optimum for some agent with characteristic γ_k under tax schedule T: $\{y : \exists i \ (\gamma^i = \gamma_k) \cap (y \in y^i_T)\}$; and*
- *$Y_{T/T'}^{\gamma_k}$ as the set of incomes that is part of the optimum for some agent with characteristic γ_k under tax schedule T, that is not optimal under tax schedule T': $\{y : \exists i \ (\gamma^i = \gamma_k) \cap (y \in y^i_T) \cap (y \notin y^i_{T'})\}$.*

Definition 5 (Identifying a Laffer effect)**.** *A characteristic* γ_k *is said to identify a Laffer effect if there exists a tax schedule* $T(y,\gamma) \le T^*(y)$ *such that, compared to tax schedule* $T^*(y)$*, we have the following for agents with characteristic* γ_k*:*

a. *Some income level* $\bar{y}$ *optimal under tax schedule* $T^*(y)$*, or optimal under tax schedule* $T(y,\gamma)$ *but not under* $T^*(y)$*, faces strictly lower taxes:*
$\exists \bar{y} \in (Y^{\gamma_k}_{T^*(y)} \cup Y^{\gamma_k}_{T(y,\gamma)/T^*(y)})\ T(\bar{y},\gamma_k) < T^*(\bar{y})$*;*

b. *While no revenue is lost from those agents (with characteristic* γ_k *at income level* $\bar{y}$*):*
$\mathbb{E}[\sum_{\{i:(\gamma^i=\gamma_k)\cap(\bar{y}\in y^i_{T^*}\cup(\bar{y}\in y^i_{T(y,\gamma)}\cap\bar{y}\notin y^i_{T^*}))\}}(T(y^i_*,\gamma^i) - T^*(y^i_*))] \ge 0.$

The reason definition 5 relies on definition 4 is twofold. First, to make the term Laffer effect only apply to incomes that are actually part of an optimum, either under the new or the old tax schedule. Second, to make the term only apply to new tax schedules that are meaningfully lower, in the sense they should make agents better off (rather than just expanding their existing optimum). This brings us to Proposition 1.

Proposition 1 (A tag is Pareto improving if and only if it identifies Laffer effects)**.** *A tax schedule* $T(y,\gamma)$ *that results in a Pareto improvement over tax schedule* $T^*(y)$ *exists if and only if some characteristic* γ_k *identifies a Laffer effect.*

Proof. I start with the *only if* part: that a Pareto Improvement implies identifying a Laffer effect. The proof starts from tax schedule $T^*(y)$, and establishes which changes to the schedule can result in a Pareto improvement. Since $T^*(y)$ is the optimum when the tax schedule does not depend on γ, any such change must apply to some generic γ_k. As it turns out, achieving a Pareto improvement requires that this γ_k identifies a Laffer effect. I proceed by demonstrating that a Pareto improvement implies parts a and b of definition 5.

1. *Some tax rate must be reduced.*
   Which changes to $T^*(y)$ result in a Pareto improvement for generic γ_k? This can be done in two ways. First, by simply reducing the tax rate where it already applies under $T^*(y)$. This results in higher consumption, and because welfare is strictly increasing in consumption, the agent has been made strictly better off. Second, a tax rate could be reduced elsewhere such that the agent's optimum changes. This only results in an welfare improvement if the new optimum and the old optimum do not overlap. Otherwise, the optimum may simply have been expanded by new income levels that result in the same welfare level. If however the new and the old optimum do not overlap, then a strict welfare improvement must have taken place: the income levels of the old optimum are still available, and taxes rates for those are the same or lower than they were.

Thus, only reducing some tax rate at income levels optimal to an agent with characteristic γ_k under $T^*(y)$, or under $T(y,\gamma)$ but not under $T^*(y)$, makes some agent better off. This directly implies part a of definition 5.

2. *No revenue can be lost.*
   The new tax schedule is only valid if expected tax revenues are increased or kept the same. Otherwise, the budget constraint would be violated and the resulting tax schedule would not be a solution to the planner's problem. This implies part b of definition 5, unless taxes received from other agents can be increased relative to $T^*(y)$. But doing so would clearly make those agents worse off.

The *if* part, that identifying a Laffer effect implies a Pareto Improvement, is obvious. If we implement the tax schedule $T(y,\gamma)$ that results from the argument above, then (by construction) no agent is worse off, and some agent is better off. Thus, if an observed γ_k exists that permits such a tax schedule, then there is tax schedule that results in a Pareto improvement over $T^*(y)$. □

Proposition 1 shows how one can improve upon an already optimal system by introducing tags: essentially, one needs to identify a group for which, on average, a lower tax rate than the previously optimal one does not result in a loss of taxes. This is similar to how one identifies Pareto improvements in any given suboptimal tax schedule without the use of tags: one identifies Laffer effects, raising the same or more taxes using lower effective tax rates. The proposition also suggests an easy test for whether tags are Pareto improving or not: this is equivalent to asking whether they identify a group for which lower taxes lead to equal or more income.

Whether such Laffer affects can be identified very much depends on the tag to be used. Do we believe there is some height group for which a Laffer effect is possible on average? Mankiw and Weinzierl (2010) do report a Pareto improving use of height as a tag. With the increasing availibility of data, finding such tags seems increasingly likely.

Suppose that the finding in Mankiw and Weinzierl (2010) is correct, and height-based taxation can indeed be used to achieve a Pareto improvement. Then, recasting the proposition above in the manner of Kaplow and Shavell (2001): *Any method of tax policy assessment violates the Pareto Principle if it does not discriminate on the basis of height.* In other words, the relation between discrimination and taxation is inherent to Welfarism. This notion is formalized in the following corollary.

Corollary 1 (Welfarism implies using any tags that identify Laffer effects)**.** *Suppose the planner maximizes the Welfarist Criterion and observes characteristic γ. Suppose also that* *some characteristic γ_k identifies a Laffer effect. Then the resulting optimal tax schedule $T^*(y, \gamma)$ is not equal to $T'(y)$ for any function T' that does not depend on γ.*

Proof. Suppose $T^*(y, \gamma) = T'(y)$ for some T' that does not depend on γ. Obviously, we must have $T'(y) = T^*(y)$, i.e. it is a tax function that solves the planner's problem without knowledge of the tag. Now, according to Proposition 1 a Pareto improvement is possible, given what we suppose on $T(y, \gamma_k)$. But then, the Welfarist Criterion cannot have been maximized because the criterion attaches a positive weight to each agent, and at the same time we can improve welfare for some agent while keeping the welfare of all others at least the same:

$$\exists i : u^i(\theta^i n^i_{T(y,\gamma_k)} - T(\theta^i n^i_{T(y,\gamma_k)}, \gamma^i), n^i_{T(y,\gamma_k)}) > u^i(\theta^i n^i_{T^*(y)} - T(\theta^i n^i_{T^*(y)}, \gamma^i), n^i_{T^*(y)}),$$
$$\forall i : u^i(\theta^i n^i_{T(y,\gamma_k)} - T(\theta^i n^i_{T(y,\gamma_k)}, \gamma^i), n^i_{T(y,\gamma_k)}) \geq u^i(\theta^i n^i_{T^*(y)} - T(\theta^i n^i_{T^*(y)}, \gamma^i), n^i_{T^*(y)}).$$

□

Because the corollary essentially just combines the proposition above with the observation that Welfarism respects the Pareto criterion, it extends to any other objective that respects the Pareto criterion. This includes the Rawlsian leximin criterion, and certain mixed objectives such as the one by Weinzierl (2014).

The main argument for Welfarism, and the main objection against horizontal equity concerns in optimal taxation, has always been the potential to violate the Pareto principle: if we can make everyone better off, why would we not do so? The above result raises the reverse question: if we really object to certain forms of discrimination, should we adhere so strictly to Welfarism and the Pareto principle?

3.3 Horizontal Equity

I now make an attempt to formalize the intuitive aversion to tags such as those based on height. In doing so, I stay as close as possible to the Welfarist objective, while at the same time introducing a minimal notion of horizontal equity.

The concept of horizontal equity, which is often casually introduced as *the equal treatment of equals*, has gained little acceptance so far. Despite its intuitive appeal, the concept suffers from the major critique that it's operationalization seems to require a choice of some natural state against which to compare allocations, a point forcefully made by Kaplow (1989). Indeed, the previous literature has made such choices. For example, to introduce the notion of Equal Sacrifice, an objective that competes with the Utilitarian criterion, Weinzierl (2014) needs some starting point from which to calculate an individual's sacrifice. He chooses the allocation in which there is no government intervention, but at the same time notes: *"A well-known conceptual issue with the idea of the laissez-faire allocation is that any economy is, in reality, inseparable from the government and state institutions that taxes fund. The laissez-faire allocation is, therefore, not well-defined, because [the absence of taxation] implies a very different economy than the status quo"* (p. 137). The seminal paper of Fleurbaey and Maniquet (2006) similarly defines a *laissez-faire allocation*, which is equivalent to one without taxation, and uses it to introduce a notion of horizontal equity. In the constrained Welfarist criterion below, I avoid this issue entirely by defining horizontal equity with respect to other individuals only, irrespective of the allocation.

For given tax function T, denote by $x^i_{i'}$ the level of choice variable x for agent i when agent i chooses the level that is optimal for agent i' (so that $x^{i'}_{i'} = x^{i'}_*$).

Definition 6 (Welfarist Criterion with Horizontal Equity)**.**

$$\begin{aligned} U(\{u^i(c^i_*, n^i_*)\}_{i\in I}) =& \mathbb{E}\sum_{i\in I} w^i u^i(c^i_*, n^i_*), \\ & \forall i: w^i > 0, \\ & \forall i, i': (\theta^i n^i_{i'} = y^i_{i'}) \implies u^i(y^i_{i'} - T(y^i_{i'}, \gamma^i), n^i_{i'}) \geq u^{i'}(y^{i'}_* - T(y^{i'}_*, \gamma^{i'}), n^{i'}_*). \end{aligned}$$

The Welfarist criterion has been restricted: Suppose a generic agent i could produce the same as agent i' using the same input. Then in doing so he is to obtain at least the level of welfare that agent i' obtains. Utility functions are still interpreted as before: cardinality can simply be a normative view the planner holds. Also note that the criterion involves the evaluation of a counterfactual: the welfare an agent obtains when counterfactually behaving like another agent chooses to behave.

The restriction is purposefully minimal. An agent of higher ability does not have the right to the same tax schedule as one of lower ability: the criterion only applies to agents who have the same ability.[7] Agents of the same ability are not necessarily treated the same: someone who cannot provide the same level of n as others with his ability is not subject to this criterion: it is only defined for those who can imitate both behavior and production. Apart from those cases, the planner may still prefer one group to another by assigning different welfare weights. At last, individuals with clearly distinct preferences (whether cardinal or relative) may be treated differentially. Nevertheless, this minimal criterion prescribes *equal treatment of equals*, as the following proposition shows.

Proposition 2 (Equal Treatment of Equals)**.** *Suppose, for two agents i and i', we have $u^i(\cdot) = u^{i'}(\cdot) = u(\cdot)$ and $\theta^i = \theta^{i'}$. Then under the Welfarist Criterion with Horizontal*

[7]One could turn the extra condition in the constrained criterion into a restriction on rank reversals, by replacing $(\theta^i n^i_{i'} = y^i_{i'})$ by $(\theta^i n^i_{i'} \geq y^i_{i'})$. Much of the following would remain unchanged.

Equity, we must have $T(y_*^i, \gamma^i) = T(y_*^i, \gamma^{i'}) = T(y_*^i)$ and $y_*^i = y_*^{i'}$, i.e. taxes paid at the optimal choice of either agent do not depend on individual characteristics γ^i and $\gamma^{i'}$.

Proof. Take two agents i and i' for which $u^i(\cdot) = u^{i'}(\cdot) = u(\cdot)$ and $\theta^i = \theta^{i'}$. Since we have that $(\theta^i n_{i'}^i = y_{i'}^i)$ and $(\theta^{i'} n_i^{i'} = y_i^{i'})$, by the restriction in our criterion we must have both $u(y_{i'} - T(y_{i'}, \gamma^i), n_{i'}) \geq u(y_{i'} - T(y_{i'}, \gamma^{i'}), n_{i'})$ and $u(y_i - T(y_i, \gamma^{i'}), n_i) \geq u(y_i - T(y_i, \gamma^i), n_i)$. (The equivalent notation $x_i^i = x_*^i$ is used. Superscripts are irrelevant other than for γ, as we are dealing with two otherwise identical agents.) At the same time, because n^i is an optimal choice for agent i and $n^{i'}$ is an optimal choice for agent i' given T, we also have $u(y_{i'} - T(y_{i'}, \gamma^{i'}), n_{i'}) \geq u(y_i - T(y_i, \gamma^{i'}), n_i)$ and $u(y_{i'} - T(y_{i'}, \gamma^i), n_{i'}) \leq u(y_i - T(y_i, \gamma^i), n_i)$. Combining results in equality of all four evaluations of u. Then, $T(y_i, \gamma^i) = T(y_i, \gamma^{i'}) = T(y_i)$, $T(y_{i'}, \gamma^i) = T(y_{i'}, \gamma^{i'}) = T(y_{i'})$, and thus $y_i = y_{i'}$. □

A direct consequence of this is that certain tags cannot be used: no matter how informative a tag is of a group's abilities, if there are equals *everywhere* along the optimal choices across the groups identified by a tag, then this tag cannot be used. By the example of height: if for every ability level there is potentially at least one tall and one otherwise equal short person, then height is ruled out as a tag altogether. I will call such tags 'diffuse'. The below definition and proposition formalize the notion.

Definition 7 (Diffuse Tags)**.** *A Diffuse Tag is a tag at any value of which an otherwise equal agent exists with a different value for that tag. In other words, observing the tag renders the planner inconclusive of an agent's $u^i(\cdot)$ or θ. More formally:*
Call Diffuse Tags any characteristics represented by γ, such that whenever there exists an agent, say i, who holds some level of γ, say γ^i, then there potentially exists another agent, say i', such that $\gamma_i \neq \gamma_{i'}$, but $u^i(\cdot) = u^{i'}(\cdot)$ and $\theta^i = \theta^{i'}$.

Proposition 3 (Exclusion of Diffuse Tags)**.** *A planner who chooses amongst allocations according to the Welfarist Criterion with Horizontal Equity will not use Diffuse Tags. I.e. when γ is a Diffuse Tag, we have $T(y, \gamma) = T(y)$ for any observed level of y.*

Proof. This follows directly from Definition 7 and Proposition 2. Take a generic level of y that is optimal to some agent i with characteristics represented by γ^i, and call this y_*^i. Then according to Definition 7, there potentially exists another agent i' with a different value for γ ($\gamma_i \neq \gamma_{i'}$), but with $u^i(\cdot) = u^{i'}(\cdot)$ and $\theta^i = \theta^{i'}$, so that his optimal income is the same ($y_*^{i'} = y_*^i$). Thus, the tax schedules that apply to them are $T(y_*^i, \gamma^i)$ and $T(y_*^i, \gamma^{i'})$, respectively. According to Proposition 2, we must have $T(y_*^i, \gamma^i) = T(y_*^i, \gamma^{i'}) = T(y_*^i)$. Since we picked a generic level of y was that was optimal to some agent, this holds for all observed y, so that we have $T(y, \gamma) = T(y)$. □

Note that this result does not depend on our interpretation of the u^i. Even if a diffuse tag is correlated with preferences, it is ruled out. At the same time, if there are surely no equals somewhere in the sense of the Horizontal Equity restriction, then a tag can be used at least there.

3.4 Dynamic Optimal Taxation

So far I have discussed a static Mirrleesian environment. A growing literature discusses optimal taxation over a life-cycle, say an agent lives from age 0 to age T, where the nature of the problem is different because an agent's earnings ability $\theta^{i,t}$ may change over time, and the agent can self-insure through asset holdings a:

$$a^{t+1} = a^t + y^t - T^t(y^t, a^t, \gamma^t, \{y^s\}_{s=0}^{t-1}, \{a^s\}_{s=0}^{t-1}) - c^t,$$

where a^0 is given. Asset holdings are typically considered observable to the planner. The planner maximizes a standard dynamic Welfarist criterion, as presented below, by taking into account the history of incomes of an individual. This is because past incomes contain information on current and future ability, unless the ability process is independent over time.

Definition 8 (Dynamic Welfarist Criterion)**.**

$$\begin{aligned} U(\{\sum_{t=0}^{T} u^{i,t}(c_*^{i,t}, n_*^{i,t})\}_{i\in I}) =& \mathbb{E}\sum_{i\in I} w^{i,t}\Big(\sum_{t=0}^{T} u^{i,t}(c_*^{i,t}, n_*^{i,t})\Big), \\ & \forall t \forall i : w^{i,t} > 0. \end{aligned}$$

Introducing horizontal equity into this criterion can be done in several ways, which is discussed below. I propose the following. (For given tax function T, denote by $x_{i',t}^{i,t}$ the level of choice variable x for agent i of age t when agent i chooses the level that is optimal for agent i' of age t (so that $x_{i,t}^{i,t} = x_*^{i,t}$).)

Definition 9 (Dynamic Welfarist Criterion with Horizontal Equity.)**.**

$$\begin{aligned} U(\{\sum_{t=0}^{T} u^{i,t}(c_*^{i,t}, n_*^{i,t})\}_{i\in I}) =& \mathbb{E}\sum_{i\in I} w^{i,t}\Big(\sum_{t=0}^{T} u^{i,t}(c_*^{i,t}, n_*^{i,t})\Big), \\ & \forall t \forall i : w^{i,t} > 0, \\ & \forall t \forall i, i' : (a^{i,t} = a^{i',t}) \& (\theta^{i,t} n_{i',t}^{i,t} = y_{i',t}^{i,t}) \implies \\ & \quad u^{i,t}(y_{i',t}^{i,t} - T^t(y_{i',t}^{i,t}, a^{i,t}, \gamma^{i,t}, \{y^{i,s}\}_{s=0}^{t-1}, \{a^{i,s}\}_{s=0}^{t-1}) - a_{i',t}^{i,t+1} + a^{i,t}, n_{i',t}^{i,t}) \geq \\ & \quad u^{i',t}(y_*^{i',t} - T^t(y_*^{i',t}, a^{i',t}, \gamma^{i',t}, \{y^{i',s}\}_{s=0}^{t-1}, \{a^{i',s}\}_{s=0}^{t-1}) - a_*^{i',t+1} + a^{i',t}, n_*^{i',t}). \end{aligned}$$

This seems to be the most straightforward extension of the idea of horizontal equity to a dynamic environment. An agent who at some age has the same wealth, behaves the same (both in his labor and asset-consumption choice), and produces the same as another agent of that age, should not be worse off than that other agent. The following corollary shows an implication of this choice.

Corollary 2 (Taxes do not depend on diffuse tags, including Past Incomes or Past Assets)**.** *Assume that whenever there exists an agent i of age t, with history and characteristics $\{\gamma^{i,t}, \{y^{i,s}\}_{s=0}^{t-1}, \{a^{i,s}\}_{s=0}^{t-1}\}$, then there exists another agent of age t, say i', such that $\gamma^{i,t} \neq \gamma^{i',t}$ but $a^{i,t} = a^{i',t}$, $u^{i,t} = u^{i',t}$, and $\theta^{i,t} = \theta^{i',t}$. What holds true for γ for some two agents, also holds for $y^{i,s}$ when $0 \leq s \leq t-1$, and $a^{i,s}$ when $0 \leq s \leq t-1$.*
Then we have $T^t(y^t, a^t, \gamma^t, \{y^s\}_{s=0}^{t-1}, \{a^s\}_{s=0}^{t-1}) = T^t(y^t, a^t)$ for any observed level of y^t and a^t.

Proof. This follows from straightforward re-application of the proofs of propositions 2 and 3. $\square$

Under the specification chosen here, age dependent taxation is still permitted. Age-dependent taxation has often been presented as a simplifying short-cut to get most of the benefit of dynamic taxation without the complicated policies that come with it. Examples include Weinzierl (2011), Farhi and Werning (2013), and Stantcheva (2017). Under the criterion we present here, such proposals actually become optimal policies.

Other specifications that carry the principle of Horizontal Equity from the static setting into a dynamic one are possible. For example, we could compare all ages at every period, instead of only those of the same age. In that case, age-dependent taxation would have been excluded. We could also compare over entire life-cycles, in which case both age-dependent and income-history dependent taxation would be acceptable. Comparing life-cycles would take some notion of life-cycle effort and production in order to define horizontal equity if we do not just want to consider different income paths incomparable (in which case Horizontal Equity would have no bearing on past incomes and assets). It is not clear how to set taxes by period in that case, and the result would still be far from the prescriptions of the dynamic optimal taxation literature. The version proposed here seems closest in nature to the original criterion, and coincidentally produces the best explanation of the policies we observe. While it is somewhat harder to operationalize Horizontal Equity in a dynamic environment, one conclusion seems worth drawing: If we accept horizontal equity as relevant for taxation, then this has implications for the dynamic optimal taxation literature.

3.5 Solving Optimal Taxation Problems under Horizontal Equity

How would one solve optimal taxation problems under the criteria suggested here? This clearly depends on the tags one wants to introduce to the environment. Nevertheless, the case for diffuse tags is clear: they should not be considered by the planner. There are generally two ways of achieving this. The first is to not make them part of the planner's information set. In that sense, many results from the literature on optimal taxation simply remain valid in the presence of diffuse tags, under the constrained criterion. The second is to take the tag as being randomly distributed over the population, so that it contains no information value.

The second approach is particularly relevant for dynamic optimal taxation, where Albanesi and Sleet (2006) provide some results for the case where abilities evolve independently over time. That problem is equivalent to solving a taxation problem with abilities that are dependent over time but diffuse under horizontal equity.[8] Unsurprisingly, Albanesi and Sleet (2006) find that optimal taxation can be implemented without making taxes dependent on past income or past assets.

4 Empirical Implications

In this section, I show that introducing horizontal equity as defined in this paper results in sharp implications that match what we observe. The correspondence between the criterion's implications and the US tax code is close, in two ways: First, there is a close link between tags that are used and not used under the criterion and in reality. Second, excluded tags are not just used to a lesser extent, but entirely excluded on the basis of principle.

Table 1 lists a large number of tags that are either part of the US or other countries' income tax code, or have been discussed in the academic literature, or have been the subject of public debate. Dynamic tags have been separated from those that are also relevant in a static setting. For each tag, I indicate whether it is used in the US tax code. Finally, I provide a reference to relevant academic literature. Under a standard Welfarist criterion, all tags are admissible as long as they provide information on any variable or function relevant to the planner's problem. The restrictions of horizontal equity, on the other hand, prescribe a usage of tags that closely corresponds to reality. The classification of tags under horizontal equity is discussed below. The discussion is informal, in that many tags imply slightly different environments from the ones presented above. Guiding principles follow from the theoretical analysis above: agents of the same ability may be treated differently if they cannot or choose not to imitate each other's effort and productivity, or if the planner assigns a different cardinality to them. Diffuse tags are ruled out.

A caveat applies to the exercise that follows. I cannot actually check whether the criterion is violated by the actual tax code, but merely whether the tag that the tax code uses would

[8]The same cannot be said about numerical results, where the evolution of ability should still resemble the actual process.

Static		
Tag	*Observed*	*Reference*
Income	Yes	Mirrlees (1971)
Consumption	Yes	Atkinson and Stiglitz (1976)
Unemployment	Yes	Saez (2002)
Disability	Yes	Diamond and Mirrlees (1978)
Blindness	Yes	Weinzierl (2014)
Household Form	Yes	Kleven, Kreiner, and Saez (2009)
Children	Yes	Domeij and Klein (2013)
Mortgage Interest	Yes	–
Charitable Contributions	Yes	Blumkin and Sadka (2007)
Health Expenditure	Yes	–
Gender	No	Alesina, Ichino, and Karabarbounis (2011)
Height	No	Mankiw and Weinzierl (2010)
Race	No	Blumkin, Margalioth, and Sadka (2009)
Genetics	No	Logue and Selmrod (2008)
Other Diffuse Tags	No	Mankiw, Weinzierl, and Yagan (2009)
Dynamic		
Tag	*Observed*	*Reference*
Income	Yes	Golosov, Kocherlakota, and Tsyvinski (2003)
Past Income	No	Kocherlakota (2005)
Assets	Yes	Golosov, Kocherlakota, and Tsyvinski (2003)
Past Assets	No	–
Income Averaging	No	Vickrey (1939)
Age	Yes	Weinzierl (2011)
Education	No	Bohacek and Kapicka (2008)
Income-contingent Loans	Yes	Stantcheva (2017)

Table 1: Tagging under Horizontal Equity

possibly be allowed under the criterion. For example, the blind or disabled may receive higher benefits or face lower taxes than those who are otherwise equal but not blind or disabled. One then has to infer which cardinal welfare levels the planner attributes to which individuals, since these are unobserved in practice. This introduces a degree of freedom in the exercise that follows. However, the problem is smaller than it seems. First, the planner's preferences may to extent correspond to shared and commonly known preferences of individuals. Second, behavior leaves clues as to these shared preferences. Consider the following observation: there do not seem to be many cases in which people attempt to blind or disable themselves because they would be better off receiving benefits as a blind or disabled person. Thus, the social planner's choices seem to provide them with more welfare than an otherwise equivalent blind or disabled person, despite the latter possibly receiving higher benefits. The tag, in short, appears to be compensation for differences in welfare at given consumption and effort levels.

4.1 Static Tags

Income is the classic basis for taxation, and can be seen as a tag for ability in the Mirrleesian setting. Consumption taxation is equivalent to income taxation in that setting, and Atkinson and Stiglitz (1976) show that Horizontal Equity (in a version somewhat different from mine) in combination with differences in tastes does not imply uniform taxation. Differential treatment requires differential choices, so that these tags are allowed under Horizontal Equity.

Unemployment benefits are typically intended for the involuntarily unemployed. Proof and reaffirmation of involuntariness are often required for unemployment benefits. This means it is intended as a precise tag of earnings inability ($\theta = 0$), which the Horizontal Equity criterion admits regardless of the height of the benefit. Discussions around the fairness of these benefits are indeed typically centered around whether beneficiaries are really unable to earn an income.

Disability and blindness benefits can increase the consumption levels of the blind and disabled above those who have the same earnings ability and work the same hours, but are clearly also a tag of circumstance: those who are disabled or blind have a natural disadvantage in life, deriving lower welfare from the same economic circumstances. As argued above, benefit levels are typically low enough so that no-one chooses blindness. A difference in assigned cardinality (differences in u^i) therefore seems credible, so that Horizontal Equity permits these tags.

Research on the role of household form and the number of children in taxation is ongoing, so that it is hard to draw definite conclusions on their implications for horizontal equity as defined here. When welfare is treated at the household level, then the welfare of a household may well be assigned different cardinality by the planner depending on its composition. Issues such as home production and household returns to scale may also affect welfare levels. Overall, it is plausible that tags for household composition are allowed under our criterion. Similar observations hold for the number of children. In addition, when treating both tags in a purely static sense, they may be seen as choices.

Mortgage interest deductions, exemptions for charitable contributions, and exemptions for health expenditures are less discussed in the academic literature, but all part of the tax code. The former two clearly relate to choices. Therefore these tags do not violate the narrow

Horizontal Equity restriction: any agent who behaves like another is treated the same. The latter should be seen similar to tags of blindness and disability: these tags are admitted because of differences in cardinality.

Finally we arrive at a large number of tags that are not used in practice, and are clearly diffuse in the sense of this paper. These include gender, height, race, genetic information, and many others. Mankiw, Weinzierl, and Yagan (2009) additionally list the following: skin color, physical attractiveness, and parents' education.

4.2 Dynamic Tags

For dynamic tags, I use the constrained criterion as above, where agents are compared to other agents of the same age on a per-period basis. As already discussed, while income and assets can be used, past income and past assets are excluded under horizontal equity. Age-based taxation is allowed, although I have already qualified earlier the extent to which this depends on how the criterion is carried into a dynamic setting.

Focusing on life-cycle comparisons instead of the criterion proposed here would address another aspect of tax systems that is sometimes considered unfair: volatile incomes are, under progressive tax systems, taxed more heavily than less volatile ones. This issue was first addressed by Vickrey (1939). Pensions also often depend on past pay-in (in addition to age), which makes them a form of lifetime income taxation. On the other hand, this choice would rule out narrowly age-dependent schemes. In fact, while the US tax code does make taxes somewhat dependent on age, it does not do so too sharply: rather, it makes exceptions for the young (be it through supplementary systems such as schooling, federal student loans, etcetera) and the elderly (in the code as well as through retirement programs). The same goes for the tax code of many other countries. On the other hand, while the US tax code no longer provides for income averaging to smooth out the effect of volatile incomes (it did in the past), such provisions are provided in the tax codes of other countries.

Seeing all this, it seems that there is indeed concern for horizontal equity that allows for a distinction by age, albeit perhaps not as sharply as a year-by-year scheme prescribes: income averaging, where it is permitted, is often very local (not over the entire life cycle but over a few consecutive years), while age-based taxation is applied to much longer phases of the life-cycle than by years. It appears that concern for horizontal equity is not quite over the entire life cycle, but also not quite limited to single years. In some sense, the latter was to be expected: there is no inherent value to age as measured in years, we just want to compare *sufficiently similar* agents.

A tag on education would be diffuse as well, and therefore excluded. Some literature on dynamic optimal taxation with human capital suggests conditioning the tax code on education. This paper suggests that we are unlikely to observe such policies in practice, while we do instead see income-contingent student loans, which would be permissible under Horizontal Equity.

5 Conclusion

The absence of the use of tags is arguably the biggest difference between the practice and the theory of taxation today. As this paper shows, an adherence to Pareto optimality implies using tags if and only if they identify Laffer effects.

This paper also shows that the concept of Horizontal Equity can reconcile the overall success of the Welfarist framework with the missing use of tags that it prescribes. In doing so, the paper has found an operationalization of Horizontal Equity in a Welfarist framework that does not rely on a 'natural state'.

As this paper shows, conclusions from the literature on optimal taxation will vary with the assumptions that are made regarding social preferences. Further positive research on taxation may therefore make important contributions by providing a better understanding of actual social preferences.

References

Akerlof, G. A. (1978). The economics of" tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *The American Economic Review*, 8–19.

Albanesi, S. and C. Sleet (2006). Dynamic optimal taxation with private information. *The Review of Economic Studies 73*(1), 1–30.

Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy 3*(2), 1–40.

Atkinson, A. B. and J. E. Stiglitz (1976). The design of tax structure: direct versus indirect taxation. *Journal of Public Economics 6*(1), 55–75.

Auerbach, A. J. and K. A. Hassett (2002). A new measure of horizontal equity. *American Economic Review 92*(4), 1116–1125.

Blumkin, T., Y. Margalioth, and E. Sadka (2009). Incorporating affirmative action into the welfare state. *Journal of Public Economics 93*(9), 1027–1035.

Blumkin, T. and E. Sadka (2007). A case for taxing charitable donations. *Journal of Public Economics 91*(7), 1555–1564.

Boadway, R. and P. Pestieau (2006). Tagging and redistributive taxation. *Annales d'Economie et de Statistique*, 123–147.

Bohacek, R. and M. Kapicka (2008). Optimal human capital policies. *Journal of Monetary Economics 55*(1), 1–16.

Cremer, H., F. Gahvari, and J.-M. Lozachmeur (2010). Tagging and income taxation: theory and an application. *American Economic Journal: Economic Policy 2*(1), 31–50.

Diamond, P. and E. Saez (2011, December). The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives 25*(4), 165–90.

Diamond, P. A. and J. A. Mirrlees (1978). A model of social insurance with variable retirement. *Journal of Public Economics 10*(3), 295–336.

Domeij, D. and P. Klein (2013). Should day care be subsidized? *The Review of Economic Studies 80*(2), 568–595.

Farhi, E. and I. Werning (2013). Insurance and taxation over the life cycle. *The Review of Economic Studies 80*(2), 596–635.

Fleurbaey, M. and F. Maniquet (2006). Fair income tax. *Review of Economic studies*, 55–83.

Golosov, M., N. Kocherlakota, and A. Tsyvinski (2003). Optimal indirect and capital taxation. *The Review of Economic Studies 70*(3), 569–587.

Hendren, N. (2014). The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function.

Immonen, R., R. Kanbur, M. Keen, and M. Tuomala (1998). Tagging and taxing: The optimal use of categorical and income information in designing tax/transfer schemes. *Economica 65*(258), 179–192.

Jacquet, L. and E. Lehmann (2015). Optimal Income Taxation when Skills and Behavioral Elasticities are Heterogeneous.

Kaplow, L. (1989). Horizontal Equity: Measures in Search of a Principle. *National Tax Journal*, 139–154.

Kaplow, L. (2000). Horizontal equity: New measures, unclear principles.

Kaplow, L. and S. Shavell (2001). Any non-welfarist method of policy assessment violates the pareto principle. *Journal of Political Economy 109*(2), 281–286.

Kleven, H. J., C. T. Kreiner, and E. Saez (2009). The optimal income taxation of couples. *Econometrica 77*(2), 537–560.

Kocherlakota, N. R. (2005). Zero expected wealth taxes: A Mirrlees approach to dynamic optimal taxation. *Econometrica 73*(5), 1587–1621.

Logue, K. and J. Selmrod (2008). Genes as Tags: The Tax Implications of Widely Available Genetic Information. *National Tax Journal*, 843–863.

Mankiw, N. G. and M. Weinzierl (2010). The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution. *American Economic Journal. Economic Policy 2*(1), 155.

Mankiw, N. G., M. Weinzierl, and D. Yagan (2009). Optimal Taxation in Theory and Practice. *Journal of Economic Perspectives 23*(4), 147–74.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, 175–208.

Musgrave, R. A. (1990). Horizontal equity, once more. *National Tax Journal*, 113–122.

Piacquadio, P. G. (2017). A fairness justification of utilitarianism. *Econometrica 85*(4), 1261–1276.

Saez, E. (2002). Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses. *Quarterly Journal of Economics 117*, 1039–1073.

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

Sausgruber, R. and J.-R. Tyran (2014). Discriminatory taxes are unpopular Even when they are efficient and distributionally fair. *Journal of Economic Behavior & Organization 108*, 463–476.

Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy 125*(6), 1931–1990.

Viard, A. D. (2001a). Optimal categorical transfer payments: The welfare economics of limited lump-sum redistribution. *Journal of Public Economic Theory 3*(4), 483–500.

Viard, A. D. (2001b). Some results on the comparative statics of optimal categorical transfer payments.

Public Finance Review 29(2), 148–180.

Vickrey, W. (1939). Averaging of income for income-tax purposes. *The Journal of Political Economy*, 379–397.

Weinzierl, M. (2011). The surprising power of age-dependent taxes. *The Review of Economic Studies 78*(4), 1490–1518.

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

Werning, I. (2007). Pareto efficient income taxation. *Massachusetts Institute of Technology, Manuscript 44*, 45.