



# Popular Acceptance of Morally Arbitrary Luck and Widespread Support for Classical Benefit-Based Taxation

## Citation

Weinzierl, Matthew C. "Popular Acceptance of Morally Arbitrary Luck and Widespread Support for Classical Benefit-Based Taxation." Harvard Business School Working Paper, No. 16-104, March 2016.

## Link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:26274606>

## Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles (OAP), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

## Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

## Share Your Story

The Harvard community has made this article openly available.

Please share how this access benefits you. [Submit a story](#)

# Popular Acceptance of Morally Arbitrary Luck and Widespread Support for Classical Benefit-Based Taxation

Matthew Weinzierl

Working Paper 16-104



# Popular Acceptance of Morally Arbitrary Luck and Widespread Support for Classical Benefit-Based Taxation

Matthew Weinzierl  
Harvard Business School

**Working Paper 16-104**

Copyright © 2016 by Matthew Weinzierl

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

# Popular Acceptance of Morally Arbitrary Luck and Widespread Support for Classical Benefit-Based Taxation

Matthew Weinzierl\*

Preliminary draft as of March 25, 2016; comments very welcome

## Abstract

U.S. survey respondents' views on distributive justice are shown to differ in two specific, related ways from what is conventionally assumed in modern optimal tax theory. A large share, and in some cases a large majority, of respondents resist full equalization of economic outcomes determined entirely by luck. A similar share prefer a justification of tax progressivity that relies on a benefit-based logic rather than diminishing marginal social welfare of income, the conventional logic. Moreover, these two views are linked: respondents who more strongly resist redistribution are more likely to prefer the benefit-based principle. Together, these results raise the possibility that the American public views the allocations of taxes and pre-tax outcomes as morally relevant, a judgment that is inconsistent with conventional consequentialist objectives.

---

\*Harvard Business School and NBER, mweinzierl@hbs.edu. Thanks to Max Bazerman, Rafael DiTella, Louis Kaplow, Benjamin B. Lockwood, David Moss, Michael Norton, Julio Rotemberg, Florian Scheuer, Hannah Shaffer, and Stefanie Stantcheva for insightful discussions.

## Introduction

This paper uses novel survey evidence to identify two specific, related ways in which the U.S. public's views on distributive justice appear to differ from the normative view assumed in modern research on optimal taxation. An understanding of the moral reasoning behind public opinion is useful for explaining existing tax policy, but it is also essential for tax scholars, advisors, or policymakers who want their theoretical frameworks and recommended tax reforms to be consonant with the public's views.

First, a large share of survey respondents resist the full equalization of luck-based outcomes that conventional optimal tax models would strongly recommend. Specifically, when presented with a hypothetical situation meant to mimic the tax policy problem, between 50% and 95% of respondents—depending on the details of the setup—choose not to fully offset unequal outcomes even when there are neither efficiency costs of redistribution nor differences in desert across individuals. These choices suggest that the conventional model's allowance for two brakes on redistribution—the need to satisfy incentive constraints and the possibility that tastes rather than abilities determine income—omits an underlying reason for which a large share of survey respondents, and perhaps Americans, resist it.

Second, a similarly large share of respondents express affinity for a principle of optimal tax design fundamentally at odds with what conventional models assume. When asked to justify progressivity in the tax scenario described above, between 60% and 79% of respondents prefer a logic based not on the conventional appeal to diminishing marginal social welfare of income but rather on one tied to a centuries-old idea that was Adam Smith's first maxim of taxation and that Richard Musgrave (1959) named classical benefit-based taxation (CBBT). Under CBBT, taxes are assigned based on the benefit a taxpayer obtains from the activities of the state, with benefit being measured by the state's role in increasing the taxpayer's income. CBBT has a long history in public debate over taxes in the United States, from its use as a justification for the new personal income tax in 1913 to its use by presidents Franklin Delano Roosevelt and Barack Obama to advocate for increases to tax progressivity.

Moreover, these two findings are linked, in that those respondents more willing to accept luck-based inequality are also more likely to prefer the benefit-based principle. One interpretation of this linkage is that a large majority of respondents appear to believe that individuals are entitled—at least in part—to morally arbitrary outcomes, and that taxes ought substantially to respect that entitlement. Advocates of benefit-based taxation stress exactly this idea when asserting its normative appeal as a voluntary rather than coercive system, in that under benefit-based taxation a taxpayer funds social goods only to the extent that he or she benefits from them, paralleling the case of voluntary exchange in private markets.<sup>1</sup>

This set of views stands in stark contrast to a key normative assumption in conventional tax theory: pre-tax outcomes are morally irrelevant. As formalized first by James Mirrlees (1971),

---

<sup>1</sup>Musgrave (1959), page 74. See Musgrave's chapter 4 for his full treatment of the benefit approach.

modern optimal tax research adopts an objective put forth by John Harsanyi (1953, 1955); namely, that taxation ought to be designed to maximize (generalized utilitarian) social welfare. Though such a consequentialist objective can in principle accommodate a wide range of judgments, almost all applications of the theory embrace what Liam Murphy and Thomas Nagel (2003) deem "the right way, investigating outcomes rather than the distribution of [tax] burdens." That is, under the standard approach neither taxes nor pre-tax outcomes have any effect on welfare—or relevance to optimal policy—other than through their (mechanical) implications for after-tax allocations. In contrast, a large share of respondents to this paper's survey appear willing to grant some moral entitlement to pre-tax outcomes and support CBBT, a principle that ignores after-tax outcomes and *defines* optimality in terms of the relationship between taxes and pre-tax outcomes.

It is important to clarify up front that the results of this paper are entirely consistent with other principles, e.g., the utilitarianism commonly used in optimal tax theory, also playing important roles in Americans' appraisals of tax policy. In fact, nearly half of the respondents to this paper's survey say they agree to some extent with the conventional logic for unequal tax burdens as well as with CBBT, echoing a large body of work across a range of fields that has shown it is common for individuals to use a mixture of criteria to make policy judgments. This paper is best seen as providing further support for the idea that models of optimal policy seeking to capture prevailing public priorities ought to use an objective characterized by normative diversity in general and including the principle of CBBT in particular.<sup>2</sup>

In addition, to prevent confusion it is worth emphasizing that this paper is intended not to defend CBBT as a normative criterion but rather to establish and understand the roots of CBBT's importance as a positive matter. As I have discussed elsewhere (Weinzierl 2016), because CBBT addresses a number of the most powerful normative critiques of narrower versions of benefit-based theory and enjoys such a prominent place in public reasoning over taxes, it may merit further study from a normative perspective, but that is not the purpose of this paper.<sup>3</sup>

This paper is closely related to a voluminous modern literature in political philosophy on the role of luck, especially so-called "brute luck" in society. The influential "luck egalitarian" approach of, for example, G.A. Cohen (2011), closely resembles the normative perspective assumed by most of the recent work in optimal tax theory. That is, the objective function in modern tax theory is typically specified such that inequalities in outcomes across individuals due to factors for which individuals do not have responsibility are to be offset, while inequalities for which individuals are responsible are not to be offset (see Fleurbaey and Maniquet 2006, Lockwood and Weinzierl 2015).<sup>4</sup>

---

<sup>2</sup>I formally develop mixed normative objectives in Weinzierl (2014) and Weinzierl (2015), where I emphasize that they are consistent with the most general objectives assumed in modern optimal tax theory (i.e., in the work of Joseph Stiglitz (1987) and Iván Werning (2007), among others). Specifically, the objectives I consider satisfy Pareto efficiency.

<sup>3</sup>How to respond to the public's normative reasoning is explored in the philosophical literature on reflective equilibrium, such as in Norman Daniels (1996).

<sup>4</sup>"Choice" is of course a complicated concept, but in optimal tax theory it is usually represented through heterogeneity in utility functions. So, a person who puts a lower value on leisure may "choose" to work more than others. To some luck egalitarians, including perhaps Cohen, such preference differences ought to be offset. But to others, these preferences are qualitatively (and morally) distinct from what optimal tax theorists call "ability," the capability

What this paper's survey evidence suggests is that the most members of the public have not yet converted to luck egalitarianism. Instead, they appear to have at least some affinity for the (very different) views of Nozick (1974) that "Whether or not people's natural assets are arbitrary from a moral point of view, they are entitled to them, and to what flows from them." Rather than taking the position that pre-tax outcomes (incomes, abilities, etc.) ought to be irrelevant to after-tax outcomes, survey respondents appear to give these intermediate steps moral weight.

The paper proceeds as follows. Section 1 describes the survey and the first main finding: resistance to full equalization of luck-based inequality. Section 2 briefly presents CBBT formally to familiarize the reader with its implications and then discusses the survey's second main finding: support for CBBT as a logic for progressivity in taxation. Section 3 shows that these two findings are linked and comments on their interpretation. I refer to related research as results are presented. Section 4 concludes.

## 1 Resistance to full equalization of luck-based outcomes

This section describes the survey and then turns to the first main result from it.

### 1.1 Survey design

The survey was listed on Amazon's Mechanical Turk (M-Turk) interface<sup>5</sup> and was completed by nearly 2500 respondents (in six rounds) during late 2015 and early 2016. Respondents were paid \$3.00 for a task that took approximately ten minutes to complete. Our main findings correspond to a small set of questions from the survey that will be described in detail in subsequent sections. In addition to those questions, respondents self-reported a set of demographic traits, completed a short arithmetic quiz, and answered a series of questions designed to measure their general political opinions. The relationship of the main results to these questions is discussed below, as well.

### 1.2 Main result: acceptance of luck-based inequality

I begin with the finding that respondents resist full equalization of luck-based outcomes, contrary to what conventional optimal tax analysis would strongly recommend. After respondents start the survey by entering their M-Turk ID number and agreeing to (or rejecting) the terms of the survey, they see the following screen:

---

of an individual to produce output. Ability is treated as brute luck in optimal tax models (though recent work on human capital requires harder thinking about this assumption, too).

<sup>5</sup>M-Turk is a cost-effective and popular platform for surveys, with recent related examples being Saez and Stantcheva (2015), Kuziemko, Norton, Saez, and Stantcheva (2015), and Weinzierl (2014).

*Please consider the following situation.*

Two people are approached with the following offer.

First, a fair coin will be flipped to determine which of the two people is to be called Person A and which is to be called Person B. The results of the coin flip are kept secret until after the two people decide whether to refuse or accept the offer.

If they refuse the offer, the results of the coin flip will be revealed and Person A will receive \$600 while Person B will receive \$300.

If they accept the offer, the results of the coin flip will be revealed and Person A will receive \$60,000 while Person B will receive \$30,000. In exchange, Person A and Person B will have to pay a cost of \$18,000, in total. Person A and Person B could each have to pay part of this cost; one of them could have to pay the entire cost while the other would have to pay nothing; or one of them could have to pay more than \$18,000, in which case the extra money would be given to the other person.

*If they accept the offer, what do you think would be the best outcome? In the first text box, please enter the amount you think Person A should have to pay (enter an amount between -12000 and 60000, and do not use a \$ or a comma). The other three text boxes will fill in automatically and will show you how much Person B would have to pay as well as how much Person A and Person B would end up with. You might find it helpful to try a few numbers in the first text box and see how the results change.*

---

Person A pays \$\_\_\_\_ (please enter an amount between -12000 and 60000; do not use a \$ or comma) :

Person B pays \$\_\_\_\_:

Person A ends up with \$\_\_\_\_:

Person B ends up with \$\_\_\_\_:

Figure 1: A scenario designed to mimic the tax policy problem.

The task of the respondent is to enter an amount for "Person A pays" in the first text box. The amounts for "Person B pays" and "Person A ends up with" and "Person B ends up with" fill in automatically.

This hypothetical situation is designed to capture the essential elements of the tax policy problem for society without directly invoking the concepts of "tax" or "government." In it, Person A and Person B have the chance to (collectively) invest in a project that yields a surplus of total output over total input. Those persons differ in the share of the output they will receive from the project, and they differ in what they will receive if the project is not undertaken. The survey respondent's task is to assign to each person an amount to contribute to the project, where the contribution by either person may exceed the total cost of the project if the respondent wishes to provide a net transfer to the other person.

Thus, the main functions for the contributions by Person A and Person B are those of taxes in the real world: to fund socially productive activity and to determine the distribution of total surplus (output) across individuals, as in Musgrave's (1959) famous delineation of the allocation and distribution branches of government. Note that the respondent is not included in the situation directly, so that he or she is implicitly put into the position of the disinterested observer or social planner.

Importantly, however, this situation is also designed to neutralize two factors that complicate the tax policy problem in reality. First, the allocations to Person A and Person B are entirely due to luck, while the relative roles of luck and tastes in determining outcomes—i.e., the role of



"desert"—has been a long-standing and heated debate in both scholarly and public discussions of tax policy (see the debate over luck egalitarianism, for example). Second, there is no effort exerted in this scenario, so there are no efficiency costs from redistribution. In the jargon of modern optimal tax theory, this scenario has one dimension of exogenous heterogeneity and inelastic labor supply.

Given this design, the optimal allocation according to standard theory (i.e., with a social welfare function that is concave in income) is clear: full redistribution. That is, Person A should pay \$24,000, Person B should receive a transfer of \$6,000, and each should end up with \$36,000. With no preference heterogeneity and a concave social welfare function, equal outcomes maximize social welfare for a given amount of resources, and with inelastic effort the amount of resources is fixed. In other words, Person A's advantages are due to pure luck, not effort, and resources are no more valuable in A's hands than in B's, so full redistribution is the optimal policy.

Respondents are less egalitarian. Figure 1 shows the 2,262 responses to versions of this question for which the answer to "Person A pays \$\_" falls between \$9,000 and \$24,000. I omit the 197 respondents who have Person A pay less than Person B or more than \$24,000, and another 219 respondents to an early round were not asked a similar question. The mean response is \$16,772, with a standard deviation of \$5,267. The modal response is the cost of the offer, \$18,000, the choice under which payments are maximally progressive without providing net redistribution.

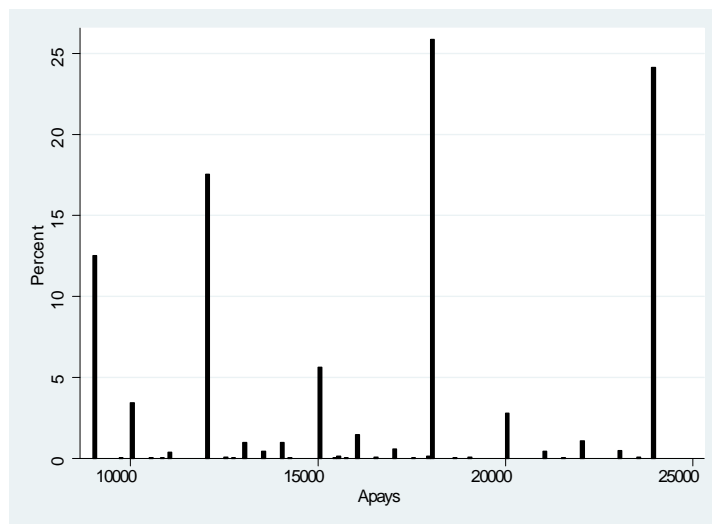


Figure 2: Distribution of responses for "Person A should pay \_\_\_\_\_" between \$9,000 and \$24,000.

The most striking result from this question is that a large majority—more than 75%—of respondents stop short of full equalization of net proceeds from the project even though redistribution is nondistortionary and the gross proceeds are explicitly determined by luck. This result is particularly challenging to the normative assumptions of conventional optimal tax theory because it suggests that the two standard reasons with which conventional theory justifies inequality in after-

tax outcomes are unsatisfying as explanations for American skepticism toward redistributive tax policy.

A plurality of respondents—42%—choose a point between full equalization of outcomes and proportional payments. A bit more than 24% choose to fully equalize the net incomes across individuals (A pays \$24,000), as conventional optimal tax models would recommend, while 18% choose to allocate the costs of the project in proportion to each individual's gross incomes (A pays \$12,000). As we will see, proportional payments are optimal, given the relationship between the outcomes when the offer is refused and accepted as described in Figure 1, according to the formalization of CBBT below in Section 2.

That a plurality of respondents chose an intermediate outcome could be interpreted as suggesting that the typical respondent feels some affinity for the principles behind each of the more extreme choices. Normative diversity of that kind has been documented by a large body of previous work outside economics and a few recent works within it (see Hochschild 1981, Frohlich and Oppenheimer 1992, Feldman and Zaller 1992, Weinzierl 2014 and 2016, Saez and Stantcheva 2015, Lockwood and Weinzierl 2015).

Further supporting these results are the respondents' answers when they are asked "what do you think the *typical American* would say is the best outcome?" in the same tax scenario. The skepticism toward redistribution that respondents attribute to the typical American is even greater than what they express themselves. The mean response is only \$14,735, and only 14% of respondents think that the "typical American" would choose full equalization in this situation, despite the lack of incentive costs or desert claims. 16% of respondents think the typical American would choose proportionality, while fully 29% think the typical American would split the costs of the offer evenly (i.e., A pays \$9,000). Consistent with their own preferences, however, the plurality of respondents think the typical American would choose an outcome between proportionality and full equalization.

Most of the demographic variables asked about in the survey (i.e., race, gender, and household income as either a child or an adult) are not significantly related to these results, though there is some evidence that older respondents and those who correctly answer all three arithmetic questions have Person A pay more on average. Two indicators of political views has predictive power, with respondents who believe an individual (rather than government) is responsible for providing basic needs and those who believe "It would be best if everyone felt the same sacrifice from paying taxes" choosing to have Person A pay significantly smaller amounts.

### 1.3 Robustness to variations in survey question

Here, I consider variations to the survey question shown in Figure 1 meant to address a variety of natural criticisms. While the specifics of the results change, the main finding of this section continues to hold.

### 1.3.1 Phrasing of the survey

First, I slightly change the wording of the question so that it avoids expressions such as "have to pay" that might prompt respondents to think of the payments as especially compulsory. Instead, the question simply uses "pay," as in "Person A and Person B could each pay part of this cost..." The results are virtually unchanged from the benchmark setup: the mean response is \$16,434 (compared to \$16,772) among the 137 respondents who chose a value between \$9,000 and \$24,000. A somewhat higher share (23%) choose proportional payments (\$12,000) and a smaller share (22%) chose the fully redistributive (\$24,000) payments, though a plurality (more than 40%) continue to choose intermediate values. In all, this particular aspect of the setup appears unimportant to the results.

### 1.3.2 The salience of payments versus outcomes

Second, I show respondents only the payments made by each person or the amounts each person ends up with, rather than both as in the benchmark setup. These variations help measure the extent to which elevating the salience of either aspect of the policy changes respondents' moral judgments, for example if emphasizing outcomes makes them more likely to equalize the amounts Person A and Person B end up with. This variation thereby links directly to the question of whether the public endorses the conventional model's assumption that only outcomes ought to matter for policy.

Whether respondents are shown only the payments or only the outcomes, the main results of this section continue to hold, but these changes do have noticeable—and informative—effects. In both cases, a large majority of respondents choose less than full redistribution, and a plurality choose an amount for Person A to pay between \$12,000 and \$24,000. However, support for the egalitarian outcome does shift with these variations on the benchmark. In the payments only variation, 5% of the 133 respondents choose to have A pay \$24,000, while in the outcomes only variation nearly 35% of the 120 respondents choose it (across all setups, 24% chose this outcome). Related, proportional payments (A pays \$12,000) are chosen by 27% and 8% of respondents in the two variations, compared to 18% across settings. One way to interpret these results is that asking respondents to engage with both payments and outcomes causes them to moderate the more extreme views they have if they consider only one or the other. For example, the mean amount Person A pays was \$16,772 across all surveys, compared to \$14,135 in the payments-only variation and \$17,988 in the outcomes-only variation.

### 1.3.3 Discrete versus continuous choices

One potential concern with the setup shown in Figure 1 is that respondents may not consider the full range of possible allocations when entering the amount Person A pays in the first text box. For example, respondents may default to having Person A pay \$18,000 because it is the cost of the project, not understanding that A could be asked to pay more than the cost so as to fund a transfer to Person B. Related, though the three text boxes fill in automatically once the respondent

chooses an amount for A to pay, the implicit mathematics behind Figure 1’s setup may be too complicated for some respondents, causing them to default to simple numbers that don’t reflect their true opinions, such as having each person pay \$9,000.

To address this concern, I modify the question to provide respondents with a discrete set of choices, each of which indicates the amounts that both persons pay (or receive) and end up with. In one variation, I provide four choices: Person A pays 9000, 12000, 18000, or 24000. In a second variation, I add two additional (intermediate) choices: Person A pays 15000 or 21000. In a third, I add two more (extreme) choices: Person A pays 6000 or 27000. Across these three versions, I obtain 266 responses.

These variations yield weaker versions of the overall results but leave intact the main finding that a large share of respondents resist costless redistribution. Figure 3 shows the respondents’ answers to the same question as in Figure 1, but for the discrete choice variations only.

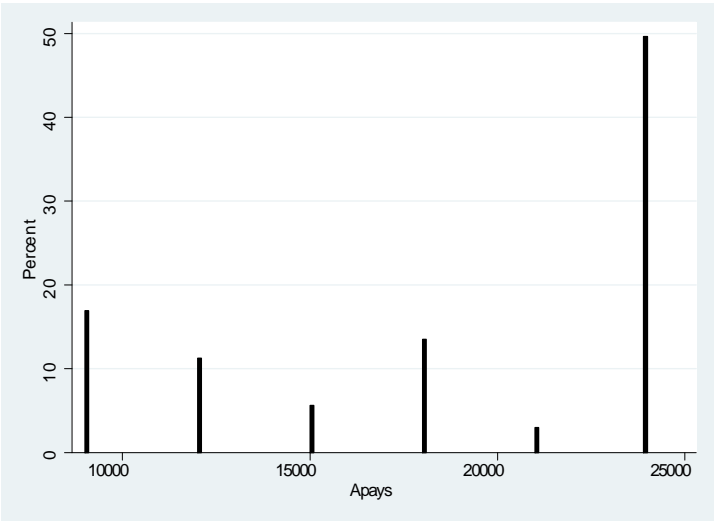


Figure 3: Distribution of 266 responses for "Person A should pay \_\_\_" between \$9,000 and \$24,000, in the discrete options variations of the survey setup.

Consistent with the concerns expressed above, these variations reduce the concentration of answers at \$18,000 and increase the share of respondents choosing full equalization of outcomes. In fact, the modal choice in these variations—chosen by nearly 50% of respondents—is Person A paying \$24,000, the fully redistributive choice recommended by the conventional model’s assumed normative view. The remaining respondents—just over 50%—continue to stop short of fully redistributing Person A’s luck-based advantage.

While these discrete choice versions of the survey clearly have advantages, there are countervailing considerations that complicate their interpretation relative to the continuous, text-entry setup of Figure 1. First, multiple-choice survey questions require less engagement by the respondent than

text-entry questions, potentially leading to less informative results. By requiring the respondent to enter an amount, text-entry questions avoid the concern that respondents may quickly scan a set of choices and choose with little deliberation. The particular technology of the Figure 1 setup has the related advantage of making the interdependence between payments and outcomes explicit, as the amounts in the three later text boxes change as the respondent enters the amount Person A pays. Second, any set of choices unavoidably introduces elements of framing, perhaps in complex non-uniform ways across the set of options, relative to a setup in which respondents have to introspect about their choice. Figure 1's text boxes are blank until the respondent enters a number, and the instructions make clear that Person A could pay any amount between -\$12,000 and \$60,000. In contrast, the multiple-choice setup makes certain choices salient, and the set of options listed determines the context in which each choice is seen by the respondent. In particular, listing the options may privilege those with "simple" features, like equal payments or equal outcomes, because they are easy for the respondent to understand. Figure 3 provides some support for that possibility, showing that both the \$9,000 and \$24,000 options for Person A's payment increase their shares relative to the benchmark setup.

## 2 CBBT as a logic for progressivity

I now turn to the second novel finding of this paper: the widespread preference among survey respondents for a benefit-based justification of progressivity rather than a justification based on diminishing marginal social welfare of income. First, however, I provide a primer on CBBT, as it and its formulation in the modern optimal tax model are not well known. The interested reader can find a more thorough treatment, including additional analytical results and a discussion of the sharp contrast between CBBT's past prominence and its present neglect in tax scholarship, in Weinzierl (2016).

### 2.1 Primer on CBBT

As mentioned in the Introduction, CBBT plays a prominent role in American rhetoric on tax policy. An important example is the following statement by President Barack Obama, who in 2011 argued for increased progressivity of the income tax:

"As a country that values fairness, wealthier individuals have traditionally borne a greater share of this [tax] burden than the middle class or those less fortunate....it's a basic reflection of our belief that those who've benefited most from our way of life can afford to give back a little bit more."<sup>6</sup>

---

<sup>6</sup>In fact, close examination of the fiery debate over the phrase "you didn't build that" in the 2012 U.S. presidential election reveals that both sides of the debate were fighting for their preferred tax policies within the same, CBBT-based, framework. Mitt Romney said, "The taxpayers pay for government....we pay for them and we benefit from them...."

Here, Obama argues that taxes ought to be based on benefit from "our way of life," and he explicitly links that benefit to the taxpayer's ability to pay. This combination of two classic principles of tax design (benefit-based and ability-based) into a "benefit-as-ability" based principle is a particularly succinct statement of CBBT. It recalls a more famous statement of it by Adam Smith (1776) as his first maxim of taxation: "The subjects of every state ought to contribute toward the support of the government, as near as possible, in proportion to their respective abilities; that is in proportion to the revenue which they respectively enjoy under the protection of the state."

Fortunately, CBBT lends itself readily to formal analysis within the apparatus of modern tax theory. In particular, we can analyze optimal tax policy according to CBBT by modifying the standard Mirrleesian setup in one small but important way: by making individuals' heterogeneous income-earning abilities endogenous functions of both endowed ability and public goods spending. Formally,  $w_i = f(a_i, G)$ , where  $i \in I$  indexes endowed ability types  $a_i$ ,  $G \geq 0$  is the level of spending on public goods, and  $f(\cdot)$  is a differentiable ability production function. An individual of type  $i$  derives utility according to

$$U(c_i, l_i) = u(c_i) - v\left(\frac{y_i}{f(a_i, G)}\right), \quad (1)$$

where  $c_i$  is private consumption for individual  $i$  and  $y_i$  is  $i$ 's income, so that  $\frac{y_i}{w_i}$  is work effort.

A social planner chooses a tax system, including an optimal  $G^*$ . Individuals take that system as given and maximize their own utility, yielding equilibrium consumption and income allocations  $\{c_i^*, y_i^*\}_{i=1}^I$  and utility levels:

$$U_i^* = u(c_i^*) - v\left(\frac{y_i^*}{f(a_i, G^*)}\right). \quad (2)$$

By applying the method of Lindahl (1919) to this setup we can determine the first-best optimal allocation under CBBT. That method has us consider a hypothetical scenario in which each individual  $i$  is allowed to choose her own consumption, work effort, and, importantly, level of public goods provision that maximize her utility subject to her personal budget constraint, taking the tax share  $\tau_i$  as given. Lindahl defined optimal policy as that in which two conditions are satisfied: first, the personalized shares cause each type to prefer the same quantity of public goods<sup>7</sup>; second, the cost of the public goods is fully covered by tax payments. I call the allocation that satisfies these conditions a *First-Best Lindahl Equilibrium*.

The feature of the resulting allocation most relevant to this paper is the taxes paid by each individual. To characterize those taxes, I first define a key elasticity term:

**Definition 1** *Define the Hicksian partial elasticity of complementarity between public goods and*

---

<sup>7</sup>It is this step that lends, according to benefit-based taxation's advocates, such a system a claim to being voluntary rather than coercive. Of course, benefit is unobservable, so that the second-best CBBT tax system will be coercive in a sense. Nevertheless, there remains an essential difference between the benefit-based system and, for example, a utilitarian one. In the first-best allocation of the former but not the latter, an individual pays a "price" for the activities of the state that is determined by his or her willingness to pay (i.e., marginal rate of substitution).

endowed ability,  $\theta_i^{G,a}$ , as:

$$\theta_i^{G,a} = \frac{f_{G,a}(a_i, G) f(a_i, G)}{f_G(a_i, G) f_a(a_i, G)}, \quad (3)$$

at a given  $G$ .

The Hicksian partial elasticity of complementarity captures the degree to which public goods and endowed ability magnify each other in determining income-earning ability. If  $\theta_i^{G,a} \leq 0$ , endowed ability and public goods are not complements in the production of income-earning ability. If  $\theta_i^{G,a} \in (0, 1)$  the elasticity of income-earning ability with respect to the level of public goods spending is positive but decreasing in endowed ability; if  $\theta_i^{G,a} > 1$ , the elasticity of income-earning ability with respect to the level of public goods spending is increasing in endowed ability.

As shown formally in Weinzierl (2016), this elasticity of complementarity determines the progressivity of tax rates under CBBT. If  $\theta_i^{G,a} > 1$ , so that those high in endowed ability benefit more than proportionally from the activities of the state, average tax rates are progressive (i.e., they increase in endowed ability). If  $\theta_i^{G,a} < 1$  taxes are regressive, and if  $\theta_i^{G,a} = 1$  taxes are proportional to income. This last case, which Smith (1776) appears to endorse, obtains if we assume a multiplicative form for the ability production function, i.e.,  $f(a_i, G) = h(a_i)g(G)$  for some functions  $h(a_i)$ ,  $g(G)$ . In that case, the flat tax rate on income equals the elasticity of income-earning ability with respect to public goods spending. For example, if  $g(G) = g^\gamma$  for some  $\gamma > 0$ , then the CBBT-optimal tax policy is a uniform tax rate of  $\gamma$ .<sup>8</sup>

It is important to note that optimal taxes under CBBT do not depend on the distribution of after-tax outcomes. Instead, they are defined by the relationship between individuals' innate abilities and pre-tax incomes (equivalently, their eventual ability levels), in stark contrast to the conventional normative approach in optimal tax.

## 2.2 Survey evidence of support for CBBT

Immediately after respondents make their choices in the hypothetical tax-like situation described in Section 1, the following screen asks them to consider the reasoning behind their choices:

---

<sup>8</sup>As noted earlier, the setup of the scenario in Figure 1 implies that the ability production function  $f(a_i, G^*)$  is multiplicative and the elasticity  $\theta_i^{G,a} = 1$ . In principle, we could use that tax-like scenario to vary the implicit functional form  $f(a, G)$  from the model above, for example by varying the relationships between the individuals' proceeds with and without the project.

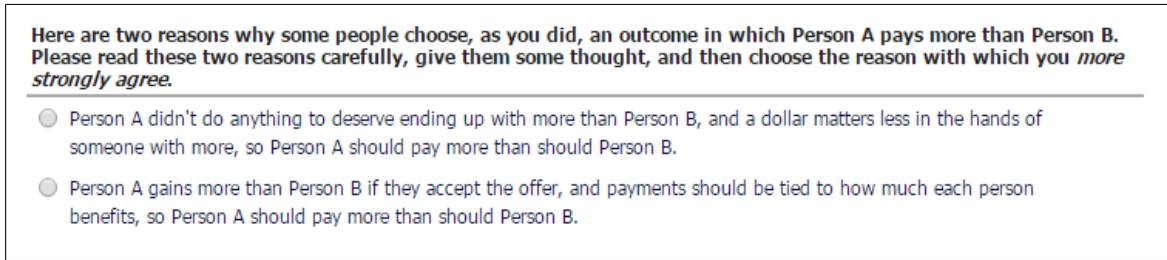


Figure 4: Respondents are asked to choose a logic for progressive payments in the tax scenario.

The first of the two reasons refers to the conventional optimal tax literature’s logic of diminishing marginal social welfare of income, while the second reason refers to the CBBT principle. The order of these two reasons was randomized.<sup>9</sup> After respondents make their choice on the question in Figure 4, the survey asks about the strength of their opinions on both reasons. Specifically, respondents are asked whether they strongly disagree, disagree, agree, or strongly agree or disagree with each of the reasons in Figure 4.

The results of these questions suggest that, when reasoning over tax policy, Americans are more comfortable with the logic of CBBT than with the conventional logic assumed in modern optimal tax theory. More than 70% of 2,209 respondents who answered both the question in Figure 4 and the question in Figure 1 preferred the benefit-based justification to the conventional one. Moreover, 70% of the 30% who preferred the conventional logic either agree or strongly agree with CBBT, bringing the total share of respondents expressing support for CBBT to around 90%. In contrast, less than half of those who preferred CBBT express agreement with the conventional logic.

### 2.2.1 Robustness to survey design

Variations in the overall survey design (as described in Section 1) and variations in the wording of this question in particular generate variation around the 70% overall figure, but the share of respondents choosing the CBBT-based option lies between 60% and 79% in all variations. The smallest shares are obtained in the discrete options variations described in Section 1.3.3. (Note that this suggests the discrete options sample pool, which was relatively small, may have been systematically more comfortable with the conventional normative assumptions of the optimal tax model). The largest shares are obtained when we use either the outcomes-only variation described in Section 1.3.2 or when we use a variation that excludes the phrase "didn't do anything to deserve ending up" and replaces it with "ends up" at the start of the first option in Figure 4. The former phrase emphasizes that Person A did not affirmatively deserve such a lucky outcome, but it is just

<sup>9</sup>In later rounds of the survey, I had respondents who did not choose to have Person A pay more than Person B see a similar screen, though they were asked which of these two reasons they "think would be the better reason for having Person A pay more than Person B."



as appropriate to include in the description of the second reason in Figure 4, so this variation is, arguably, the more neutral setup.

### 2.2.2 Robustness across respondent traits

Support for CBBT is essentially uniform across demographic and socioeconomic traits that the respondents self-report, as shown in the following tables for all variations of the survey.

Table 1: Shares of respondents preferring CBBT logic for progressivity, by demographic trait

	Gender		Age		Race			Education		Income, child		Income, adult	
	M	F	25-39	40+	Bl.	Wh.	Oth.	≤A.B.	A.B.+	Low	Upp	Low	Upp
Share	0.70	0.71	0.70	0.71	0.70	0.70	0.69	0.69	0.71	0.71	0.72	0.71	0.69
Obs	1150	1051	1626	578	180	1826	199	998	1208	1335	212	921	344

Table 2: Shares of respondents preferring CBBT logic for progressivity, by political view

	Political position			Libertarianism			Basic needs		Dist. of Sacrifice	
	Left	Center	Right	Opp	Unsure	Supp	Govt	Individ	Prog.	Equal
Share	0.70	0.70	0.71	0.74	0.69	0.69	0.70	0.70	0.70	0.70
Obs	1043	583	499	621	602	983	1403	803	1629	580

Statistical analysis in Section 3 will largely support these simple cross-tabulations, though with some suggestive exceptions. One interpretation of these patterns is that support for CBBT is largely universal, consistent with the substantial support it receives among those who prefer the conventional logic.

## 3 Link between the results, and the moral relevance of pre-tax outcomes

In this section, I show that the two main results are linked: that is, respondents who more strongly resist equalization of luck-based outcomes are more likely to prefer the benefit-based logic for taxation. After presenting the evidence for this linkage, I discuss one potential interpretation of it that relates to the debate over whether pre-tax outcomes ought to be considered morally relevant for tax design.

Visual evidence of this linkage is shown in Figure 5, which gives the share of respondents preferring the CBBT principle in the question from Figure 4 for five ranges of answers to how much

Person A should pay in the question from Figure 1.




Figure 5: Share of respondents choosing CBBT, by their response to "Person A pays \_\_\_".

This figure shows a substantial decline in the share supporting CBBT as affinity for redistribution rises. Note, however, that even among those respondents choosing full equalization (A pays \$24000), 63% prefer CBBT to the conventional logic. This pattern holds across variations of the survey questions, including in the multiple-choice versions described in section 1.3.3 (where support for redistribution was substantially greater).

Table 3 presents statistical evidence of this linkage. The dependent variable is whether the respondent prefers CBBT in the question from Figure 4. The first explanatory variable is the respondent's answer to how much Person A pays in the question from Figure 1. The remaining explanatory variables are those shown in Tables 1 and 2, where cross-tabulations showed little variation in support for CBBT across demographic traits or political opinions.

Table 3: Explaining preference for CBBT logic for progressivity, benchmark setup

	Coefficient	Standard error
Person A pays	$-3.3E^{-5}$	$(0.6E^{-5})^{**}$
Political position (L to R)	0.02	(0.05)
Support for libertarianism	-0.09	(0.04)
Govt responsible for basic needs	-0.04	(0.04)
Support for equal sacrifice	0.10	$(0.05)^{**}$
HH's income status at age 45	-0.01	(0.02)
HH's income status at age 10	-0.02	(0.02)
Gender (M=0,F=1)	0.01	(0.06)
Age	0.04	(0.05)
Race (White=0, Black=1)	-0.01	(0.05)
Education level	0.02	(0.05)
Score on math questions	0.19	$(0.10)^*$
N	1,888	

Notes: The dependent variable in this probit regression is the respondent's 0-1 choice between the conventional logic (0) and the CBBT logic (1) for progressivity, as shown in Figure 4. The mean value for the dependent variable is 0.70. A positive coefficient on an explanatory variable therefore indicates that a higher value for it is related to a higher likelihood the respondent prefers the CBBT logic for progressivity. The symbol \*\* denotes significance at the 5% level; \* at the 10% level.

These results indicate that the relationship between resistance to redistribution and support for CBBT is significant and substantial. Calculating marginal effects from these results shows that, for example, a \$1,000 decrease in how much the dictator ought to give to the recipient implies an increase of 1.1 percentage points in the likelihood that the respondent prefers the CBBT logic to the conventional logic for progressivity (the mean value is 70%). Extrapolating this effect across the range of values for "Person A pays" would explain more than two-thirds of the gap between the share of respondents preferring CBBT shown in Figure 6 across the range from \$12,000 to \$24,000.

The only other significant relationship in Table 3 shows that respondents who believe the "sacrifice" from paying taxes ought to be borne equally (rather than more by the rich) are more likely to support CBBT. The conceptual links between the principle of Equal Sacrifice and CBBT were hinted at in Feldstein (1976) and are discussed at greater length in Weinzierl (2016).

### 3.1 Are pre-tax outcomes morally irrelevant to tax policy?

How might we understand the link between resistance to redistribution of luck-based outcomes and support for a benefit-based principle of taxation suggested by the results of this paper's survey?

One interpretation of these results relates to a core normative assumption in optimal tax theory:

namely, that distributions of pre-tax income and taxes are morally irrelevant other than through their mechanical connection to the outcomes that ought to be the target for policymakers. Murphy and Nagel (2003) forcefully defend this assumption by pointing out that taxes are just one part of a suite of policy institutions determining outcomes: "Pretax income, in particular, has no independent moral significance. It does not define something to which the taxpayer has a prepolitical or natural right, and which the government expropriates from the individual in levying taxes on it."

Respondents to this paper's survey appear to disagree with this assumption. A majority effectively grant Person A some entitlement to a purely luck-based advantage, bestowing on that pre-tax outcome a moral relevance it is denied by a conventional welfarist objective. Consistent with that unconventional position, respondents who are more willing to accept luck-based outcomes prefer to have taxes be based on a principle such as CBBT, which seeks to implement a more "voluntary" tax system and defines optimality in terms of the relationship between taxes and pre-tax outcomes.

In other words, a large share of this paper's survey respondents appear to put the burden of proof for desert on the opposite side of where most modern egalitarian political philosophers do. For egalitarians, unequal outcomes due to morally arbitrary luck are unacceptable unless proven otherwise and therefore should be offset. Our results suggest that Americans, in contrast, are sympathetic to the view that unequal outcomes due to morally arbitrary luck are acceptable unless proven otherwise and therefore not the proper object of redistribution.

Though the idea that pre-tax incomes and taxes are morally relevant to the public sharply contradicts the standard approach, this paper is not the first to find evidence supporting it. Charité, Fisman, and Kuziemko (2015) demonstrate that M-Turk respondents are less likely to equalize random allocations across individuals if those individuals know the results of the randomization (as they do in this paper's survey) than if they do not. They interpret this finding as evidence that respondents both assume the individuals are loss-averse and take the welfare implications of that loss aversion into account. A complementary interpretation is that the knowledge of the allocations causes individuals to grant each other entitlements to them. While such resistance to redistributing entitlements may be described through the mathematics of loss aversion, it may represent a more fundamental social judgment based on a non-consequentialist normative criterion. Our findings are also consistent with those of Saez and Stantcheva (2015), who find "...evidence showing that both disposable income and taxes paid matter and hence that subjects are neither pure utilitarians (for whom only disposable income matters) nor pure libertarians (for whom only taxed paid matter)." In the formal terms of their analysis, the marginal social welfare weight that the public appears to grant to an individual, which determines the optimal allocation for that person, depends positively on the taxes that person pays. Finally, Weinzierl (2014, 2016) has shown evidence that two unconventional principles, J.S. Mill's (1871) principle of Equal Sacrifice and CBBT (as described in this paper), both capture an aspect of public reasoning over tax policy in which pre-tax outcomes are relevant to optimal policy.

In fact, the moral relevance of pre-tax income is a source of controversy in political philosophy. Critiquing Murphy and Nagel (2003), Brennan (2005) writes: "The problem with Murphy and

Nagel's argument, as I see it, is that it takes an entirely defensible claim—namely that individuals do not have an incontestable moral claim to their individual gross incomes—and replaces it with a much stronger claim—that they have no moral claim to their individual incomes at all...I think there is a middle turf. I think it's obvious that there's a middle turf." This paper presents evidence on the position taken in this debate by the American public.

## 4 Conclusion

This paper's results raise many questions. How robust are the respondents' answers to alternative ways of framing the scenarios? Do these responses reflect their considered preferences or their gut reactions that would change if they gave more time to the questions? Would "education" in these issues change their preferences? Do their stated preferences for these hypothetical scenarios translate to their votes for policy and policymakers?

That said, the main contribution of this paper is to present survey evidence that the normative views of most Americans appear to include ambivalence toward the egalitarianism that has been so influential in contemporary political philosophy and implicitly adopted by modern optimal tax theory. To the extent that this basic finding is valid and represents sincere normative diversity, optimal tax theorists ought to consider capturing that ambivalence in their work, as well.

## References

- [1] Brennan, Geoffrey, "The Myth of Ownership," *Australian Journal of Legal Philosophy*.
- [2] Charité, Jimmy, Raymond Fisman, and Ilyana Kuziemko (2015), "Reference Points and Redistributive Preferences: Experimental Evidence," NBER Working Paper 21009.
- [3] Cohen, G.A. (2011). *On the Currency of Egalitarian Justice, and Other Essays in Political Philosophy*. Princeton.
- [4] Daniels, Norman (1996). *Justice and Justification*. Cambridge.
- [5] Feldman, Stanley and John Zaller (1992), "The Political Culture of Ambivalence: Ideological Responses to the Welfare State," *American Journal of Political Science* 36(1).
- [6] Feldstein, Martin (1976). "On the Theory of Tax Reform," *Journal of Public Economics*, 6, pp. 77-104.
- [7] Fleurbaey, Marc and Francois Maniquet (2006). "Fair Income Tax," *Review of Economic Studies*, 73, pp. 55-83.
- [8] Frolich, Norman and Joe A. Oppenheimer (1992). *Choosing Justice: An Experimental Approach to Ethical Theory*. Berkeley, CA: UC Press.
- [9] Harsanyi, John C. (1953) "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking," *Journal of Political Economy*, 61(5), (October), pp. 434-435.

- [10] Harsanyi, John C. (1955). "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility," *Journal of Political Economy*, 63(4), (August), pp. 309-321.
- [11] Hochschild, Jennifer (1981). *What's Fair? American Beliefs about Distributive Justice*. Cambridge, MA: Harvard University Press.
- [12] Kuziemko, Ilyana, Michael Norton, Emmanuel Saez, and Stefanie Stantcheva, (2015). "How Elastic are Preferences for Redistribution: Evidence from Randomized Survey Experiments," *American Economic Review*
- [13] Lindahl, Erik (1919). "Just Taxation – A Positive Solution" (E. Henderson, transl.) in Musgrave, Richard A. and A.T. Peacock, Eds, *Classics in the Theory of Public Finance*. Macmillan.
- [14] Lockwood and Weinzierl (2015). "De Gustibus non est Taxandum," *Journal of Public Economics*.
- [15] Mirrlees, J.A., (1971). "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies* 38, 175-208.
- [16] Murphy, Liam and Thomas Nagel (2003). *The Myth of Ownership*, Oxford.
- [17] Musgrave, Richard A. (1959). *The Theory of Public Finance*. McGraw-Hill.
- [18] Nozick, Robert (1974). *Anarchy, State, and Utopia*. Basic Books.
- [19] Rawls, John (1971). *A Theory of Justice*. Cambridge: Harvard University Press.
- [20] Saez, Emmanuel and Stefanie Stantcheva (2015). "Generalized Social Welfare Weights for Optimal Tax Theory." Forthcoming, *American Economic Review*.
- [21] Smith, Adam (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*.
- [22] Stiglitz, Joseph (1987). "Pareto Efficient and Optimal Taxation and the New Welfare Economics," *Handbook of Public Economics*, Chapter 15.
- [23] Sugin, Linda (2011). "A Philosophical Objection to the Optimal Tax Model," *NYU Tax Law Review*
- [24] Weinzierl, Matthew (2014). "The Promise of Positive Optimal Taxation: Normative Diversity and a role for Equal Sacrifice," *Journal of Public Economics*, 118.
- [25] Weinzierl, Matthew (2016). "Revisiting the Classical View of Benefit Based Taxation," *Economic Journal*.
- [26] Werning, Ivan, (2007). "Pareto Efficient Income Taxation," Working Paper, MIT, April.