

NBER WORKING PAPER SERIES

POPULAR ACCEPTANCE OF INEQUALITY DUE TO BRUTE LUCK AND SUPPORT FOR
CLASSICAL BENEFIT-BASED TAXATION

Matthew Weinzierl

Working Paper 22462
<http://www.nber.org/papers/w22462>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2016

Thanks to Max Bazerman, Rafael DiTella, Louis Kaplow, Benjamin B. Lockwood, Francois Maniquet, David Moss, Eric Nelson, Michael Norton, Julio Rotemberg, Florian Scheuer, Hannah Shaffer, and Stefanie Stantcheva for insightful discussions. The author declares that he has no relevant or material financial interests that relate to the research described in this paper. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Matthew Weinzierl. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Popular Acceptance of Inequality due to Brute Luck and Support for Classical Benefit-Based Taxation

Matthew Weinzierl

NBER Working Paper No. 22462

July 2016

JEL No. D63,H21

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 research. A large share of respondents, and in some cases a large majority, resist the full equalization of inequality due to brute luck that standard analyses would recommend. Related, a similar share prefer a classical benefit-based logic for the assignment of taxes over the conventional logic of diminishing marginal social welfare. Moreover, these two views are linked: respondents who more strongly resist equalization are more likely to prefer the classical benefit-based principle. Together, these results suggest that a large share of the American public views the allocation of pre-tax incomes as relevant to optimal tax policy and—at least in part—justly deserved unless proven otherwise, judgments that are inconsistent with standard welfarist objectives.

Matthew Weinzierl

Harvard Business School

277 Morgan

Soldiers Field

Boston, MA 02163

and NBER

mweinzierl@hbs.edu

Introduction

According to novel survey evidence presented in this paper, the U.S. public's views on distributive justice are at odds with key features of the normative view typically applied in modern optimal tax research. This evidence suggests that the conventional approach is likely to disappoint tax scholars, advisors, and policymakers who want their theoretical frameworks and recommended tax reforms to be consistent with the public's underlying policy preferences. The results of this analysis can be organized into three main findings.

First, a large share of survey respondents resist full equalization of after-tax incomes even when conventional optimal tax analyses would strongly recommend it. In a hypothetical situation meant to mimic the tax policy problem, between 50% and 95% of respondents choose not to fully offset inequality due to brute luck even when there are neither efficiency costs of redistribution nor differences in desert across individuals. These choices suggest that the two reasons why conventional optimal tax analyses tolerate after-tax inequality—the importance of encouraging effort and the possibility that some people "choose" to have low incomes—are not the only reasons why survey respondents, and perhaps Americans in general, accept it.

The second finding offers an explanation of the first: a large share of survey respondents prefer an alternative logic for taxation than that which is typically used in optimal tax analyses. The conventional logic stems from the use of a social welfare function that exhibits diminishing marginal social welfare of income. When presented with two possible justifications for their choices in the tax problem, between 62% and 79% of respondents prefer, instead of this logic, one tied to a centuries-old idea 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 economic opportunities. In addition to being Adam Smith's first maxim of taxation, 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 progressivity. In that context, finding support for CBBT among the American public is natural, despite its absence from modern optimal tax theory.

The third finding of this paper is that the first two are linked, in that those respondents more willing to accept inequality due to brute luck are significantly more likely to prefer CBBT as an optimal tax principle. In other words, a large share of respondents appear to support—at least in part—the ideas that individuals are entitled to pre-tax incomes and that taxes ought to respect that entitlement. Advocates of benefit-based taxation stress exactly these ideas 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.

These results therefore speak to a conceptual debate within optimal tax theory, and political philosophy, over whether pre-tax incomes have any moral significance for policy design. As formalized first by James Mirrlees (1971), modern optimal tax research typically adopts an objective put

forth by John Harsanyi (1953, 1955); namely, to maximize a social welfare function that depends only on individual utility levels. Though such a consequentialist objective can in principle accommodate a wide range of judgments, almost all applications of it embrace what Liam Murphy and Thomas Nagel (2002) deem "the right way, investigating outcomes rather than the distribution of [tax] burdens." Under the standard approach, therefore, pre-tax incomes and taxes paid have no effect on welfare or relevance to optimal policy. In contrast, a large majority of respondents to this paper's survey support CBBT, a principle that ignores after-tax incomes (i.e., "outcomes") and *defines* optimality in terms of the relationship between pre-tax incomes and taxes paid.

It is important to clarify up front that the results of this paper are entirely consistent with there being some role for other principles in Americans' appraisals of tax policy. In fact, nearly two-thirds of the respondents to this paper's survey say they agree to some extent with the conventional logic for assigning taxes 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 that include the principle of CBBT in particular. In Weinzierl (2014, 2015), I formally develop mixed normative objectives and emphasize their consistency 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).

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.¹

This paper is closely related to a voluminous modern literature in political philosophy on the role of luck in economic outcomes, especially the so-called brute luck that is not the result of an individual voluntarily accepting risk. 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).² This paper's survey

¹How to respond to the public's normative reasoning is explored in the literature on reflective equilibrium, such as in Norman Daniels (1996).

²"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 of an individual to produce output. Ability is treated as brute luck in optimal tax models (though recent work on human capital, such as Stantcheva 2016, complicates this assumption).

evidence suggests that most members of the American public have not yet converted fully 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."³

The paper proceeds as follows. Section 1 describes the survey and the first main finding: resistance to full equalization of inequality due to brute luck. Section 2 introduces CBBT and then discusses the survey's second main finding: support for CBBT as a principle of taxation. Section 3 shows that the first two findings are linked and comments on their interpretation. I refer to related research as results are presented. Section 4 concludes. In the Appendix, I provide brief formal statements of the standard optimal tax model and a CBBT modification to it for reference.

1 Resistance to full equalization of inequality due to brute luck

First, a note on the survey behind this paper's evidence, which I have reproduced in full in the Appendix. I listed the survey on Amazon's Mechanical Turk (M-Turk) interface⁴ in six rounds during late 2015 and early 2016. Nearly 2500 respondents were paid \$3.00 to complete the survey, which took approximately ten minutes. The main findings correspond to a small set of questions from the survey that I will describe below. 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. I discuss the relationship of the main results to these questions below, as well.

1.1 Main result: acceptance of inequality due to luck

After respondents start the survey by entering their M-Turk ID number and agreeing to the terms of the survey, they see the following screen:

³Importantly, the findings here do not imply that Americans accept inequality regardless of its cause. For example, unjust acquisition leading to inequality would be rejected by even those who fully embrace the libertarian view. Rhetoric in the 2016 U.S. presidential election emphasizing that "the system is rigged" is consistent with this reason for opposing inequality.

⁴M-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). See Horton, Rand, and Zeckhauser (2010) for an analysis of the reliability of online labor markets for experiments in economics.

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 respondents' task is to enter an amount for "Person A pays" in the first text box. The amounts for "Person B pays," "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.⁵ 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, due to brute luck they cannot avoid, in the share of the output they will receive if the project is undertaken and in what they will receive if it is not. 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. The respondent is not included in the situation directly, so he or she is implicitly put into the position of the disinterested observer or social planner.

⁵The situation is presented without directly invoking the concepts of "tax" or "government" so as to avoid causing respondents to answer based on their experience with specific political institutions.

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 importance of luck and tastes in determining incomes—i.e., the role of "desert"—has inspired a long-standing and heated debate in both scholarly and public discussions of tax policy. 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.⁶

Given this design, the optimal allocation according to the standard optimal tax objective (i.e., a social welfare function that is concave in income) is clear: full equalization. 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 after-tax incomes maximize welfare for a given amount of resources, and with inelastic effort the amount of resources is fixed.

Respondents are less egalitarian. Figure 1 shows the 2,037 responses to versions of this question for which the answer to "Person A pays \$_" falls between \$9,000 and \$24,000.⁷ The mean 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 a net transfer to Person B.




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 of respondents—more than 75%—stop short of full equalization of the net proceeds from the project even though redistribution

⁶In the notation of the model in the Appendix, these simplifications amount to assuming all individuals have the same θ and the incentive constraints are ignored in the tax authority's optimization.

⁷I omit the 197 respondents who have Person A pay less than Person B or more than \$24,000. Another 219 respondents to an early round were not asked a similar question.

is nondistortionary and the gross proceeds are explicitly determined by luck. This result suggests that the two reasons emphasized in conventional optimal tax analyses for allowing inequality in after-tax incomes are unsatisfying as explanations for American skepticism toward redistribution.

A substantial share 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 analyses 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 is shown in the Appendix, proportional payments are optimal under CBBT given the relationship (described in Figure 1) between the gross proceeds when the offer is refused and accepted.

That more than two-fifths of respondents, including the median respondent, choose progressivity but not equalization is consistent with the idea 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, 2016, Saez and Stantcheva 2015). It also appears in the second result of this paper, discussed below.

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.⁸ 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, a substantial share of respondents (33%) think the typical American would choose an outcome between proportionality and full equalization.

1.2 Robustness across respondent traits

The survey gathers several indicators of respondents' personal traits. It asks about three demographic indicators: their age in four ranges (18-25, 26-40, 41-64, and 65+); race (white, black, or other); and gender (male, female). It also asks them to report their education level completed (some high school, high school graduate, some college, or college graduate). To obtain an estimate of their economic status, it shows them the current CBO income distribution (four lower quintiles and then four finer-grained quintiles within the top quintile) in a column chart and asks them for their household's position in that distribution when they were 10 years old and 45 years old. Finally, it tests respondents' numeracy with three multiple-choice arithmetic questions.

⁸This result is not due to the over-representation of those who identify with the political "left" among the respondent population (see Table 2), as the mean response for the "typical American" is substantially less than for those respondents who identify with the political "right."

Table 1 shows the mean answer (to how much Person A should pay) among the respondents by their answers to these questions. The mean overall was \$16,772.

	Gender		Age		Race		
	M	F	25-39	40+	Black	White	Other
Mean	16,710	16,874	16,613	17,262	16,263	16,832	16,846
Obs	1057	972	1504	528	169	1685	179

	Education		Income, child		Income, adult		Numeracy	
	≤A.B.	A.B.+	<80th pct	top 20th pct	<80th pct	top 20th pct	<100%	100%
Mean	16,574	16,949	16,684	16,622	16,798	16,685	15,513	16,935
Obs	917	1117	1218	192	831	332	233	1804

The subgroup means are similar, with standard deviations within answers of approximately 5,000. The largest gaps appear in the numeracy and age categories, with respondents who give more correct answers to the arithmetic questions and older respondents having Person A pay more.

The survey also has respondents self-report several aspects of their political opinions. It asks them to describe their political perspective on economic issues (left-leaning or liberal, centrist or moderate, right-leaning or conservative, not sure); to say whether they strongly or somewhat oppose or support libertarianism with regard to economic issues (libertarianism is not explicitly defined); to say whether the government or individuals are responsible for people having their basic needs met; and to say whether they think that the "sacrifice" from paying taxes ought to be borne more by the rich than the poor or borne equally by everyone.

Table 2 show the mean answer to how much Person A should pay among the respondents who chose each answer to these political opinion questions.

	Political position			Libertarianism			Basic needs		Dist. of Sacrifice	
	Left	Center	Right	Opp	Unsure	Supp	Govt	Individ	Prog.	Equal
Share	17,126	16,535	16,392	16,909	16,971	16,586	17,057	16,294	17,036	16,046
Obs	948	548	459	547	563	924	1294	740	1494	543

The differences across mean answers by political opinion are somewhat larger than across personal traits. As might be expected, respondents who identify with the political right, who think individuals ought to be responsible for meeting their basic needs, and who think the sacrifice from taxes ought to be borne equally (rather than progressively) have Person A pay substantially less.

The patterns visible in these tables are confirmed in a simple OLS regression. The only statistically significant demographic predictors are numeracy and age, with those who got all three

arithmetic questions and older respondents having Person A pay more. Two of the indicators of political views—belief in individual responsibility for basic needs and support for equal sacrifice—have significant predictive power in the directions suggested by Table 2. While statistically significant, none of these variables have economically substantial effects on the mean answers.

In sum, resistance to full equalization of inequality due to luck appears to be widespread across subgroups of the survey population. It also appears to be related, not surprisingly, to respondents' general views on the proper role of government and tax policy.

1.3 Robustness to variations in survey question

Here, I show how two variations to the survey question from Figure 1 provide further insight into the result just described.⁹

1.3.1 The salience of payments versus outcomes

First, 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 affects respondents' moral judgments. For example, emphasizing after-tax incomes may make respondents 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 approach's assumption that only after-tax, not pre-tax, incomes 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: 95% in the payments-only version and 65% in the outcomes-only version. Approximately two-fifths choose an amount for Person A to pay between \$12,000 and \$24,000: 43% and 40% respectively. 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.

These results suggest that asking respondents to engage with both payments and outcomes (i.e., after-tax incomes) 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.

⁹I do not discuss the effects of minor wording and framing changes across the rounds of the survey that had negligible effects on the results. For example, replacing "have to pay" with "pay" in the text of the scenario from Figure 1 may make the payments seem less compulsory, but that change had no noticeable effect on the results.

1.3.2 Discrete versus continuous choices

Second, I modify the scenario in Figure 1 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 \$9,000, \$12,000, \$18,000, or \$24,000. In a second variation, I add two additional (intermediate) choices: Person A pays \$15,000 or \$21,000. In a third, I add two more (extreme) choices: Person A pays \$6,000 or \$27,000. Across these three versions, I obtain 266 responses.

These variations are intended to address two concerns. First, respondents to the setup shown in Figure 1 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. Second, 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 A pay the precise cost of the project.

Consistent with the first of these concerns, these variations reduce the concentration of answers at \$18,000, but contrary to the second concern they increase the shares of respondents choosing the simplest options: full equalization of after-tax incomes and equal absolute tax payments. The modal choice in these variations—chosen by just under 50% of respondents—is Person A paying \$24,000, the full equalization recommended by the conventional approach's assumed normative view. The remaining respondents—just over 50%—continue to stop short of fully offsetting Person A's luck-based advantage. In sum, these variations yield weaker versions of the overall results but leave intact the main finding that a large share of respondents resist costless redistribution.

While the discrete choice versions of the survey clearly have advantages, these results make clear that there are countervailing considerations complicating 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 (especially than those in which the interdependence of allocations is made clear by the amounts automatically adjusting), potentially causing respondents to choose quickly and without consideration. Second, any set of choices unavoidably introduces elements of framing, and (as the data suggest) listing the options may privilege those with "simple" features such as equal payments or equal after-tax incomes because they are easier to understand.

1.3.3 Interpretation

The results of this section show a robust resistance among a large share of survey respondents to equalizing outcomes even when the policy problem is unconstrained by the practical concerns of efficiency and desert that typically matter in taxation. One possible explanation is that respondents may disagree with the objective assumed in conventional optimal tax analyses. The next section turns to evidence on that possibility.

2 Preference for CBBT as a principle of optimal taxation

I now turn to the second novel finding of this paper: a widespread preference among survey respondents for a classical benefit-based logic for taxation over the conventional approach's logic of 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 brief formal treatment in the Appendix and 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

CBBT is the combination of two ideas: taxes ought to be based on the benefit an individual obtains from the activities of the state; and the best measure of that benefit is how much the state's activities increase the economic opportunities (i.e., the income-earning ability) of the individual.

As mentioned in the Introduction, CBBT has played and continues 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."¹⁰

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."

The normative appeal of CBBT, according to its advocates, is due in part to its avoidance of coercive taxation. Of course, benefit as defined here is unobservable, so that the second-best CBBT tax system will be coercive in one sense. Nevertheless, CBBT's supporters argue that there remains an essential difference between the benefit-based system and one that maximizes a consequentialist social welfare function. The goal of the former, but not the latter, is for an individual to pay an amount for the activities of the state that is determined by his or her willingness to pay (i.e., marginal rate of substitution).

For the purposes of this paper, it is important to note that optimal taxes under CBBT do not depend on the distribution of after-tax incomes. Instead, they are defined by the relationship

¹⁰See the debate over the phrase "you didn't build that" in the 2012 U.S. presidential election, as discussed in Weinzierl (2016).

between individuals' innate abilities and pre-tax incomes, in stark contrast to the conventional normative approach in optimal tax.

While there exists evidence for CBBT playing a role in elite rhetoric and thinking on tax policy, no direct evidence on its appeal to the public has been gathered. I turn to that evidence next.

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:

Here are two reasons why some people choose, as you did, an outcome in which Person A pays more than Person B. Please read these two reasons carefully, give them some thought, and then choose the reason with which you *more strongly agree*.

- Person A didn't do anything to deserve ending up with more than Person B, and a dollar matters less in the hands of someone with more, so Person A should pay more than should Person B.
- Person A gains more than Person B if they accept the offer, and payments should be tied to how much each person benefits, so Person A should pay more than should Person B.

Figure 3: Respondents are asked to choose a logic for progressive payments in the tax scenario.

The first of the two reasons refers to the logic of diminishing marginal social welfare of income applied in conventional optimal tax analyses, while the second reason refers to the CBBT principle.¹¹ After respondents make their choice on the question in Figure 3, 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 with each of the reasons in Figure 3.

In response to this question, 71% of respondents prefer the classical benefit-based justification to the conventional one. Moreover, as Table 3 shows, nearly 70% of the 29% who prefer the conventional logic state that they either agree or strongly agree with CBBT, bringing the total share of respondents expressing support for CBBT to over 90%. In contrast, less than half of those who prefer CBBT express agreement with the conventional logic.

¹¹The order of these two reasons was randomized. 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."

Table 3: Widespread support for CBBT over conventional logic

	CBBT logic	Conventional logic
Preferred logic	0.71	0.29
Opinion on logic if not preferred		
Strongly agree	0.10	0.07
Somewhat agree	0.59	0.40
Somewhat disagree	0.26	0.42
Strongly disagree	0.04	0.11
Total expressing agreement	0.91	0.63

Notes: The final row is a sum: e.g., 0.91 is the share preferring CBBT (0.71) plus the share of those preferring the conventional logic who either strongly or somewhat agree with CBBT ($0.29 \times (0.10 + 0.59)$)

These results suggest that, when reasoning over tax policy, Americans are more comfortable with the logic of CBBT than with the logic typically applied in modern optimal tax analyses.

2.2.1 Robustness to survey design

Changes to the overall survey design (as described in Section 1) and to the wording of this question generate variation around the 71% overall figure, but the share of respondents preferring CBBT lies between 62% and 79% in all versions. The smallest share (62%) is obtained in the discrete options version described in Section 1.3.3. The largest shares are obtained in either the outcomes-only version described earlier (78%), or in a version that excludes the phrase "didn't do anything to deserve ending up" and replaces it with "ends up" in the first option in Figure 3 (79%).¹²

2.2.2 Robustness across respondent traits

Table 4 and Table 5 summarize support for CBBT across self-reported demographic traits and political views. Statistical analysis in Section 3 will largely support these simple cross-tabulations, though with some suggestive exceptions.

Table 4: Shares of respondents preferring CBBT logic for progressivity, by demographic trait

	Gender		Age		Race			Education		Inc as child		Inc as 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

¹²The former phrase emphasizes that Person A did not affirmatively deserve such a lucky outcome, so its lowering expressed support for CBBT is not surprising. Note that this clause would have been appropriate to include in the description of the second reason, so the version yielding 79% support for CBBT is, arguably, the neutral setup.

Table 5: 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	0.68	0.69	0.70	0.70	0.70	0.70
Obs	1043	583	499	621	602	983	1403	803	1629	580

These patterns suggest that support for CBBT is largely universal, consistent with both the substantial support it receives among those respondents who prefer the conventional logic and its use by elite political figures in the United States.

3 Link between results: the moral significance of pre-tax income

In this section, I show that the two previous sections' findings are linked: that is, respondents who more strongly resist equalization of inequality due to brute luck are more likely to prefer the classical benefit-based logic for taxation. After presenting evidence of this linkage, I discuss one potential interpretation of it that relates to the debate over whether pre-tax incomes ought to be considered morally relevant for tax design, and I address a few interpretive questions.

Visual evidence of this linkage is shown in Figure 4, which gives the share of respondents preferring the CBBT principle for six ranges of answers to how much Person A should pay.




Figure 4: 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. This pattern holds across variations of the survey questions, including in the discrete choice versions described in section 1.3.3 (where support for redistribution was greater).

Table 6 presents the results of a simple regression analysis revealing that this relationship between resistance to redistribution and support for CBBT is sizeable and statistically significant.

	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 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.

Calculating marginal effects for these results implies that a \$1,000 decrease in "Person A pays" is associated with an increase of 1.1 percentage points in the likelihood that the respondent prefers the CBBT logic (the mean value is 71%). 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 4 across the range from \$12,000 to \$24,000.

The only other significant relationship in Table 5 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 relationship between the principle of Equal Sacrifice and benefit-based taxation was hinted at in Feldstein (1976) and is discussed at greater length in Weinzierl (2016).¹³ But an important aspect of this relationship is worth highlighting here: both CBBT and Equal Sacrifice give moral weight to pretax incomes. I turn to that common feature next.

¹³Weinzierl (2014) presents evidence for the role of Equal Sacrifice in public views of tax policy, and Scheve and Stasavage (2016) present extensive evidence on the importance of Equal Sacrifice as a principle guiding tax policy across a range of countries over the last century.

3.1 Pretax income's moral significance and the burden of proof for desert

One way to interpret the linkage between respondents' views of luck-based inequality and CBBT relates to a current debate in both optimal tax theory and political philosophy over whether pre-tax incomes are relevant to optimal tax policy. Murphy and Nagel (2002) forcefully argue that they are not: "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 Murphy and Nagel's view. A majority effectively grant Person A some entitlement to a purely luck-based advantage, bestowing on pre-tax incomes a moral relevance it is denied by a conventional welfarist objective. Consistent with that position, respondents more willing to accept inequality due to luck also prefer to assign taxes based on CBBT, a principle that seeks to implement a more "voluntary" tax system and defines optimality in terms of the relationship between pre-tax incomes and taxes paid.

In other words, a large share of this paper's survey respondents appear to put some of the burden of proof for desert on the opposite side of where luck egalitarian political philosophers do. For luck egalitarians, inequality due to luck is unacceptable unless proven otherwise and therefore should be offset. Our results suggest that most Americans, in contrast, at least partially endorse the view that inequality due to luck is 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 part of a recent body of work finding evidence for 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 "individuals, placed in the position of a social planner, do in fact respect the reference points of others." Saez and Stantcheva (2015) 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, both capture an aspect of public reasoning over tax policy in which pre-tax incomes are relevant to optimal policy.

In fact, some moral philosophers also resist that claim that pre-tax incomes are morally irrelevant. 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's evidence suggests that a large share of Americans agree.

3.2 Interpretative questions

The results in this paper raise a number of interpretive questions. Here, I address a few of them.

First, would respondents choose differently if the hypothetical situation in Figure 1 had put the unlucky person in a much worse absolute position? In particular, would respondents be more likely to make a net transfer to Person B in that case? These questions highlight an important point. If putting Person B in a state that mimics poverty were to likely yield more support for net transfers from Person A, such a result would be entirely consistent with the conclusions above. As noted, the choices of respondents, both over equalization and their preferred principle of taxation, reinforce the evidence from a wide range of sources that most people balance competing normative principles when making judgments such as these. If Person B were put in a dire position, the force of the egalitarian principles in most people's calculations would increase dramatically, so respondents having Person A pay more would not imply general support for equalization of luck.

Second, what if the resistance to equalization reflects empathy with loss aversion (as in Charité, Fisman, and Kuziemko, 2015), so that respondents' hesitancy to take away from Person A and give to Person B can be explained with a conventional utilitarian logic? Though this possibility may seem like a challenge to this paper's conclusions, it can be seen as another way to state them. The evidence presented here shows that survey respondents are willing to grant that pre-tax incomes have some moral significance and, therefore, that individuals have some justifiable claims with regard to them. This willingness means that a reshuffling of those pre-tax incomes would change the set of claims respondents would grant to individuals.¹⁴ Such a change seems illogical to someone convinced that what matters (from a moral perspective) are after-tax incomes—after all, how could a random reassignment of pre-tax incomes change the optimal assignment of outcomes? But respondents are not convinced that what matters are after-tax incomes; in fact they seem to hold the view that even entirely luck-based pre-tax incomes are reasonable starting points for the determination of taxes (after all, they endorse CBBT as a principle of tax design). So long as respondents view pre-tax incomes as meaningful, explaining their responses as reflecting empathy with loss aversion is simply one way of describing the judgments this paper's results highlight.

Third, how are these results consistent with the value most individuals evidently place on insurance? After all, one of the most familiar justifications for tax policy that offsets innate inequality is that it provides insurance for otherwise uninsurable risk. The debate in political philosophy over this argument is vast, but for the purposes of this paper the important point is that most survey respondents appear to be at least somewhat hesitant to embrace the insurance analogy when it comes to risk at the "starting point," contrary to the arguments not only of Harsanyi but also, famously, of John Rawls (1971). Instead, they appear to treat such risk, even though it is clearly "brute luck" and therefore outside the control of individuals, as just or at least not unjust, consistent with the counterarguments of Nozick (1974).

¹⁴One way in which empathy with loss aversion would be morally distinct is if it continued to apply in a situation where advantages were acquired unjustly, an open and interesting question for future research.

4 Conclusion

The main contribution of this paper is to present new survey evidence of ways in which a large share of Americans—arguably a majority—are ambivalent toward key features of the normative framework that has been generally adopted by modern optimal tax analyses. To the extent that these findings indicate sincere normative diversity in most people’s attitudes toward distributive justice, and to the extent that optimal tax theorists want their models to be consonant with public priorities for taxation, researchers ought to consider capturing that ambivalence in their work, as well.

These results raise many questions that will require substantial time and further study to answer. Do respondents’ answer 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 respondents’ stated preferences for these hypothetical scenarios translate into votes for specific policies and policymakers? Exploring these questions will further improve our understanding of popular reasoning on distributive justice.

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 (2002). *The Myth of Ownership*, Oxford.
- [17] Musgrave, Richard A. (1959). *The Theory of Public Finance*. McGraw-Hill.
- [18] Narveson, Jan. (2010). "Cohen's Rescue," *Journal of Ethics*, pp. 263-334.
- [19] Nozick, Robert (1974). *Anarchy, State, and Utopia*. Basic Books.
- [20] Rawls, John (1971). *A Theory of Justice*. Cambridge: Harvard University Press.
- [21] Saez, Emmanuel and Stefanie Stantcheva (2015). "Generalized Social Welfare Weights for Optimal Tax Theory." Forthcoming, *American Economic Review*.
- [22] Scheve, Kenneth and David Stasavage (2016). *Taxing the Rich: A History of Fiscal Fairness in the United States and Europe*. Princeton.
- [23] Smith, Adam (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*.
- [24] Stantcheva, Stefanie (2016). "Optimal Taxation and Human Capital Policies over the Life Cycle," Forthcoming, *Journal of Political Economy*
- [25] Stiglitz, Joseph (1987). "Pareto Efficient and Optimal Taxation and the New New Welfare Economics," *Handbook of Public Economics*, Chapter 15.
- [26] Sugin, Linda (2011). "A Philosophical Objection to the Optimal Tax Model," *NYU Tax Law Review*
- [27] Weinzierl, Matthew (2014). "The Promise of Positive Optimal Taxation: Normative Diversity and a role for Equal Sacrifice," *Journal of Public Economics*, 118.
- [28] Weinzierl, Matthew (2016). "Revisiting the Classical View of Benefit Based Taxation," *Economic Journal*.
- [29] Werning, Ivan, (2007). "Pareto Efficient Income Taxation," Working Paper, MIT, April.

5 Appendix

5.1 Conventional optimal tax model

The standard modern optimal tax model, modified to allow for preference heterogeneity, has the following structure.¹⁵

A population of individuals differ in two unobservable ways, income-earning ability $w \geq 0$, and preferences for leisure $\theta \geq 0$, jointly distributed according to the density $f(w, \theta)$. Each individual has utility of after-tax income c and pre-tax income y that depends on the product of w and θ according to the utility function $u(c, y, w\theta)$.

A tax authority specifies bundles of pre-tax and after-tax income to maximize a function of individual utilities. The authority's normative judgments may depend on w and θ , not just the product $w\theta$. The objective is:

$$\max_{\{c(w\theta), y(w\theta)\}} \int_0^\infty \int_0^\infty G(U(c(w\theta), y(w\theta), w\theta), w, \theta) f(w, \theta) dw d\theta, \quad (1)$$

where $G(\cdot)$ is assumed to be an increasing and concave function of utility. The tax authority faces a feasibility constraint:

$$\int_0^\infty \int_0^\infty (y(w\theta) - c(w\theta)) f(w, \theta) dw d\theta \geq 0, \quad (2)$$

and incentive compatibility (IC) constraints that guarantee individuals choose labor supply optimally:

$$U(c(w\theta), y(w\theta), w\theta) \geq U(c(w'\theta'), y(w'\theta'), w\theta), \quad \forall w, w', \theta, \theta'. \quad (3)$$

To solve for the optimal tax policy, (1) is maximized subject to (2) and (3).

5.2 CBBT optimal tax model

To solidify ideas, it may help to show how CBBT can be represented within the apparatus of modern tax theory. As in the standard setup, suppose a social planner chooses taxes and the level of public spending. Individuals are differentiated by ability w , indexed with i , and derive utility according to

$$U(c_i, l_i) = u(c_i) - v\left(\frac{y_i}{w_i}\right), \quad (4)$$

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. Individuals take the tax 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^* .

To capture CBBT in this setup requires two novel steps. First, we make individuals' heterogeneous income-earning abilities endogenous functions of both endowed ability and public goods spending. Formally,

$$w_i = f(a_i, G),$$

¹⁵This follows Mirrlees (1971), Saez (2001), and Lockwood and Weinzierl (2016).

where $i \in I$ now 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.

Second, we apply the method of Lindahl (1919) to determine the first-best optimal allocation under CBBT. That method has us consider a hypothetical scenario in which each individual i is assigned a share of total taxes to be paid, τ_i , and then allowed to choose the level of public goods provision that maximize her utility subject to her personal budget constraint taking τ_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 goods¹⁶; 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 endowed ability, $\theta_i^{G,a}$, as:

$$\theta_i^{G,a} = \frac{f_{G,a}(a_i, G) f(a_i, G)}{f_G(a_i, G) f_a(a_i, G)}, \quad (5)$$

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 γ . Note that the setup of the scenario in Figure 1 implies that the ability production function $f(a_i, G^*)$ takes this multiplicative form.

¹⁶It 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).