CES ifo Working Papers

12645
2026

April 2026

Revealed Preference for Redistribution and the Role of Government

Matthias Rodemeier, Gregory Sun

CES ifo

Imprint:

CESifo Working Papers

ISSN 2364-1428 (digital)

Publisher and distributor: Munich Society for the Promotion
of Economic Research - CESifo GmbH

Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740

Email office@cesifo.de
https://www.cesifo.org

Editor: Clemens Fuest

An electronic version of the paper may be downloaded free of charge
· from the CESifo website: www.ifo.de/en/cesifo/publications/cesifo-working-papers
· from the SSRN website: www.ssrn.com/index.cfm/en/cesifo/
· from the RePEc website: https://ideas.repec.org/s/ces/ceswps.html

Revealed Preference for Redistribution and the Role of Government

Matthias Rodemeier and Gregory Sun

March 2026

Abstract

We study how Americans trade off the welfare of the poor versus the rich using incentivized transfer experiments. Combining this with estimates of the Elasticity of Taxable Income, we quantify optimal income tax rates in the US. Revealed preferences show strong concern for the poor across the political spectrum, implying tax rates far more progressive than existing political agendas. This creates a puzzle: individuals vote for policies that are less progressive than their distributive preferences imply. We trace this puzzle to aversion toward government-mediated redistribution. Liberals' support for very progressive tax rates is dampened by misuse of public funds, while conservatives object to taxation on principled grounds tied to coercion and property rights. Our paper illustrates that disagreement over redistribution mostly reflects disagreement over *institutions* rather than over helping the poor.

JEL Codes: D63, D72, D90, H21, H30, P35

Keywords: Income Taxation, Redistributive Preferences, Social Welfare Weights, Government Aversion, Political Economy

---

Rodemeier: Bocconi University, CESifo, IZA, BAFFI and IGIER. Sun: Olin Business School, Washington University. We thank Sara Bernardoni for her exceptional dedication as a research assistant. We are grateful to Nathan Hendren, Claus Kreiner, Timm Gries, Ben Lockwood, Benjamin Moll, Krishna Srinivasan, Guido Tabellini, Dmitry Taubinsky, and Edoardo Teso for very helpful comments.

1 Introduction

Redistributive policy requires society to confront a fundamental normative question: how should gains to the poor be traded off against losses to the rich? This question lies at the core of debates over taxation, transfers, and the scope of government intervention. Understanding these preferences matters because modern normative frameworks, such as optimal income taxation, translate distributive views into concrete policy prescriptions through social marginal welfare weights. These weights summarize how much society values an additional dollar accruing to individuals at different income levels. Identifying to what extent people disagree over trading off the welfare of poor versus rich individuals is essential for interpreting observed policy choices and political discourse.

This paper provides evidence using a series of incentivized transfer experiments conducted on a large sample of over 9,000 U.S. citizens, representative on key demographic observables. Participants make real allocation decisions that trade off income between low- and high-income recipients. These choices allow us to recover revealed preference estimates of how individuals value welfare across the income distribution. To map these estimates into a standard income tax framework, we collect over 1,200 estimates of the elasticity of taxable income (ETI) from the literature, which tells us how people across the income distribution respond to taxation. This brings together the two pillars of income taxation: the normative dimension (distributive preferences) and the positive one (behavioral responses to taxes).

The first main finding is that support for the poor is ubiquitous across the political spectrum. Both liberals and conservatives place substantially higher weight on the welfare of the poor than on the welfare of the rich. When translated into optimal tax formulas, these revealed preferences imply levels of redistribution and top tax rates far higher than those observed in practice. Indeed, when the same individuals are asked directly about tax policy, they support much lower taxes on high-income earners than what we recover from their revealed preferences from transfer choices. In other words, *people do not support the taxes that their generosity towards the poor would imply*. For instance, revealed preferences from transfer choices imply average tax rates of around 60% for households earning $500,000, yet the same subjects state they prefer rates of only 21-33% on these households when asked directly about tax policy. This distributive gap constitutes a central puzzle of this paper. Crucially, the importance of this gap can only be established through the quantitative framework we develop: without translating incentivized allocation choices into implied optimal tax schedules, we would not know that stated and revealed tax preferences diverge in an economically meaningful way. Measurement is therefore at the core of this paper.

A natural question is whether this gap reflects incorrect beliefs about the efficiency costs

of taxation or misperceptions of societal norms. In a second experiment (N=2,456), we randomly provide subjects with information about the elasticity of taxable income from our meta-analysis or about the redistributive preferences revealed in our first experiment. In line with the literature, information about the ETI reduces redistributive preferences by raising attention to efficiency costs, while information about social norms increases them as people underestimate other's generosity towards the poor. However, quantitatively, the resulting optimal tax schedules remain in a similar range and far above the rates subjects say they support when asked directly. The gap between distributive and political preferences thus cannot be explained by incorrect beliefs about behavioral responses or social norms.

The gap can also not be explained by narrow self-interest whereby subjects desire lower taxes across the income distribution simply in order to avoid paying higher taxes themselves. Even among older, low-income respondents who would be net beneficiaries of progressive taxation, stated tax preferences are far below what their incentivized transfer choices imply. Rather, we show that our key puzzle is resolved by accounting for aversion to government institutions. Through additional experiments, we demonstrate that opposition to taxation is not primarily driven by disagreement about the desirability of helping the poor, but by resistance to the government as the implementing institution. In these experiments, some funds that go to poor recipients are tied to giving money to the US Treasury, and we truthfully implement these donations to the Treasury in our experiment. This change in experimental protocol reveals that subjects actively avoid transfers tied to government revenue, which translates into substantially lower implied tax rates that are close to the rates people vote for. Aversion to government can explain over 80% of the gap between transfer-based revealed preferences and self-reported ideal tax rates. As a result, support for less progressive tax rates does not reflect lower welfare weights towards the poor, but mostly an aversion to government taxation per se.

The sources of government aversion differ across the political spectrum. In a final experiment, we ask subjects to assume the government would be perfectly efficient. This treatment eliminates government aversion for liberals but does not change behavior at all for conservatives. Liberals seem to be driven by concerns about inefficiency, or misuse of public funds. When redistribution is presented as perfectly efficient, support for government-mediated transfers closely aligns with the redistributive preferences revealed in private transfer tasks. For conservatives, by contrast, government aversion reflects a more fundamental objection that persists even under assumptions of perfect efficiency. Additional results illustrate that these people have principled views that emphasize property rights, voluntary exchange, and opposition to coercion, and that view taxation as intrinsically problematic rather than merely imperfectly implemented.

Consistent with this interpretation, conservatives express a strong preference for *non-governmental* approaches to addressing inequality. When asked how they would choose to help low-income

families, conservatives disproportionately favor voluntary charity over government programs, while liberals are more likely to endorse public redistribution. These divergent institutional preferences explain why even a hypothetical, perfectly efficient government would fail to attract broad support for higher taxes among conservatives.

To provide a portable framework for quantifying these institutional preferences, we introduce the *Equivalent-Transfer Metric of Government Funds* that measures how much money would need to go directly to a household with a specific income to make a voter indifferent to instead giving $1 to the government. Under assumptions of perfect government efficiency, liberals are roughly indifferent between $1 going to the US government versus $1 going directly to a household that earns below $250,000/year. As recipient income rises, liberals increasingly prefer government revenue over direct transfers: Liberals would require $86 to go to a millionaire to feel indifferent to $1 going to the government. For conservatives, the metric remains essentially flat near zero across the entire income distribution. Even for very wealthy recipients, conservatives derive virtually no utility from tax revenue going to the government.

By combining incentivized transfer choices with policy preferences and institutional framing, this paper clarifies the *difference between distributive values and support for redistributive policy*. The results suggest that disagreement over taxation reflects not only differences in how people value the poor relative to the rich, but also deep differences in how they view the role of government itself.

Literature This paper ties into the literature that uses experiments and surveys to understand the drivers of redistributive preferences, such as normative judgments, and fairness concerns (e.g., Alesina and Angeletos, 2005; Alesina and La Ferrara, 2005; Alesina et al., 2018; Mengel and Weidenholzer, 2023; Cappelen et al., 2024). Closely related, Kuziemko et al. (2015) and Stantcheva (2021) document in surveys that people who are concerned about inequality may not support inequality-reducing policies and that this relates to distrust of the government. We differ from this important work by directly measuring the underlying distributive trade-offs using incentivized allocation choices that recover revealed preferences and then map them into welfare weights. Only this new approach allows us to recover people's implied optimal tax rates and quantitatively contrast them with the tax rates the same people would vote for. This allows us to document the dramatic gap between distributive and political preferences. Further, by changing the allocation technology (private transfer vs. government-mediated transfer) we can infer how much people value $1 in tax revenue relative to $1 in a direct transfer to a family. This approach is new in that it puts a "price on public funds" and provides a portable quantitative framework for future work.

Another important addition to prior work is that we uncover the motives behind people's

distrust and institutional resistance. Our results establish that these motives are where most of the partisan polarization regarding redistribution stems from, and that they can be classified into efficiency-based versus principled foundations. This provides an explanation for resistance to redistribution that differs from prior explanations, such as whether inequalities are caused by luck versus merit (Alesina and Angeletos, 2005), intergenerational mobility (Alesina et al., 2018), or whether redistributive policies help in-group versus out-group members (Stantcheva, 2024).

The complementary working paper by Capozza and Srinivasan (2025) uses a related elicitation method to quantify social marginal welfare weights from transfer choices. Our paper advances this work by showing that transfer-choice–based welfare weights overstate support for progressive tax schedules because people are averse to government. The unique contribution of our paper is, therefore, the finding that estimated welfare weights do not generally produce *politically feasible* income tax schedules.[1]

Our work also relates to alternative measures of welfare weights that infer weights from observed policies and tax elasticities (e.g., Bargain et al., 2011, 2014; Lockwood and Weinzierl, 2016; Jacobs et al., 2017; Hendren, 2020). The idea behind this "inverse optimum approach" is to find the set of social welfare weights that would make the status quo tax policy welfare-maximizing. However, researchers rarely interpret resulting estimates as society's true preferences, as the approach relies on the strong assumption that tax policy has been set optimally, which is likely violated due to political economy constraints and other frictions. Saez and Stantcheva (2016) develop a theoretical framework of Generalized Social Marginal Welfare Weights (GSMWW), which directly incorporates political economy concerns and even non-utilitarian principles such as Libertarianism. Their framework calls for a direct measurement of the primitives rather than an indirect one via observed policies and tax elasticities. We offer such direct measurement in a revealed preference framework and point to a new challenge: while the GSMWW framework of Saez and Stantcheva (2016) is in principle flexible enough to encode institutional and philosophical considerations into the weight function, our results show that welfare weights recovered from private transfer experiments do not coincide with the weights that would rationalize observed tax policy. The gap arises because transfer-based elicitation does not activate the institutional preferences that shape support for actual policies. Our approach takes a first step toward closing this identification gap by introducing experimental variation in which redistribution is mediated through actual donations to the U.S. Treasury. This extension allows institutional considerations to be reflected in the recovered welfare weights and, consequently, in implied optimal income tax schedules.

[1] Our collection of over 1,200 elasticity estimates that we use for the estimation of the optimal income tax schedule represents another difference to their work.

Finally, our paper complements work that studies policymakers' uncertainty about the key parameters of income taxation. Lockwood et al. (2021) incorporate uncertainty over the elasticity of taxable income into optimal tax design and measures it through expert surveys with economists. We complement this work by studying the normative side of income taxation and show how disagreement over distributive preferences maps into uncertainty about optimal tax rates.[2]

2 The Mirrleesian Framework: A Roadmap

To motivate the rest of this paper, it is useful to briefly revisit Mirrlees (1971)'s optimal income taxation framework, which is arguably the workhorse model in public finance for thinking about redistributive taxation. Under this framework, all questions about redistribution can be boiled down into two empirical ingredients: i) how much do people care about the rich versus the poor, and ii) what are the behavioral responses of taxpayers to changes in tax policy. A key insight of the upcoming results is that this model misses other ingredients that are *quantitatively* important for optimal income taxes.

Individuals differ by a scalar type θ and have quasi-linear preferences $U(c, z; \theta) = c - \theta v(z)$ over consumption c and pre-tax income z. Given a tax schedule $T(z)$, each type chooses income to solve $\max_z z - T(z) - \theta v(z)$, generating an endogenous income distribution with density $h(z)$ and CDF $H(z)$. Our specification ensures that at most one type generates a given income level z; denote this type by $\theta(z)$. The social planner sets $T(z)$ to maximize

$$\max_{T(\cdot)} W\big(T(\cdot)\big) = \int_0^\infty \Gamma\big(z - T(z) - \theta(z)v(z)\big)\, h(z; T(\cdot))\, dz, \quad \text{s.t.} \quad \int_0^\infty T(z)\, h(z; T(\cdot))\, dz \geq E,$$

where $\Gamma(\cdot)$ is strictly increasing and concave, and we normalize average welfare weights to one. The solution decomposes as a product of three factors (see Diamond, 1998; Ferey et al., 2024):

$$\frac{T'(z)}{1 - T'(z)} = \underbrace{[1 - \tilde{G}(z)]}_{\text{redistributive preferences}} \times \underbrace{\frac{1}{\varepsilon(z)}}_{\text{behavioral responses}} \times \underbrace{\frac{1}{\alpha(z)}}_{\text{pareto parameter}}. \tag{1}$$

The first factor contains $\tilde{G}(z) \equiv \frac{\int_z^\infty \Gamma'(s - T(s) - \theta(s)v(s))h(s)\,\mathrm{d}s}{1 - H(z)}$, the average welfare weight for individuals earning above z. This isolates the contribution of normative preferences: the integrand $\Gamma'(\cdot)$ captures the marginal social value of transferring one dollar to a household with income z. In Section 3, we design transfer experiments that directly estimate this object. The second factor, $\varepsilon(z) \equiv \frac{v'(z)}{zv''(z)}$, is the elasticity of taxable income at z, summarizing behavioral responses

[2] Our large-scale meta-analysis of ETI estimates provides an additional scientific benchmark for behavioral responses to taxation that complements Lockwood et al. (2021).

to taxation. We estimate this function through a large-scale meta-analysis in Section 4. Through an additional experiment, we also test whether subjects' beliefs about the ETI are accurate and whether this matters for their willingness to redistribute. The third factor, $\alpha(z) \equiv \frac{zh(z)}{1-H(z)}$, is the local Pareto parameter, which can be computed directly from observed income data. Section 5 combines all three inputs to compute optimal tax schedules.

The empirical tension. The results of this exercise reveal that individuals across the political spectrum place substantially greater weight on the welfare of the poor than on the welfare of the rich. Taking the Mirrleesian framework to the data then implies optimal marginal tax rates that are far more progressive than existing U.S. tax policy. Yet the same individuals, when asked directly about their preferred tax rates, endorse levels of taxation well below what their revealed distributive preferences would imply. This gap between choice-implied and stated tax preferences constitutes the central puzzle of the paper.

We provide evidence that this resistance is driven by an explicit aversion to government. To formalize this, we augment the planner's problem: for a given tax policy $T(\cdot)$, let $\mathcal{T}_+(T) \equiv \{z : T(z) > 0\}$, and modify the objective to

$$\int_0^\infty \Gamma\big(u\big(z - T(z)\big) - \theta v(z)\big)\, h(z)\, dz - \lambda \int_{\mathcal{T}_+(T)} T(z) h(z)\, \mathrm{d}z$$

where $\lambda \geq 0$ captures *government aversion*: a deadweight loss agents associate with taxation per se, beyond the efficiency costs in $\varepsilon(z)$. In Section 6, we estimate λ by introducing transfer choices where part of the funds are donated to the U.S. Treasury, and show that this can resolve most of the gap between distributive preferences and support for redistributive policy.

3 Measuring Revealed Distributive Preferences

Conceptual Framework We model subjects in our experiments as making choices about the distribution of available resources according to marginal social welfare weights. This requires us to parameterize the welfare weight function. Following the Mirrleesian framework, we specify a CARA function in other people's baseline income.[3] Let i index an individual subject in our experiment that can choose to transfer resources to family j. Then, the marginal social welfare weights of individual i are assumed to be of the form

$$g_i(C_j) = \exp(-\gamma_i C_j) \tag{2}$$

[3] We also explore a CRRA function in the Appendix H for robustness.

where C_j is (pre-transfer) income of family j. The parameter γ_i is the absolute risk aversion parameter of individual i, measuring how i's preferences for household j change as j's income level increases. If $g_i(\cdot)$ is concave, subject i puts larger welfare weights on poorer households. Note that we condition only on household income in our elicitation, abstracting from other household characteristics such as family size, age, or employment status. The welfare weights we estimate therefore represent average social valuations for each income group, integrating across the heterogeneity of households within that income level.

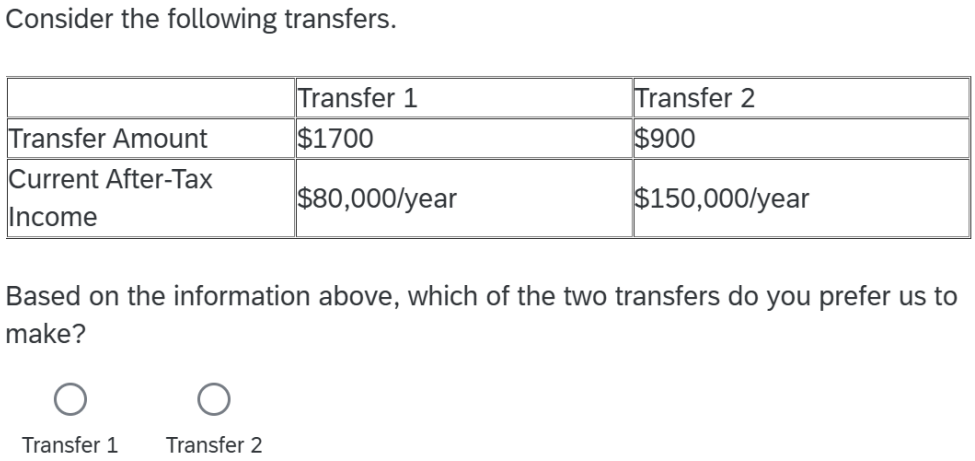
Identifying Welfare Weights from Choice Data The core of our methodology relies on pairwise choice tasks that reveal how individuals trade off transfers to households at different income levels. Each participant completed 10 pairwise choice scenarios. In each scenario, respondents were presented with two potential transfers to two different U.S. households and asked to choose which transfer they preferred to implement. The scenarios varied along two dimensions:

1. **Transfer amounts**: T_{iq1} and T_{iq2} for households 1 and 2, respectively. We randomly draw transfers that are multiples of \$100 from the interval \$100 to \$2,000 with equal probability.

2. **Recipient income levels**: The current after-tax income C_{iq1} and C_{iq2}. These values are drawn from the interval \$5,000 to \$500,000 per year, in \$5,000 increments.

An example of a choice scenario is given in Figure 1 where a subject could either allocate a transfer of \$1,700 to a household with net income of \$80,000 per year, or a smaller transfer of \$900 to a richer household earning \$150,000 per year. Transfers and income were randomly drawn for each choice set and could also include larger transfers to richer households.

Figure 1: Example of Choice Screen


Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	\$1700	\$900
Current After-Tax Income	\$80,000/year	\$150,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

Notes: This figure shows an example of one of the 10 transfer choices subjects could make in the experiment.

If the subject answers this question according to her redistributive preferences, she would transfer to household 1 over household 2 if and only if

$$T_1 \exp(-\gamma_i C_1) - T_2 \exp(-\gamma_i C_2) > 0$$

We allow for the possibility that there is some noise in choices, so in reality, the subject chooses redistribution to household 1 if and only if

$$T_1 \exp(-\gamma_i C_1)\varepsilon_{iq1} > T_2 \exp(-\gamma_i C_2)\varepsilon_{iq2}.$$

We assume that for each individual i, $\log \varepsilon_{iq1}, \log \varepsilon_{iq2} \overset{i.i.d.}{\sim} T1EV(\sigma_i)$. Here, σ_i is a scale parameter on the T1EV distribution, which may vary across individuals.

Taking logs and dividing both sides of the inequality by σ_i, the above is equivalent to individuals choosing to redistribute to family 1 if and only if

$$-\frac{\gamma_i}{\sigma_i}(C_1 - C_2) + \frac{1}{\sigma_i}[\log(T_1) - \log(T_2)] + \frac{1}{\sigma_i}[\log \varepsilon_{iq1} - \log \varepsilon_{iq2}] \geq 0.$$

We divide by σ_i in deriving the above expression because the distributional assumption on the ε's implies that the last term above follows a standard logistic distribution, hence our model boils down to a fairly standard random coefficients logit model without an intercept. The absolute value of the ratio between the random coefficient in front of $(C_1 - C_2)$ and the random coefficient in front of $[\log(T_1) - \log(T_2)]$ represents the individual-level distributive preference, γ_i.

Incentive Compatibility To make choices consequential, participants were informed that we would randomly select two families from the sample of the survey provider with income levels matching those in one of the scenarios. One of the selected families would receive the exact transfer amount chosen by a randomly selected participant, while the other would receive nothing (see also Capozza and Srinivasan, 2025, for a similar design). Paying a random subset of choices is a well-established and widely used approach in experimental methods and commonly produces the same results as if all choices are deterministically paid out (Charness et al., 2016). Importantly, subjects were informed that we would only transfer the money to people who are not participating in the current survey, so there was no chance that subjects could benefit from the transfers themselves. To assure that subjects understood that their choices had real consequences in expectation, we included i) an extensive explanation of the incentive mechanism, and ii) a comprehension check about the consequential nature of the transfer choices (see Appendix

J).[4]

Additional Questions The experiment further elicited a wide range of characteristics such as income, employment, education, age, and, most importantly, political preferences.

We also included attention checks throughout the experiment. Subjects who failed these checks were disqualified from the experiment and are not included in the analysis. Only 2 subjects failed any attention check.[5]

All experiments were pre-registered at the AEA RCT Registry.[6]

3.1 Results

Sample We designed four experiments with a total of 9,212 subjects to elicit social welfare weights from a representative sample of the U.S. population. Data collection started in April 2025 and ended in January 2026. We recruited subjects through Prolific, a widely used survey provider in social sciences (Peer et al., 2022). Our sample size of the first experimental wave consists of 2,421 subjects. Before illustrating the estimation of social welfare weights, we provide a set of reduced-form analyses.

Sensitivities to Recipient Income and Transfer Size Our first question is to what extent differences in income of and transfers to families matter to subjects. To answer this, we estimate the linear probability model:

$$Y_{i1} = \alpha + \beta_1 \Delta C_i + \beta_2 \Delta T_i + \epsilon_i \tag{3}$$

where $Y_{i1} \in \{0, 1\}$ is one if subject i chose the transfer for family 1 rather than family 2. The variables $\Delta C_i = C_{i1} - C_{i2}$ and $\Delta T_i = T_{i1} - T_{i2}$ are the differences in income and transfers between the two families.

Panel A of Figure 2 reports the coefficient on the transfer amount, β_2, for the average subject and also across the political spectrum. On average, a \$1,000 increase in the transfer amount to

[4]We constructed a comprehension check that was relatively difficult to pass as it involved four possible answers that were long to read and where some answers only differed slightly. While this may cause some subjects that understood instructions (but were too impatient to read through all options) to fail the test, our conservative test allows us to condition on the subsample that truly engaged and understood the instructions. 73% of all subjects chose the correct answer, which is typical on these platforms for comprehension (rather than attention) checks.

[5]We further included an open-text question asking participants to describe the main factor influencing their transfer decisions. We do not report these results here, as they provide limited additional insight beyond the choice data.

[6]See trial ID AEARCTR-0015041.

family 1 increases the likelihood for that family to receive the transfer by 4 percentage points (pp). The most liberal survey participants respond the least to transfer variation, with a sensitivity of 2pp. As we move further to the political center, transfer sensitivity increases to 4.8pp for Liberals, and 3.8pp both for Moderate and Conservatives. Subjects identifying as very conservative are most responsive to transfer size, with a sensitivity coefficient of 6.7pp.

Panel B shows income sensitivity, β_1. The likelihood for family 1 to receive a transfer decreases on average by 0.12pp if their income increases by \$1,000. Interestingly, we observe the reversed pattern across the political spectrum that we found in Panel A: Sensitivity decreases as we move from the political left to the right. The most liberal subjects have a sensitivity of –0.16pp. Moving further to the right politically reduces this negative income sensitivity to –0.13pp (Liberals), –0.12pp (Moderate), –0.1pp (Conservative), and –0.08pp (Very Conservative). These results imply that the political left cares relatively less about transfer size while being very sensitive to income. The opposite is true for the political right. Philosophically, this insight means that those on the political left prioritize redistributive concerns based on need (as measured by income levels) over maximizing the economic pie for recipients (transfer size), and view transfers primarily as a tool for addressing inequality.

These results provide the first confirmation that choices in the experiment are meaningful. Subjects choose broadly in line with the political agenda of the party they support. Another takeaway is that redistribution is preferred across *all* political positions, as every income coefficient is statistically different from zero.

Implied Welfare Weights We now show how these reduced-form results translate into implied social welfare weights. We first derive the sampling distribution of the data conditional on individual level parameters γ_i and σ_i from our CARA specification in equation 2. The experiment randomly samples income and transfer amounts for family 1 and 2, (C_1, C_2, T_1, T_2). Since we have 10 questions per subject, let the corresponding 10 draws be denoted by

$$\mathbf{X}_i \equiv (C_{1i}^1, C_{2i}^1, T_{1i}^1, T_{2i}^1, \ldots, C_{1i}^{10}, C_{2i}^{10}, T_{1i}^{10}, T_{2i}^{10}).$$

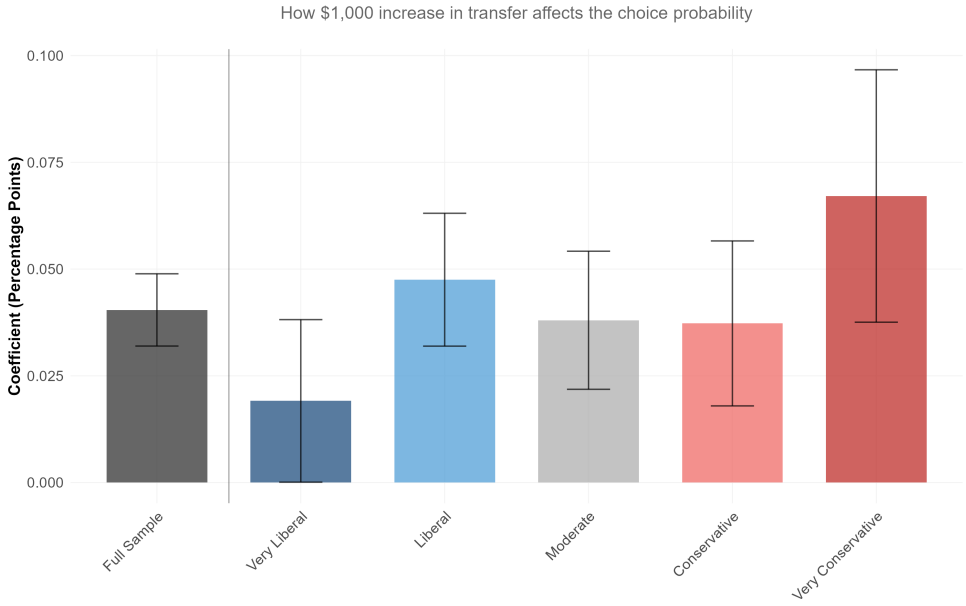
In each scenario, the respondent decided whether to transfer to family 1, Y_{iq}, yielding a vector of dummy variables

$$\mathbf{Y}_i = (Y_{i1}, \ldots, Y_{i10}),$$

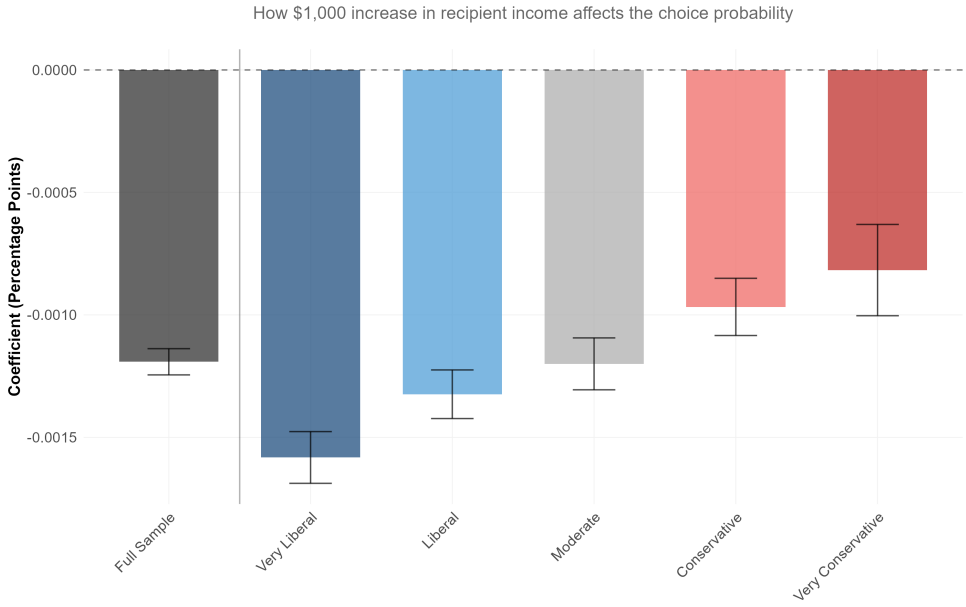
The likelihood function conditional on $\mathbf{X}_i$ and the individual-level parameter values is given by

$$\mathcal{L}(\mathbf{Y}_i|\mathbf{X}_i, \gamma_i, \sigma_i) = \prod_{q=1}^{10} \frac{\exp\left(\frac{-\gamma_i(C_{1i}^q - C_{2i}^q) - \log(T_{1i}^q - T_{2i}^q)}{\sigma_i}\right)}{1 + \exp\left(\frac{-\gamma_i(C_{1i}^q - C_{2i}^q) - \log(T_{1i}^q - T_{2i}^q)}{\sigma_i}\right)}$$

Figure 2: Redistributive Preferences



(a) Transfer Sensitivity



(b) Income Sensitivity

Notes: This figure plots two dimensions of redistributive preferences. In Panel A, we plot the effect of raising the transfer amount to family 1 by $1,000 on the probability that subjects in the experiment choose the transfer for family 1. We plot regression coefficients in percentage points and plot both the effect on the full sample (purple), as well as on the political spectrum. Panel B does the same exercise for an exogenous increase in the income of family 1 by $1,000.

To estimate the distribution of γ_i as flexibly as possible, we approximate the nonparametric joint distribution of (γ_i, σ_i) as being supported on a fine grid of values, and estimate this model using a variant of the approach proposed in Fox et al. (2011) (FKRB).

The basic idea behind the FKRB estimator is to parameterize the joint distribution of (γ_i, σ_i) as being supported on some grid of the form $\boldsymbol{\gamma} \times \boldsymbol{\sigma}$. Given this support restriction, the joint distribution can be represented by the probability mass placed on each support point. Specifically, let $\boldsymbol{\theta}$ be a $|\boldsymbol{\gamma}| \times |\boldsymbol{\sigma}|$ grid of values, where each entry $\theta_{a,b}$ represents the probability $\Pr[\gamma_i = \boldsymbol{\gamma}_a, \sigma_i = \boldsymbol{\sigma}_b]$. The laws of probability imply that $\theta_{a,b} \geq 0$ and $\sum_{a,b} \theta_{a,b} = 1$, but no further constraints on the entries of $\boldsymbol{\theta}$. For a given guess of $\boldsymbol{\theta}$, the likelihood function for a given participant is thus given by

$$\mathcal{L}(\mathbf{Y}_i|\mathbf{X}_i, \boldsymbol{\theta}) = \sum_a \sum_b \theta_{ab} \mathcal{L}(\mathbf{Y}_i|\mathbf{X}_i, \gamma_a, \sigma_b).$$

The procedure behind our estimation is thus as follows. First, we construct a data-driven set of gridpoints, $\boldsymbol{\gamma}$ and $\boldsymbol{\sigma}$, as we explain next. Second, we estimate $\boldsymbol{\theta}$ by maximizing a penalized log-likelihood, subject to the constraints on $\boldsymbol{\theta}$ described above:[7]

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}, \sigma} \sum_{i=1}^{N} \log \mathcal{L}(\mathbf{Y}_i|\mathbf{X}_i, \boldsymbol{\theta}) + \lambda \sum_{a=2}^{41} \left(\sum_b \theta_{a,b} \right)^2 \quad \text{s.t.} \quad \theta_{a,b} \geq 0, \forall a, b, \ \sum_{a,b} \theta_{a,b} = 1.$$

Choice of Grid To apply an FKRB-style estimator, we must choose the grid of points which (γ_i, σ_i) is supported on. There is little formal guidance in the literature on how to do this, so we pick a data-dependent procedure based on what worked well in a pilot version of our study. In the first step, we fit a standard logistic regression without an intercept. The outcome variable is an indicator for whether or not redistribution to family 1 was chosen, and the regressors are $(C_1 - C_2)$ and $[\log(T_1) - \log T_2]$.

Let $\bar{\gamma}$ be the ratio of the coefficients from this regression, and let $\bar{\sigma}$ be the reciprocal of the coefficient in front of $[\log(T_1) - \log T_2]$. Let $\tilde{\boldsymbol{\gamma}}$ be a grid of 40 points that is evenly spaced in log units between $\bar{\gamma}/10$ and $40\bar{\gamma}$. Further, let $\boldsymbol{\gamma} = (0, \tilde{\boldsymbol{\gamma}}, 100)$. Following our pre-registration we restrict our analysis to $\gamma_i \notin \{0, 100\}$. The case of $\gamma = 0$ corresponds to "efficiency maximizing" preferences, and someone with $\gamma = 0$ will always choose the largest transfer. The case of $\gamma = 100$ corresponds to Rawlsian redistributive preferences in that the respondent will always transfer

[7] The intuition behind the penalty term is that the standard FKRB specification has difficulty distinguishing random coefficients which are too similar to one another. Heiss et al. (2022) propose the above penalty term to help smooth over these difficulties by encouraging the estimator to pick estimated distributions of γ_i that place similar mass on similar γ's. This is akin to the use of kernels to smooth out estimated densities when performing kernel density estimation. In line with this intuition, we pick our penalty parameter to decay with sample size according to $\lambda \propto N^{-1/5}$.

to the poorer household. An issue with interpreting these choices as true preferences is that they are also observationally equivalent to choices corresponding to simple "rules of thumb" rather than preferences. The $\gamma = 0$ case corresponds to only paying attention to transfer amount, while the $\gamma = 100$ case corresponds to only paying attention to income. Our downstream analysis takes the distribution of γ_i conditional on $\gamma_i \notin \{0, 100\}$ as our measure of true preferences. In Appendix G, we show that our results are robust to including the entire sample, but we focus on the pre-registered sample restriction in the main part of the paper for consistency.

Finally, let $\boldsymbol{\sigma} = (\bar{\sigma}/3, (\bar{\sigma}/3 + 2\bar{\sigma})/2, 2\bar{\sigma})$ be a grid of error variances. Then we assume that the joint distribution of (γ_i, σ_i) is supported on $\boldsymbol{\gamma} \times \boldsymbol{\sigma}$. This discrete distribution can be represented by a matrix $\boldsymbol{\theta}$ where $\theta_{ab} \equiv \Pr[\gamma_i = a, \sigma_i = b]$, where $\theta_{ab} \geq 0$ for all a, b and $\sum_{a,b} \theta_{ab} = 1$.

Recovering the Welfare Weight Function From the estimated distribution of γ_i, we can recover the average SMWW at any income level z:

$$\hat{G}(z) = \mathbb{E}[\exp(-\gamma_i z)] \tag{4}$$

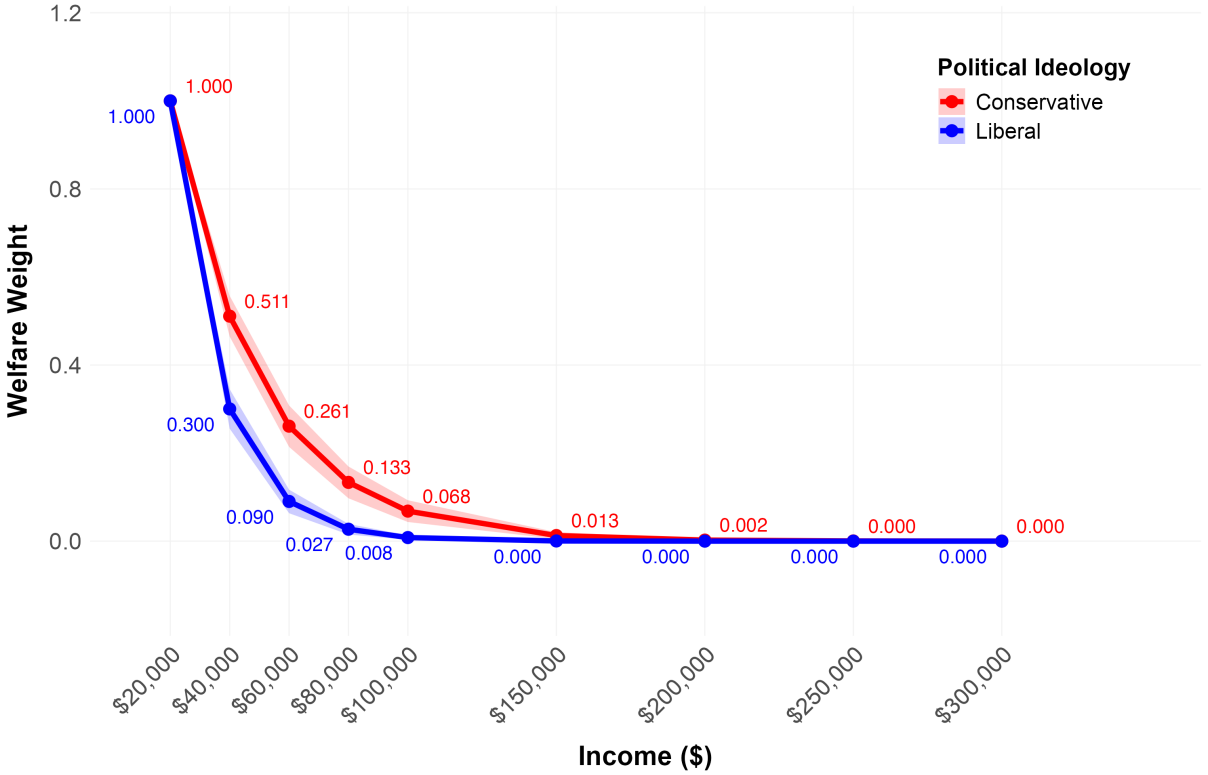
To offer a more intuitive interpretation of these weights, we normalize them so that a low-income individual earning \$20,000 per year receives a weight of 1. Figure 3 plots this normalized function, separately for Liberals and Conservatives. For completeness, we provide the original weights without normalization in Appendix A, Figure 14.

In Figure 3, we include bands for both political groups representing the 5th and 95th percentile of γ_i's within each of the groups. Both functions have negative slope, reflecting that as households become richer, their welfare becomes less important to the decision-maker. Liberals value the welfare of a household earning \$40k/year only 21% as much as a household earning \$20k/year. As expected, for conservatives this number is larger but still as low as 35%. Thus, even republicans quickly discount the welfare of a recipient as income increases.[8] As we further increase income to \$60,000 and \$80,000, weights drop to 10.9% and 6.3% for liberals, and to 19.7% and 11.8% for conservatives. Annual incomes of \$100,000 already have only a weight of 3.8% and 7.4% for liberals and conservatives, respectively.

Above \$150,000, the function becomes very flat and reaches approximately zero at \$300,000 across the political spectrum. This is also where our non-normalized weights, $\hat{G}(z)$, become equal to zero as shown in the Appendix. A welfare weight $\hat{G}(z) = 0$ means that the optimal marginal tax rate is revenue maximizing (the top of the Laffer curve) and proportional to the inverse income elasticity. Thus, normative preferences have a large effect on the marginal tax

[8] We exclude welfare weights at an income of zero in this plot.

Figure 3: Social Marginal Welfare Weights



Notes: This figure plots social marginal welfare weights across the income distribution. All weights are normalized so that an individual earning $20,000/year receives a welfare weight of 1. Weights below an income of $20,000/year and above an income of $300,000/year are excluded in this plot.

rate at the bottom and middle of the income distribution, but matter little at the top.

In conclusion, while we do see the expected differences in welfare weights across political ideology, even the political right assigns substantially larger import to poorer households. As we will see later, these weights imply income tax rates far above what the political agenda of each party would suggest.

4 Measuring How People Respond to Income Taxes

To inform the elasticity term $\varepsilon(z)$ in our tax formula, we conduct a comprehensive Bayesian meta-analysis of the existing empirical literature. We synthesize a large number of scattered point estimates into a coherent, smooth elasticity function $\varepsilon(z)$ defined over the income distribution, together with a full quantification of uncertainty.

4.1 Data Collection

We compile a hand-coded database of 1,205 elasticity estimates from 74 empirical studies. The reason one study has multiple estimates is both because they can come from different parts of the income distribution and because many papers analyze different subsamples (e.g., self-employed workers). We first instructed research assistants to collect the universe of studies that exploit nonlinearities in tax schedules—tax kinks and notches—for identification of the ETI. These estimates have the benefit that they can be directly linked to a position in the income distribution. Of our total sample, 871 elasticities are identified of these nonlinearities. Next, we instructed research assistants to search for papers that exploit tax reforms for identification. This added 334 elasticity estimates. Many of those estimates add important information about the ETI to the far right part of the income distribution where there is no change in the marginal tax rate (i.e., no kink or notch). All studies had to report an elasticity estimate $\hat{\varepsilon}_{si}$, a standard error $\hat{\sigma}_{si}$, and a clear mapping to a point or bin in the income distribution z_{si}.[9]

For each estimate, we record contextual variables including the country, tax system, policy reform exploited, year, estimation method, and whether the elasticity refers to taxable income or broader income measures. We also record a variable indicating which of the estimates at a given income bin is the authors' preferred estimate, which we later use as our main estimate for that bin within a study. We will make this database of tax elasticities publicly available to the academic community once this paper is published, and aim to enrich it further with additional detailed information that can be used by other scholars in future work.

We plot the raw data in Figure 22 in Appendix D and turn to the estimation of the elasticity function next.

4.2 From Scattered Elasticities to a Global Elasticity Function

The end product of our meta-analysis is a function $\varepsilon(\cdot)$, where for each fixed income level z, $\varepsilon(z)$ represents the best guess of the ETI conditional on an individual having income z, given estimates from the data. Unlike a more standard meta-analysis, the object of interest that we aim to recover here is not just a *number*, but instead, an entire *continuous function*. In order to accomplish this, we adopt a Bayesian approach that borrows from an extensive machine learning literature on Gaussian processes (GP).[10]

[9]The dataset was manually collected by research assistants long before AI-based extraction tools became widely available. Manual collection is subject to potential human coding error but avoids errors arising from AI-automated extraction and hallucinations. We mitigated human error by having multiple assistants independently code overlapping sets of papers and resolving discrepancies via reconciliation.

[10]For more background on the statistical theory behind Gaussian Processes, see, e.g., the review article by MacKay et al. (1998). There is precedent outside of economics for using Gaussian processes as a basis for specifying a prior for the purpose of conducting a meta-analysis. For example, see Salimi-Khorshidi et al. (2011) for an

Here, we provide intuition about what our Bayesian meta-analysis approach does and defer a full discussion of technical details to Appendix C. Just like more standard Bayesian frameworks allow the researcher to place a prior on random variables/vectors, Gaussian processes parameterize priors over random functions. Importantly, GP priors have hyperparameters that encode prior knowledge that the underlying function we wish to estimate is *smooth*. Intuitively, smoothness of the underlying elasticity function $\varepsilon(\cdot)$ implies that getting a signal about $\varepsilon(z)$ is also informative about $\varepsilon(z')$ for a point z' close to z. In our dataset, we observe ETIs at specific income levels, so this idea gives us a principled way to "glue together" evidence from a discrete set of income levels to form conclusions about the global elasticity function.

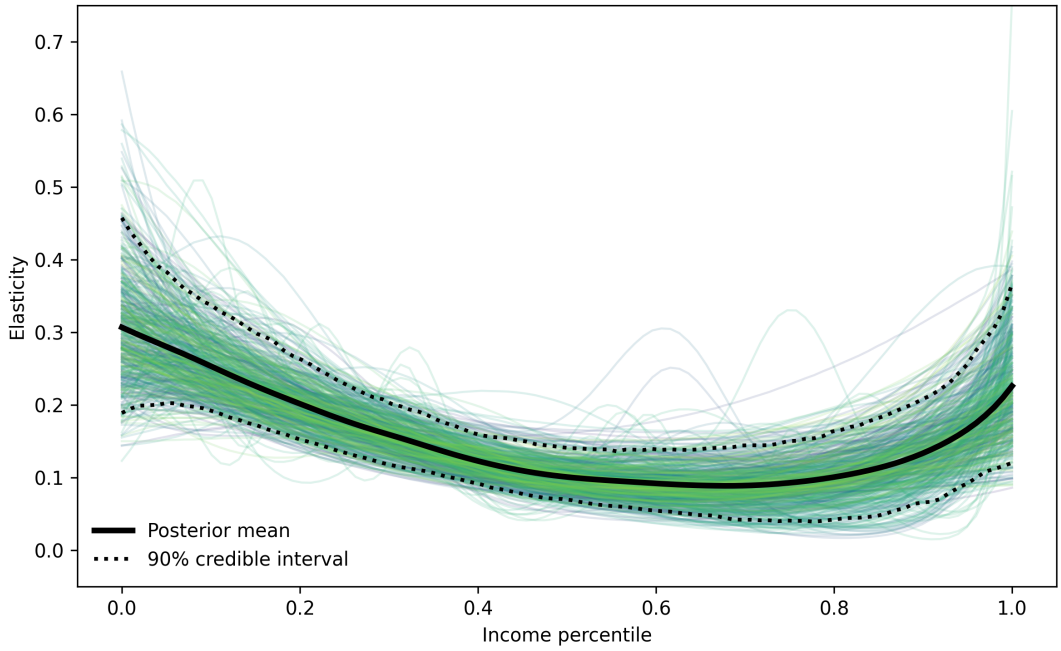
Our basic approach is as follows. First, we specify a diffuse GP prior on the space of possible ETI functions. Second, we derive the likelihood function of a given observation in our dataset. With a specification of the prior and the likelihood, we then use a standard software implementation of GP models to (numerically) apply Bayes Rule and sample draws from the posterior distribution over possible ETI functions. The average elasticity function across these draws provides us with a point estimate of the underlying ETI function, while the variance in the resulting elasticity function across different draws allows us to quantify uncertainty.

The results from doing this are contained in Figure 4. To construct this figure, we draw 100 realizations of the ETI function from the posterior distribution given our sample of ETI estimates. Each draw from the posterior is depicted as a colored curves. The black solid curve is the point-wise mean of the elasticity function across the posterior sample. Dotted lines show 90% Bayesian credible intervals of this mean function. At the bottom of the income distribution, the elasticity is around 0.3 and decreases slowly to 0.1 as we reach the median income. Above the median, the function increases once again to 0.21. The Bayesian credible intervals reject the typical assumption in the literature that the elasticity is constant.

Elasticity estimates are most uncertain in the tails. We cannot exclude with 90% certainty that the elasticity of 0.4 at the bottom could also be 0.19 or 0.5. In the right tail, above the 75th percentile, elasticities might be as low as 0.05 or as large as 0.37. The right tail is particularly noisy and calls for more empirical work on how the 75th percentile of the income distribution responds to taxation. By contrast, we have substantially more certainty over elasticities between the 20th and 60th percentile of the income distribution, which also represents the majority of the US population. Here, the credibility bounds move by only +/- 0.04 around the mean.

example from neuroscience.

Figure 4: Posterior of Elasticity of Taxable Income Function



Notes: This figure summarizes the posterior distribution of the elasticity function $\varepsilon(r)$ given the full sample. The colored lines represent 100 draws from the posterior distribution. The solid black curve corresponds to the pointwise mean of the elasticity function across our posterior sample while the dotted lines represent pointwise 90% Bayesian credible intervals.

5 The Optimal Income Tax Schedule

Given our empirical estimates on welfare weights and elasticities, we solve for a tax schedule $T(z)$ mapping pre-tax income z to the amount of taxes owed by the individual. Let $v(\cdot)$ be a fixed value of the individual-level utility function, and let F_θ be the corresponding distribution of θ types in the population. Given a tax schedule $T(z)$, an individual of type θ endogenously chooses labor supply to solve $z^*(T;\theta,v) \equiv \arg\max_z z - T(z) - \theta v(z)$, which yields utility $U(T;\theta,v) \equiv \max_z z - T(z) - \theta v(z)$. The social planner has redistributive preferences of the same CARA form used to parameterize participants' preferences: $\mathcal{W}(U;\gamma) = -\frac{1}{\gamma}\exp\left(-\gamma(U-\bar{z})\right)$, where $\bar{z}$ is the average level of income at current tax levels.[11] Following the standard setup, we assume that the social planner wishes to maximize average welfare, subject to raising revenue at least equal to $\underline{R}$. We calibrate $\underline{R} \approx$\$15,000, which constitutes the current government proceeds per household in the US *after transfers*. Thus, this can be thought of as the amount currently collected for the purpose of financing public goods, and other government projects.

[11] Note that subtracting $\bar{z}$ from U does not affect the optimal policy, but is useful for maintaining numerical stability of the exponential function.

The social planner's objective is

$$\max_{T} \int \mathcal{W}(U(T;\theta,v);\gamma)\,\mathrm{d}F(\theta) \quad \text{s.t.} \quad \int T(z^*(T;\theta,v))\,\mathrm{d}F(\theta) \geq \underline{R} \tag{5}$$

Let $H(z)$ denote the distribution of endogenously chosen incomes under the optimal tax schedule, let $\bar{G}(z) \equiv \int_{\theta:z(T;\theta)\geq z} \mathcal{W}'(U(T;\theta,v);\gamma)\,\mathrm{d}F(\theta)$ be the social marginal welfare weight for individuals with income above z, and let $\alpha(z) \equiv zh(z)/(1-H(z))$ be the local Pareto parameter. The general solution to social planner's problems like (5) was discussed in Section 2, and rearranging the solution in Equation (1) stated there, the optimal tax policy can alternatively be characterized by the first-order condition:

$$T'(z) = \frac{1-\bar{G}(z)}{1-\bar{G}(z)+\alpha(z)\varepsilon(z)}. \tag{6}$$

While this formula is useful in gaining intuition about the nature of the optimal income taxation, it cannot be directly implemented, because the distribution of incomes $H(z)$, as well as the social marginal welfare weights $\bar{G}(z)$, are endogenous to the tax schedule. In Appendix F, we describe in detail how, for each given draw of the elasticity function $\varepsilon(z)$ and redistributive preference γ, we compute the optimal income tax schedule characterized by Equations (5) and (6). Intuitively, we first recover the structural primitives v and F_θ from the observed income distribution and current tax schedule, then directly optimize over the type-to-income mapping $Z(\theta)$, parameterized as a cubic B-spline, subject to incentive compatibility and the revenue constraint.

5.1 Results

To provide an intuitive comparison, we do the following exercise. We compute two sets of optimal tax schedules: one that would be optimal if redistributive preferences were all low, equal to the 10^{th} percentile of the distribution of γ; and another schedule that would be optimal if redistributive preferences were all high, equal to the 90^{th} percentile. This illustrates differences in tax rates driven by differences in *normative preferences*. Next, for each of these two extreme values of γ, we compute a 90% confidence interval that represents the remaining uncertainty solely from variation in the estimated posterior distribution of ETIs. This tells us how much tax rates vary due to *positive uncertainty*. Figure 5 plots the results for the entire income distribution. To make the plot more intuitive, we plot average rather than marginal tax rates. The thick blue curve represents optimal tax policy among the set of people with strong redistributive tastes, as represented by the 90^{th} percentile of γ (the "political left"). The thick red curve represents optimal tax policy for the subset of people with low redistributive preferences, given by the 10^{th}

percentile γ (the "political right"). Meanwhile, the thin dotted lines demarcate the 10^{th} and 90^{th} confidence intervals due to positive uncertainty arising from variation in the ETI.

The first interesting result is that variance in the optimal tax rates due to preference heterogeneity (i.e., the normative dimension) is substantially larger than variance due to uncertainty over the ETI (i.e., the positive dimension). The dotted 90% confidence intervals are very tight, due to a rich literature that measures the ETI and yields our elasticity function. By contrast, the outer points of the 90% preference interval are much farther apart from each other, suggesting that normative disagreement is the key determinant that drives uncertainty for a social planner on how to design income taxes.

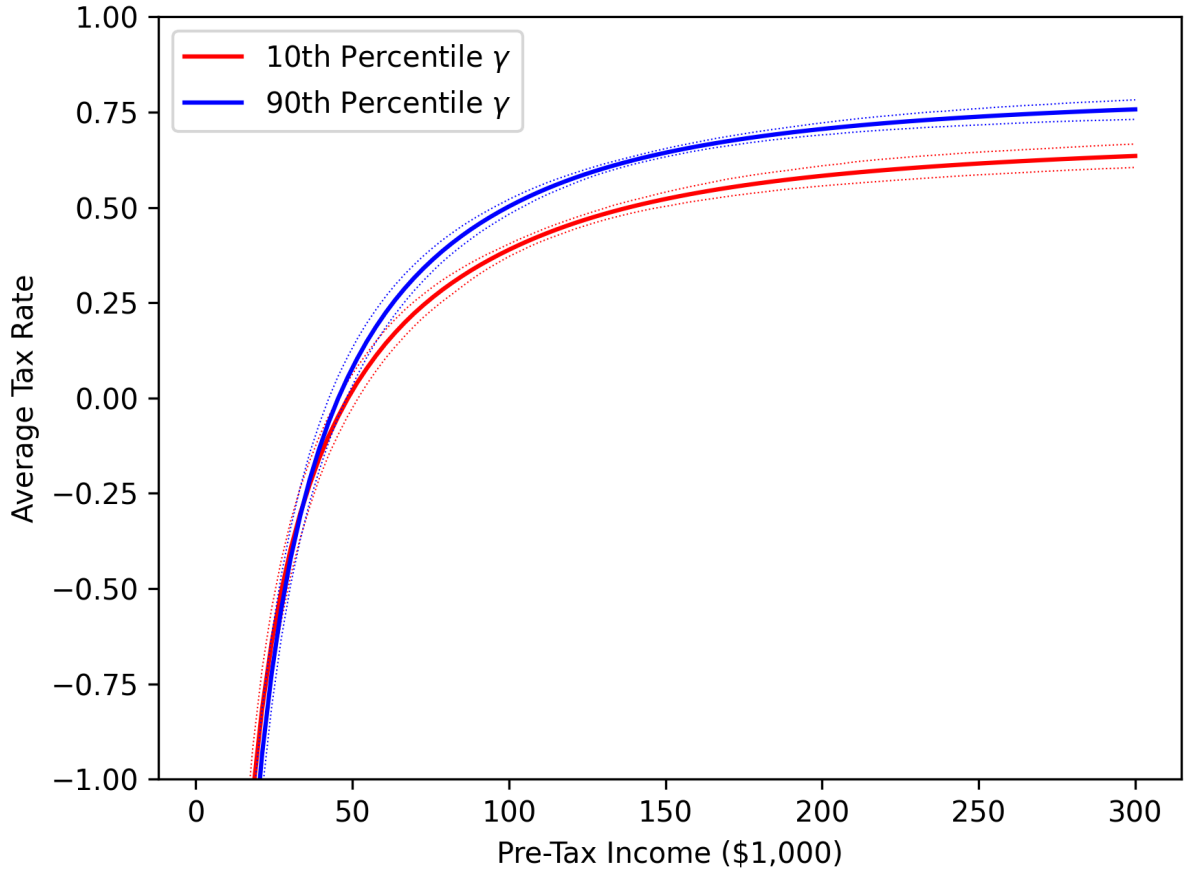
The second result is that optimal average tax rates implied by our estimates reveal substantial progressivity and sizable transfers at the bottom of the income distribution. For the least redistributive individuals (10th percentile γ, red line), the poorest households receive net transfers of around $27,000 per year. As income rises, these transfers decline, crossing zero (where taxes paid equal transfers received) at approximately $45,000 in annual income. Beyond this threshold, households face positive net taxation, with average tax rates rising to around 39% at $100,000 and stabilizing near 63% for incomes above $300,000.

For the most redistributive individuals (90th percentile γ, blue line), even more generous transfers and higher tax rates are implied throughout. For the poorest households, transfers exceed $35,000, while the break-even point where net taxes turn positive occurs around $50,000, roughly $10,000 higher than for the less redistributive group. Average tax rates then climb steeply, reaching approximately 50% at $100,000 and approaching 76% for top earners above $300,000. These implied tax schedules are dramatically more progressive than current U.S. policy, where average federal tax rates are only around 25% for earners making $300,000/year. The difference between the red and blue lines illustrates that normative disagreement translates into differences of 10-15 percentage points in preferred average tax rates across the income spectrum, with somewhat larger gaps at the bottom due to differing generosity of transfers.

The tax schedules imply that households that do not have any income would receive $27,000 in annual transfers based on the red curve, and $35,000 based on the blue curve. While this constitutes a sizable difference, both tax schedules would be substantially more progressive than the current US tax schedule. In Appendix G, we show that the current income tax schedule in the US would be close to optimal if society assigned *equal* welfare weights to each income level (i.e., the utilitarian case). This is puzzling because, given our welfare weights, even Republicans should support much higher taxes than are currently in place.[12]

[12]In Appendix A, we further illustrate which tax schedule would come out of democratic preference aggregation. Specifically, we compare the policy which maximizes *average* welfare (averaging the welfare function across the distribution of γ) to the policy which maximizes the preferences of the median voter (the median γ's welfare). We find that the two policies are identical: maximizing the median voter's preferences gets us to the policy that would

Figure 5: The Optimal Income Tax Schedule



Notes: This Figure shows optimal average tax rates as a function of pre-tax income. In red, we illustrate the optimal tax schedule using the 10th percentile of the distribution of CARA parameters. These are the parameters preferred by people with the least redistributive preferences. In blue we show the 90th percentile of CARA parameters, reflecting the most redistributive preferences. Negative values are transfers.

We next explore explanations for why our estimated tax rates are so much higher than the policies people vote for.

6 Why are Optimal Tax Rates so Much Higher than Current US Income Taxes?

We investigate this question through a sequence of experiments that test three candidate explanations. First, prior work has documented that beliefs about the efficiency costs of taxation and perceptions of social norms are associated with support for redistribution (Kuziemko et al., 2015; Stantcheva, 2021; Cruces et al., 2013). If these beliefs also shape the normative preferences re-

maximize average welfare.

covered from private transfer choices, then misperceptions could inflate or deflate the welfare weights we estimate and, consequently, the implied optimal tax rates. We test this through an information experiment that provides subjects with scientific evidence on behavioral responses to taxation and on the redistributive preferences of their fellow citizens.

Second, the Mirrleesian framework may be missing an important ingredient. Our transfer experiments ask subjects to allocate private windfall money between households, but actual redistribution is mediated through government taxation. If people are averse to government intervention beyond its efficiency costs, the standard framework would overstate support for progressive tax policy. To test this, we directly elicit subjects' preferred income tax rates and compare them to the rates implied by their transfer choices. We then introduce transfer choices tied to real donations to the U.S. Treasury to isolate the role of government aversion, and decompose its sources across the political spectrum.

6.1 The Role of Beliefs about Efficiency and Equity

A sample of 2,456 subjects is randomly assigned to one of three information treatments about the ETI, or to a control group. Treatment 1 informs subjects about the average ETI from our meta-analysis. Treatment 2 reports an unusually high ETI of 1.6, drawn from the highest single estimate by Saez (2010). The objective of this treatment is to induce a high belief about the ETI without deceiving participants.[13] Treatment 3 communicates the ETI for poorer and richer households, with poorer households responding more to taxation than richer ones.

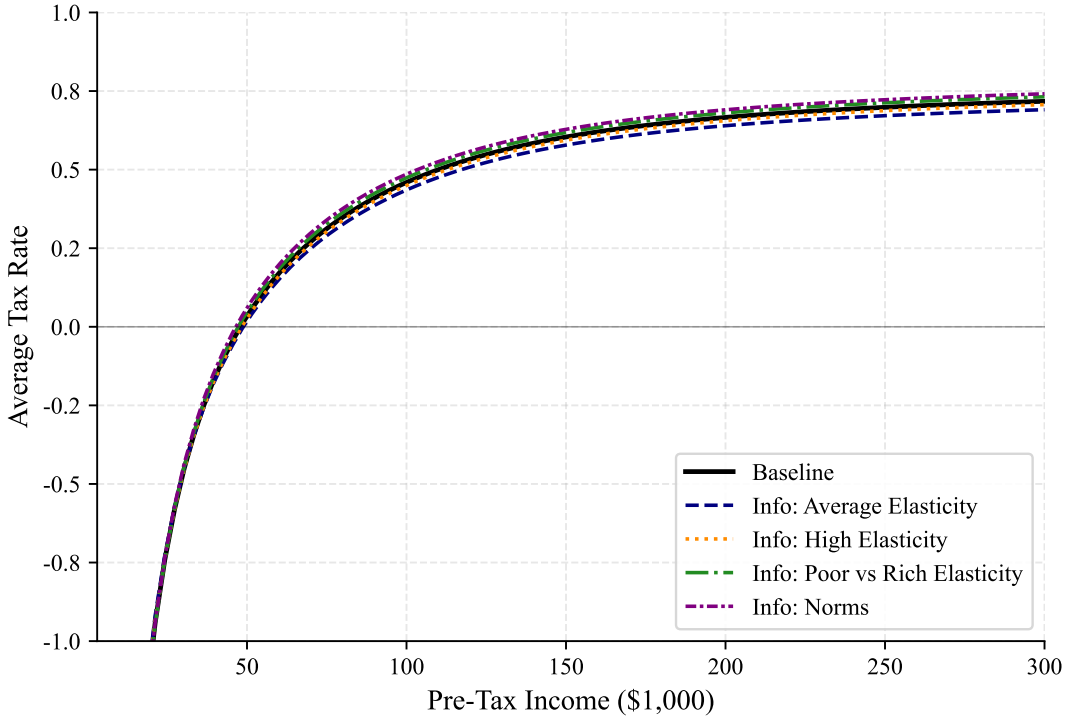
Orthogonally to these treatments, half of all subjects also receive information about the social welfare weights revealed in our first experiment. For instance, we inform subjects that, based on our estimates, Americans value a dollar transferred to a family earning $20,000 approximately 8 times more than a dollar transferred to a family earning $80,000.

Subjects then complete the same 10 pairwise transfer choices as in Wave 1. Screenshots of all treatments can be found in Figure 20 in the Appendix.

Results We find that information about the ETI reduces redistributive preferences, while information about others' redistributive norms increases them. We report treatment effects on optimal tax rates directly in Figure 6, delegating the full set of reduced-form results on choices and belief distributions to Appendix B. The norms treatment produces the largest shift, increasing implied average tax rates by approximately 2–3 percentage points across the income distribution relative to the control group. The Appendix shows that this occurs because subjects systematically

[13] The approach of selectively reporting different values to induce different belief distributions is widely used in experiments and considered a deception-free way of calibrating beliefs (see, for instance, Bottan and Perez-Truglia, 2025).

Figure 6: Optimal Tax Policy Implied under ETI and Norms Treatments



Notes: This figure plots optimal average tax rates as a function of pre-tax income, computed separately for each Wave 2 treatment group. For each treatment, we re-estimate the distribution of welfare weights and solve for the optimal tax schedule using the elasticity function from our meta-analysis. The black line corresponds to the baseline (control) group. Colored lines correspond to the three ETI information treatments and the social norms treatment. Negative values indicate net transfers.

underestimate others' generosity, and correcting this misperception raises their own redistributive behavior. The ETI treatments shift optimal rates downward by smaller magnitudes, and we illustrate that this reflects increased salience of the efficiency costs of transfers. Our findings are consistent with earlier *qualitative* evidence (Stantcheva, 2021; Cruces et al., 2013), but our methodology permits us to *quantify* how much these belief corrections actually move the implied optimal income tax schedule. Figure 6 highlights that, across all treatment arms, implied tax schedules remain far above current U.S. tax policy. Correcting beliefs about behavioral responses or social norms thus is important but cannot close the gap between transfer-implied and stated tax preferences.

6.2 Stated Tax Preferences and The Role of Government

We run two additional experimental waves with over 5,000 subjects. The first experiment (Wave 3), randomizes subjects into one of two groups. Subjects in the first group make the same transfer choices as in wave 1 and 2 but, we also ask them directly about their preferences for income tax

rates in a separate question. This question is shown either before or after the transfer choices, in random order. Specifically, for incomes in $\{\$20K, \$40K, \$60K, \$100K, \$500K, \$1M\}$, we ask subjects what, in their ideal tax policy, would be the *average* tax rate of individuals at that income. We elicit average tax rates because they are cognitively easier to process than marginal tax rates. The inclusion of this question allows us to understand if our transfer-choice-based preferences for redistribution systematically diverge from the tax schedule subjects can directly select in a hypothetical scenario.

One reason why stated and revealed preferences may differ, is that people dislike giving money to the government and our transfer choices do not incorporate any notion of taxation: subjects choose to allocate windfall gains to one of two families, but there is no income tax raised by the government.[14] This motivates our second experimental group. In this group, we design transfer choices that involve giving money to the government through a new incentive-compatible mechanism. We modify our 10 transfer choices by telling subjects that there is a baseline transfer to each family, and a tax on the transfer which goes to the government. Let B_{iqj} denote the baseline transfer in the experiment, and let r_{iqj} be the tax rate. Then, the recipient household j in question q of decision-maker i gets an after-tax transfer of $T_{iqj} = (1 - r_{iqj})B_{iqj}$, while the remaining amount $\tau_{iqj} = r_{iqj}B_{iqj}$ goes to the government. We randomize $B_{iqj} \in \{400, 800, \ldots, 4000\}$ and $r \in \{0\%, 25\%, 50\%, 75\%\}$ across choices and, as in previous waves, also the income of each family.

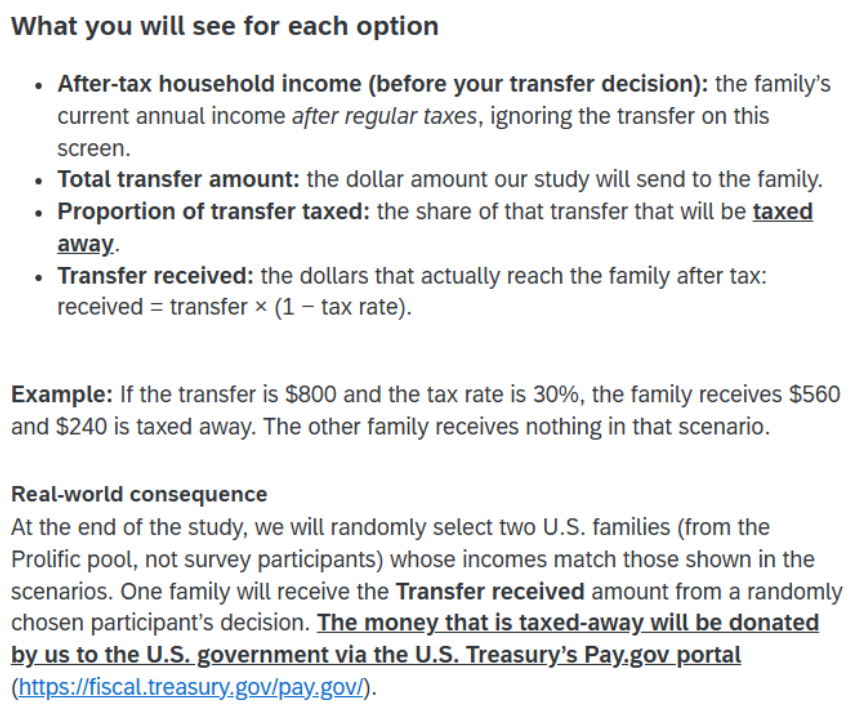
Incentive Compatibility Importantly, choices are fully incentivized: the transfer amount of a randomly selected choice goes to the selected family and the corresponding tax revenue is *donated to the US government* via the U.S. Treasury's www.Pay.gov portal. This portal allows anyone to give money to the Treasury. We tell subjects that, if their choice is randomly selected, we will donate the collected tax revenue to the US government. Figure 7 provides a screenshot of the relevant instructions.

Results We start by comparing stated versus revealed preferences of the group that made our standard transfer choices without taxes to the Treasury. In Figure 8, dotted lines are stated preferences, and solid lines are revealed preferences from transfer choices; the black line is the current US federal tax code. Red and blue lines represent the subsamples that self-identify as Republican and Democrats; moderates are excluded from the plot for simplicity.[15] Revealed preference

[14] Another reason why stated and revealed preferences diverge could be because hypothetical choices are often more pro-social than true behavior, e.g. due to social desirability bias (Rodemeier, 2025; Bursztyn et al., 2025). However, this would mean that hypothetical tax rates would be even more progressive than our revealed preference estimates, which is not what we find.

[15] Stated tax rates by moderates lie, as expected, between the blue and red dotted lines.

Figure 7: Instructions Explaining Real Transfers to the U.S. Treasury



What you will see for each option

- **After-tax household income (before your transfer decision):** the family's current annual income *after regular taxes*, ignoring the transfer on this screen.
- **Total transfer amount:** the dollar amount our study will send to the family.
- **Proportion of transfer taxed:** the share of that transfer that will be **taxed away**.
- **Transfer received:** the dollars that actually reach the family after tax: received = transfer × (1 − tax rate).

Example: If the transfer is $800 and the tax rate is 30%, the family receives $560 and $240 is taxed away. The other family receives nothing in that scenario.

Real-world consequence

At the end of the study, we will randomly select two U.S. families (from the Prolific pool, not survey participants) whose incomes match those shown in the scenarios. One family will receive the **Transfer received** amount from a randomly chosen participant's decision. **The money that is taxed-away will be donated by us to the U.S. government via the U.S. Treasury's Pay.gov portal** (https://fiscal.treasury.gov/pay.gov/).

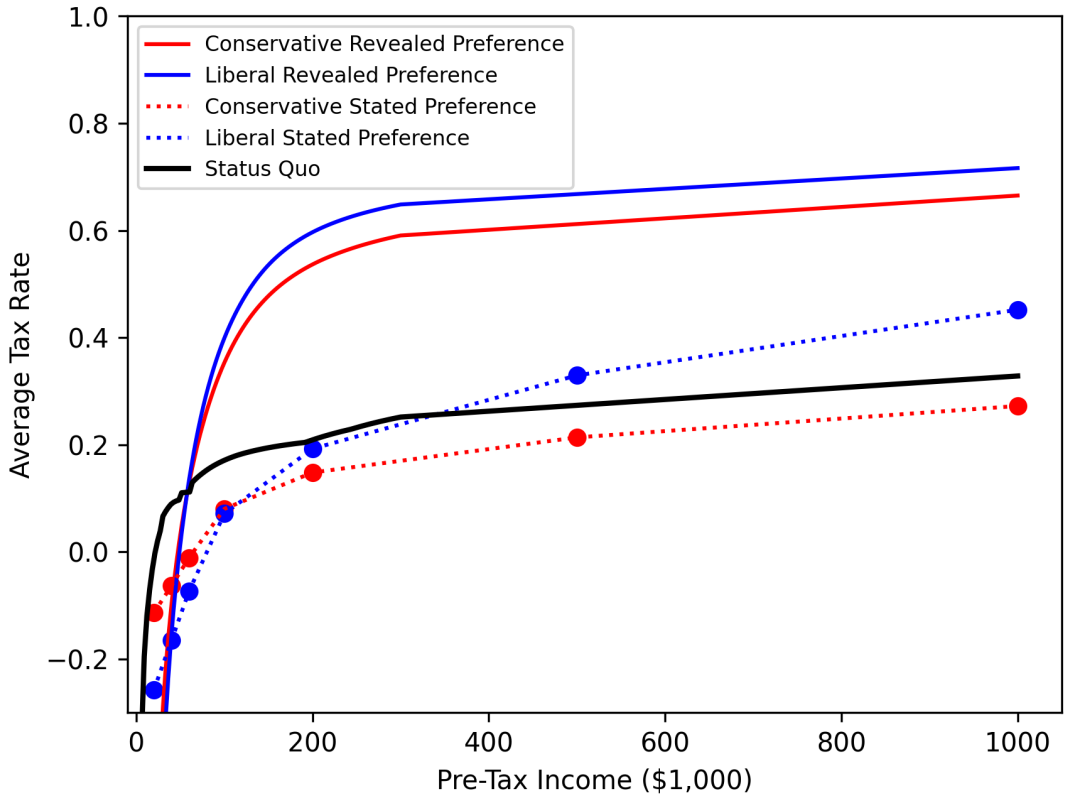
Notes: This Figure illustrates the relevant excerpt of the experimental instructions for subjects that faced transfer choices with taxes to the US Treasury.

estimates that are implicit in the transfer choices of subjects are again close to our previous results, and, dramatically more progressive than stated preferences. For incomes of $100k annually, implied *average* tax rates are around 35%, around 15% above the current US tax code. By contrast, both political camps state that they want lower taxes (of around 10%) for this income group in their stated preferences. Generally, we find that, stated preferences across the political spectrum imply lower taxes on families with incomes below $300k than the current US tax code implements. Above that level, Republicans prefer lower tax rates for richer households than are currently in place, while Democrats prefer higher taxes. For instance, at an income of $500k annually, the US tax code implements an average tax rate of around 29%, but Democrats say they prefer a rate of 33% and conservatives a rate of 21%. Yet, revealed preferences would imply much higher rates of above 60% for both political camps.

These results illustrate that when subjects directly choose income taxes in hypothetical elicitations, the resulting preferences are much closer to the current US tax code, and far below the rates we obtain from incentivized transfer choices.

The Role of Narrow Self-Interest Another explanation for the gap between distributive preferences and support for tax policy is self-interest: when voting for tax policy, individuals may

Figure 8: Revealed vs Stated Tax Preferences



Notes: This Figure illustrates revealed preference estimates of optimal average income taxes (colored solid lines), stated preferences (dashed lines), and the current US income tax schedule (black line). Blue and red indicate people's self-reported political affiliation to either Liberals or Conservatives.

shade their preferred rates downward to protect their own after-tax income, either today or in the future. This channel would not operate in our transfer experiments, where subjects allocate windfall money to others with no consequences for their own tax burden. If narrow self-interest drives the puzzle, it should do so most strongly for individuals who stand to lose from higher taxes. High-income respondents face a direct cost from progressive taxation, and younger respondents may anticipate upward income mobility that would expose them to higher rates in the future. Conversely, individuals who are both currently low-income and near retirement have little material stake in the tax rates applied to high earners: they are unlikely to reach those income levels themselves, and they stand to benefit from the transfers that progressive taxation finances. We exploit this logic to test the self-interest hypothesis. We restrict the sample to respondents who are older than 50 and report annual income below $50,000. For this subgroup, the self-interest motive for understating preferred tax rates on the rich is minimal. If anything, narrow self-interest predicts that these individuals should favor *more* progressive taxation than the full sample, since they would be net beneficiaries of redistribution. In Figure 16 in the Appendix,

we show that this subgroup prefers even *lower* taxes according to their stated preferences, while their tax rates based on transfer choices are identical to the ones of the full sample. This rules out narrow self-interest as the primary driver of the divergence between stated and revealed preference.

The Role of Government We turn to our second group whose incentivized transfer choices were tied to giving money to the government. As before, let C_{iqj} be the post-tax income of the recipient household, and let Y_{iq} be an indicator for giving the transfer to household 1. We run a regression of the form

$$Y_{iq} = \underbrace{\beta_1[\tau_{iq1} - \tau_{iq2}]}_{\text{effect of tax difference}} + \underbrace{\beta_2 \frac{C_{iq1} + C_{iq2}}{2}}_{\text{average effect of income}} + \underbrace{\beta_3[C_{iq1} - C_{iq2}]}_{\text{effect of income difference}} + \underbrace{\beta_4 \frac{T_{iq1} + T_{iq2}}{2}}_{\text{average effect of transfer}} + \underbrace{\beta_5[T_{iq1} - T_{iq2}]}_{\text{effect of transfer difference}} + \varepsilon_{iq}. \quad (7)$$

In this regression, β_1 measures the isolated effect of higher tax amounts paid to the government on subjects' choice, while holding constant income and net-of-tax transfer size across receiving families. Coefficients β_2 and β_4 represent, respectively, the effect of the average income level and the average transfer level across the two families on the subject's choice. Coefficients β_3 and β_5 tell us how the *difference* in income and transfer size between the two families matter for the choice. A comparison of β_1 with β_5 quantifies how much more subjects value giving money to the receiving family than to the government.

We run the regression separately for conservatives, moderates, and liberals. The results from this regression are given in Table 1. For conservatives, the first row shows that, holding all else constant, increasing the amount received by the government by \$1,000 *decreases* the probability of making that transfer choice by 4.7 pp. This means that conservatives actively dislike when the Treasury receives money. It is interesting to compare this to the last row, which shows that increasing the amount given to a household *increases* the probability of making the transfer choice by 7.2 pp. To put this number into perspective, this implies that a \$1,500 transfer taxed at 33% (so the household receives \$1,000 and the government receives \$500) is, on average, as attractive as making a tax-free \$700 transfer to the household. In other words, *conservatives in our sample are willing to reduce net generosity to recipients in order to avoid giving money to the government*. The results illustrate a strong aversion to providing funds to the government generally, and are even more striking in the light of the current US federal government being Republican.

The point estimate for taxes among liberals is small and statistically insignificant. Thus, while liberals in our sample do not seem to actively avoid giving the government money, they derive no incremental utility from doing so. It is still possible that these subjects positively value government revenue *generally*, but the presence of the Trump administration may induce

concerns about the way public funds are used, drawing this coefficient toward zero. We turn to this issue again later. Moderates appear about halfway between conservatives and liberals, and while not statistically significant, appear to directionally also dislike transferring money to the government.

These results imply that observed tax policy deviates from distributive preferences in part because people are averse to government mediation.

Table 1: Sensitivities to Taxes, Income, and Transfer Size

	Subgroup:			
	All	Conservative	Moderate	Liberal
	(1)	(2)	(3)	(4)
$\tau_{iq1} - \tau_{iq2}$	-0.014467^{**}	-0.047192^{***}	-0.011121	0.000401
	(0.005701)	(0.011578)	(0.010731)	(0.008202)
$\frac{C_{iq1}+C_{iq2}}{2}$	0.001137^{***}	0.000951^{***}	0.001209^{***}	0.001184^{***}
	(0.000042)	(0.000085)	(0.000078)	(0.000061)
$[C_{iq1} - C_{iq2}]$	-0.000640^{***}	-0.000568^{***}	-0.000407^{***}	-0.000805^{***}
	(0.000031)	(0.000062)	(0.000061)	(0.000044)
$\frac{T_{iq1}+T_{iq2}}{2}$	0.117584^{***}	0.138473^{***}	0.111931^{***}	0.111746^{***}
	(0.007319)	(0.014865)	(0.013512)	(0.010656)
$[T_{iq1} - T_{iq2}]$	0.057541^{***}	0.072480^{***}	0.060652^{***}	0.048438^{***}
	(0.004629)	(0.009414)	(0.008790)	(0.006619)
N	6,040	1,440	1,760	2,840

Notes: This Table reports regression results from specification 7 for the whole sample as well as separately for self-reported conservatives, moderates, and liberals. Statistical significance at *$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$*

6.2.1 Optimal Taxes with Government Aversion

We quantitatively assess the importance of our estimated measures of government aversion in shaping optimal tax policies. For a given tax policy $T(\cdot)$, define $\mathcal{T}_+(T) = \{z : T(z) > 0\}$ as the set of incomes where individuals face a net tax burden. We parameterize the idea of government aversion by modifying the planner's problem so that the social planner trades off the utility effects of tax reform against an explicit penalty on the overall level of tax revenue collected. Formally,

we compute optimal tax policies of the form

$$\max_{T(\cdot)} \int_0^\infty \Gamma\big(u\big(z - T(z)\big) - \theta v(z)\big)\, h(z)\, dz - \lambda \int_{\mathcal{T}_+(T)} T(z) h(z)\, \mathrm{d}z,$$
$$\text{s.t.} \quad \int_0^\infty T(z)\, h(z)\, dz \geq E.$$

Here, $\lambda \geq 0$ reflects the degree of government aversion, and the case where $\lambda = 0$ corresponds to the standard Mirrlees problem. This represents an additional deadweight loss from government involvement itself, not included in the standard Mirrleesian framework.

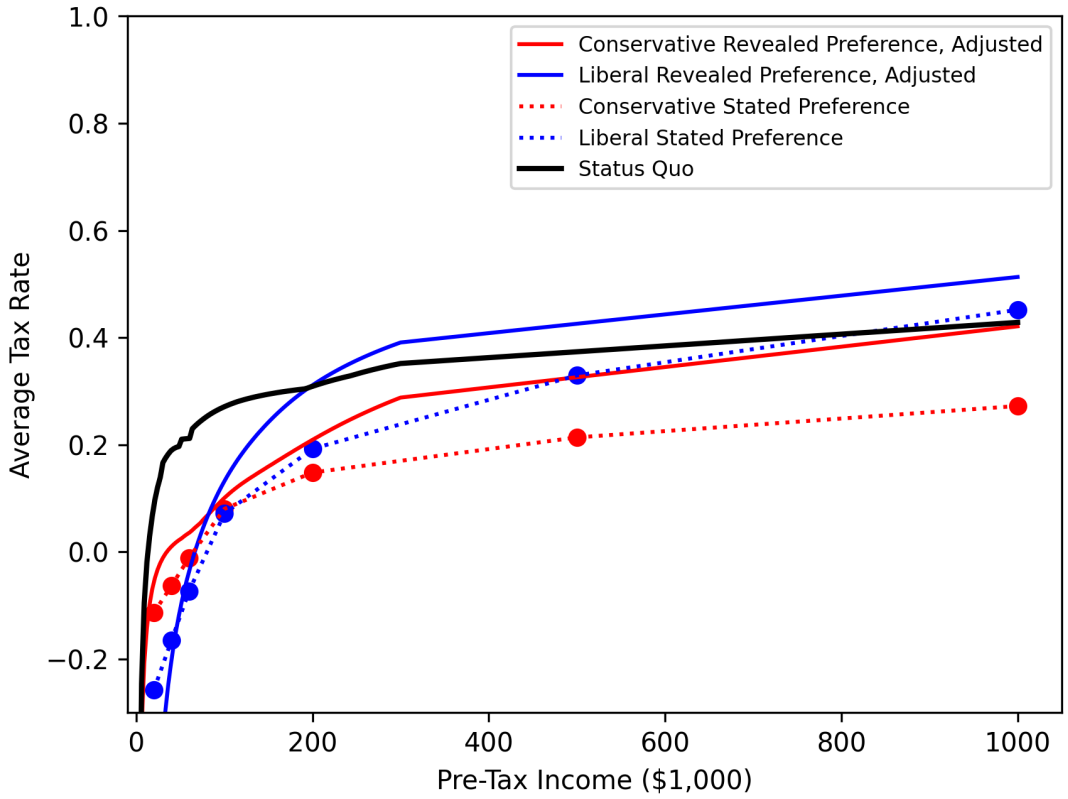
To identify λ, we use the coefficients from Table 1. Note that if preferences over tax policy conform to the standard Mirrlees model, a single dollar of tax raised for the government should be valued equivalently to a transfer to a randomly chosen member of society.[16] Thus, the coefficient on $\tau_{iq1} - \tau_{iq2}$ should be positive and equivalent in magnitude to the coefficient on $T_{iq1} - T_{iq2}$. Since liberals do not actively value tax revenue (at least not to the current Republican administration), their level of government aversion must be just enough to offset the positive benefits of the transfer. Thus, $\lambda_{lib} \approx \mathbb{E}[G_{lib}(z)]$ for liberals, where $G_{lib}(z)$ is the welfare weight placed on an individual with income z, given the redistributive preferences of the average liberal. Conservatives value a \$1 increase in transfers at 0.072 but disvalue a \$1 increase in taxes at –0.047. Thus, $\lambda = (1 + 0.047/0.072) \times \mathbf{E}[G_{cons}(z)] \approx 1.65 \times \mathbf{E}[G_{cons}(z)]$, where analogously, $G_{cons}(z)$ is the welfare weight on an individual with income z, given the redistributive preferences of the average conservative.

Using the estimated values of λ, we re-compute optimal tax schedules and report the results in Figure 9. Accounting for government aversion substantially reduces the implied level of progressivity. For households earning \$500,000, implied optimal rates fall from previously 60% to now 36% and 46% for conservatives and liberals, respectively. For households earning \$100,000, implied rates drop from 35–40% to 13% for liberals and 10% for conservatives. Implied optimal tax schedule that account for aversion government also have implications for transfer policies: the new tax schedule now provides transfers to all households earning \$60K/year or less.

Our revised revealed preference rates now align much more closely with people's stated preferences for taxation. On average, about 84% of the gap between transfer-based revealed prefer-

[16] In the standard Mirrlees setup, any windfall increase in government revenue relaxes the budget constraint and is optimally recycled to citizens via a uniform lump-sum rebate. By the planner's envelope condition, the social value of this rebate equals the average welfare weight across the income distribution, which is normalized to one. Hence a marginal dollar of tax revenue is socially equivalent to a dollar transferred to a randomly drawn member of society. This equivalence is robust to explicitly modeling public goods provision. If the government jointly optimizes redistribution and a public good G with social benefit $B(G)$, the optimality condition equates the marginal value of redistribution spending to $B'(G)$; a windfall dollar relaxes both margins equally, so its social value remains the average welfare weight. Under the Mirrlees null, therefore, the coefficient on $\tau_{iq1} - \tau_{iq2}$ in our regression should equal the coefficient on $T_{iq1} - T_{iq2}$. Any negative deviation from this equality identifies $\lambda > 0$.

Figure 9: Optimal Taxes with Government Aversion



Notes: This Figure illustrates revealed preference estimates of optimal average income taxes (colored solid lines), stated preferences (dashed lines), and the current US income tax schedule black line). Blue and red indicate people's self-reported political affiliation to either Democrats or Republicans.

ences and the reported ideal income tax schedule is explained just by accounting for government aversion.[17] The remaining 16% gap may reflect additional factors we can only speculate about, such as misunderstanding of the tax system or political economy constraints not captured in our framework. Future work can make important contributions here by leveraging our methodology.

The documented partisan difference in government aversion helps explain why tax policy remains contentious even when both political camps strongly favor the welfare of the poor over the rich.

[17] We measure the gap between revealed and stated tax preferences at a given income level as the difference between the tax burden implied by solving the optimal tax problem and the tax burden implied by the average tax rates stated by survey respondents. For each political affiliation, and for each income level where survey respondents were asked to state a preferred tax rate, we computed the magnitude of the gap, both under the current model (with government aversion) as well as under the baseline model (without government aversion). We then computed by what percent the average gap under the current model is smaller than the average gap under the baseline model.

6.2.2 The Sources of Government Aversion

Our findings suggest that the large gap between choice-implied optimal tax rates and actual policy support may stem from a deep skepticism about government mediation. There are two competing explanations for this aversion: First, respondents may distrust that the government efficiently allocates tax revenue. Under this interpretation, people care about helping lower-income households but believe that government programs waste resources, create perverse incentives, or fail to reach intended beneficiaries. The aversion to government taxation would then reflect *instrumental concerns* about implementation rather than fundamental opposition to redistribution. Second, respondents may have *principled objections* to the coercive nature of taxation itself, regardless of efficiency considerations. This view emphasizes property rights and the moral significance of earned income, suggesting that taking money from those who earned it violates important ethical principles even when used for beneficial purposes.

Understanding which source drives behavior is important because the policy implications differ substantially. If government aversion stems from efficiency concerns, institutional reforms might unlock latent support for redistribution. If instead it reflects deeper philosophical commitments about property and coercion, such reforms would likely prove ineffective. We run a final follow-up experiment (Wave 4) that speaks to these mechanisms.

A final concern with our prior experiment in Wave 3 is that government aversion might not reflect attitudes toward government per se, but rather respondents' beliefs about the alternative use of research funds in the experiment. Specifically, when participants choose smaller transfers that generate less tax revenue, they implicitly preserve funds to our research budget. If participants believe these retained funds could be productively used for future research or other valuable purposes, their apparent "government aversion" might actually reflect a preference for research funding over government revenue, rather than general skepticism toward taxation. To address this concern, we include treatments that manipulate attention and beliefs about whether retained research funds have any productive use.[18]

Experimental Design In January 2026, we ran our final experiment, in which subjects again choose between two transfer options that differ in recipient income, net-of-tax transfer amounts, and the amount of money transferred to the government. Unlike Wave 3, all choices in Wave

[18] We note that this concern applies far beyond our paper. For instance, a longstanding, and still active, literature in economics measures prosocial behavior in experiments, where subjects can allocate research funds to other experimental participants (e.g., Güth et al., 1982; Kahneman et al., 1986). A famous example is the dictator game where scholars ignore the (arguably unlikely) possibility that the subject playing the dictator has preferences to save money to the researcher and therefore proposes overly unfair offers to the recipient in order to provoke a rejection of the offer. This is just one of many examples in which preferences to save the researcher's budget would bias results. Our additional treatments reassuringly show that subjects' choices do not exhibit such preferences.

4 are hypothetical, which allows us to implement strong framing assumptions that would be impossible to enforce under full incentive compatibility.

In total, 3,103 participants are randomly assigned to one of four conditions. The control condition replicates the Wave 3 framing without additional information. In the remaining conditions, we manipulate how government revenue and the alternative use of foregone transfers is framed:

1. *Efficient Government*: Subjects are instructed to assume that the government uses tax revenue with perfect efficiency, such that every dollar transferred to the government generates one dollar of valuable public services that directly benefit citizens, including the recipient households.

2. *Research Budget Reminder*: Subjects receive a salient reminder informing them that choosing smaller transfers saves money for the researchers, which would otherwise be used for future research projects.

3. *No Alternative Use for Research Budget*: Subjects are asked to assume that any money saved by choosing smaller transfers has no productive use and simply disappears.

Screenshots of the instructions for these treatments are shown in Figure 10. The *Efficient Government* treatment is the key intervention, as it explicitly removes efficiency-based objections to taxation. The *Research Budget Reminder* tests whether attention to alternative uses of retained funds changes transfer choices. If it does, subjects were (at least partially) inattentive to this possibility in prior waves. The *No Alternative Use for Research Budget* treatment eliminates this channel entirely. If this treatment affects behavior, it would imply that subjects were previously mindful of alternative uses.

We estimate the same linear probability model as in Wave 3 (equation 7), relating transfer choices to differences in tax payments to the government, recipient incomes, and transfer amounts.

Results We report the full regression results in Table 2. In Figure 11 we illustrate the key finding. Panel A shows the coefficient on tax revenue, $\tau_{iq1} - \tau_{iq2}$, as well as the coefficient on transfer size to the family, $T_{iq1} - T_{iq2}$, for the control group and the Efficiency treatment. In the control condition, we replicate the core finding from Wave 3: the coefficient on taxes paid to the government is negative and statistically significant, indicating government aversion.

The Efficient Government treatment reverses the sign of the tax coefficient, making the marginal valuation of government revenue positive and statistically significant from zero. This

implies that government aversion is largely explained by efficiency concerns and a wasteful government. The coefficient on tax revenue is 0.012 and around 23% of the coefficient on transfer size (0.052). Under an efficient government, the average subject values a dollar going directly to the average family 4.3x as much as a dollar going to the Treasury.

Turning to Table 2, we find that reminding people that spending less in our experiment reduces subject's transfer elasticity, making them more mindful of other productive uses. Asking subjects to assume that saved research funds simply disappear does not change behavior and produces the same results as in the control group. Reassuringly, these results imply that people in our control condition (and in prior waves) likely did not pay attention to the possibility that any saved funds to the researcher may have productive use in the future.

Finally, comparing the control group in Wave 4 (Table 2) with the sample from Wave 3 (column 1 of Table 1) shows that they barely differ in their behavior. This means that the hypothetical nature of the transfer choices in Wave 4 did not change choices relative to the consequential behavior in Wave 3.[19]

In Panel B, we report coefficients when we run the regressions separately for Liberals and Conservatives. The effect of our efficiency framing treatment is almost entirely driven by subjects identifying as Liberals. The effect on Conservatives is small and statistically insignificant. These results suggest that government aversion has different foundations across the political spectrum. For liberals, aversion to taxation appears largely instrumental and driven by concerns about inefficiency or misuse of funds.[20] Once these concerns are removed, support for government-mediated redistribution aligns closely with revealed redistributive preferences. For conservatives, government aversion appears to reflect a more fundamental objection that is not eliminated by assumptions of perfect efficiency. This is consistent with principled views that consider taxation a form of coercion and theft, and that emphasize property rights.

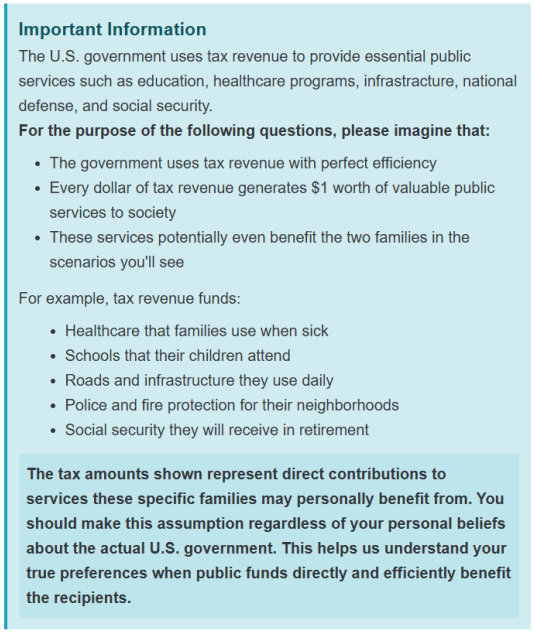
Since even conservatives were generous to the poor in our private donation tasks in Wave 1-3, we elicited in Wave 4 if these subjects indeed view charity as the better solution to tackle inequality. We asked subjects "If you wanted to help low-income families, which method would you prefer?" and we plot the answers in Panel A of Figure 12. The contrast across political ideology is stark. Among liberals, 39% view government programs as the appropriate way to help the poor, while 15% prefer direct donations to individuals and another 15% prefer donations via nonprofit organizations. Among conservatives, only 13% choose government programs, while 37% and 28% prefer direct donations or nonprofit organizations, respectively.

The share of people who consider all methods equally good is 30% among liberals and 20%

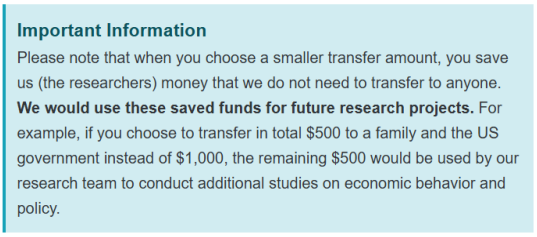
[19]This supports a recent trend in economics to rely more heavily on stated choices from surveys rather than exclusively on incentivized choices.

[20]This can also explain why in Wave 3 liberals' positive valuation of government revenue was drawn to zero due to worries that the Trump administration uses these funds for policies they oppose.

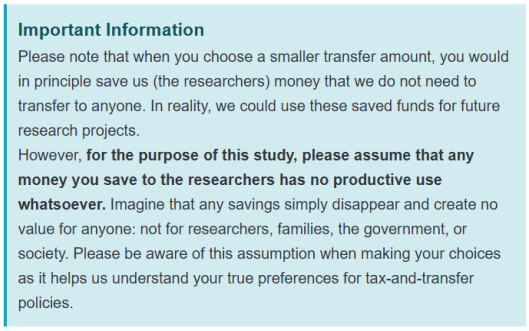
Figure 10: Treatments in Wave 4 Experiment on Government Efficiency



(a) Efficient Government



(b) Highlight Research Budget



(c) No Alternative Use for Research Budget

Notes: This figure displays screenshots of the instructions used in the Wave 4 experiment. Panel A presents the Efficient Government treatment, which asks subjects to assume that all tax revenue is used with perfect efficiency. Panel B presents the Research Budget Reminder treatment, which makes salient that choosing smaller transfers preserves funds for the researchers. Panel C presents the No Alternative Use for Research Budget treatment, which stipulates that any funds not transferred have no productive use.

among republicans. This is a non-negligible share even among the right and points to important heterogeneity within the conservative camp. Finally, across the political spectrum people are willing to help the poor. Only 0.3% of liberals and 3% of republicans answer that they would not want to help at all. This aligns with our previous result that the key partisan difference is not the willingness to help the poor but the divergence in institutional preferences.

Panel B sheds further light on Conservatives' principled objection against government redistribution. Subjects were asked to what extent they agree with the statement "The government

Table 2: Transfer Sensitivities across Treatments from Wave 4

		Treatments		
	Control	Research Budget Reminder	No Alternative Use for Research Budget	Government Efficient
$\tau_{iq1} - \tau_{iq2}$	-0.029067^{***}	-0.039436^{***}	-0.036077^{***}	0.012075^{**}
	(0.005278)	(0.005349)	(0.005272)	(0.005454)
	(0.000039)	(0.000040)	(0.000039)	(0.000041)
$C_{iq1} - C_{iq2}$	-0.000576^{***}	-0.000630^{***}	-0.000556^{***}	-0.000418^{***}
	(0.000029)	(0.000030)	(0.000029)	(0.000031)
$\frac{T_{iq1}+T_{iq2}}{2}$	0.134925^{***}	0.117495^{***}	0.132444^{***}	0.131728^{***}
	(0.006985)	(0.007005)	(0.006867)	(0.007134)
$T_{iq1} - T_{iq2}$	0.060394^{***}	0.037740^{***}	0.072898^{***}	0.052381^{***}
	(0.004332)	(0.004418)	(0.004326)	(0.004478)
N	6,830	6,810	6,910	6,800

Notes: This table presents regression results for the Wave 4 experiment, using specification 7, and separating the sample into each experimental group.

has the right to take money from those who earn more to give to those who earn less" on a 7-point scale from strongly disagree to strongly agree. The distribution of responses confirms a fundamental philosophical divide. Among liberals, responses cluster strongly toward agreement, with nearly 30% strongly agreeing and another 26% somewhat agreeing that government has this right. The modal response for liberals is strong agreement. Among conservatives, the pattern reverses. Approximately 27% strongly disagree with the statement, and responses concentrate heavily on the disagreement side of the scale. The modal response for conservatives is strong disagreement. Only about 13% somewhat agree while less than 2% strongly agree. This stark contrast in views about government legitimacy helps explain why even a perfectly efficient government would fail to attract conservative support for progressive taxation.[21] The divergence reflects competing philosophical views about property rights and the proper scope of state authority, not merely disagreement about government competence.

[21]The results also imply that current initiatives to improve government efficiency, such as the Department of Government Efficiency (DOGE), would not result in more support for collecting tax revenue among conservative voters.

Finally, we introduce a useful metric to quantify the value of public funds across the political spectrum: the "Equivalent-Transfer Metric of Government Revenue," which we define as the voter's indifference point between giving $1 to the US Treasury or alternatively $Y to a family with income $X. This metric answers the question: how much money would need to go directly to a household at a given income level to make a voter indifferent to giving $1 to the government? We estimate this quantity in the aggregate by fitting a modified logistic regression. Specifically, we assume that the utility of survey-taker i in question q from making transfer T_{ijq} to family j with income Z_{ijq} with tax τ_{ijq} is

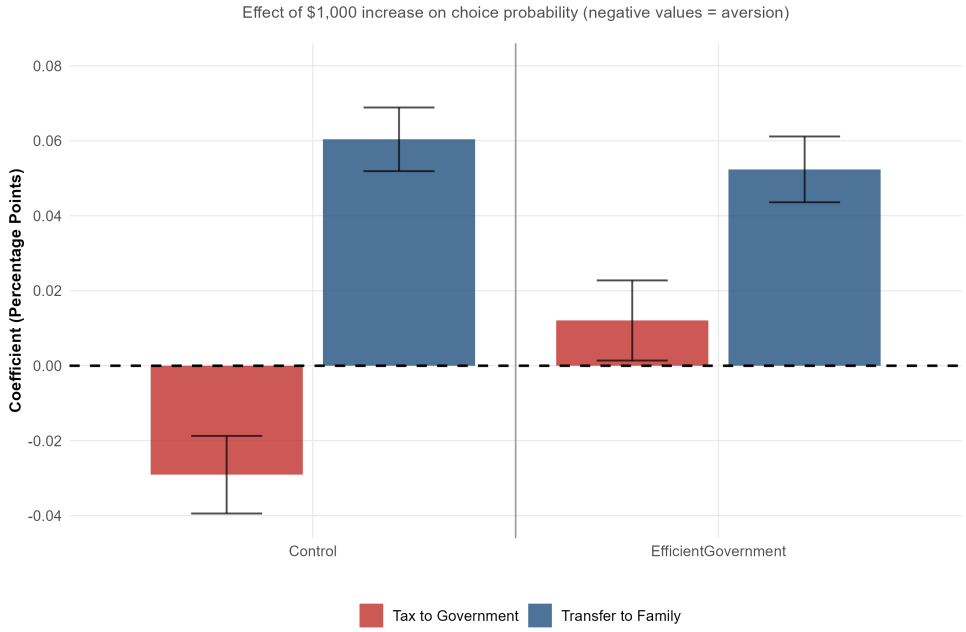
$$U_{ijq} = \lambda\tau_{ijq} + T_i \exp(-\gamma Z_{ijq}) + \varepsilon_{ijq}.$$

where $\varepsilon_{ijq} \sim T1EV(\sigma)$. This is similar to a standard logistic regression, except that utility is non-linear in parameters. We fit it using nonlinear least squares. Under this specification, $\frac{\lambda}{\exp(-\gamma Z)}$ represents the marginal value of collecting an additional dollar of tax revenue divided by the marginal value of giving that dollar to a family of income Z. When this ratio equals one, the subject is indifferent whether the dollar goes to the government or to the family.

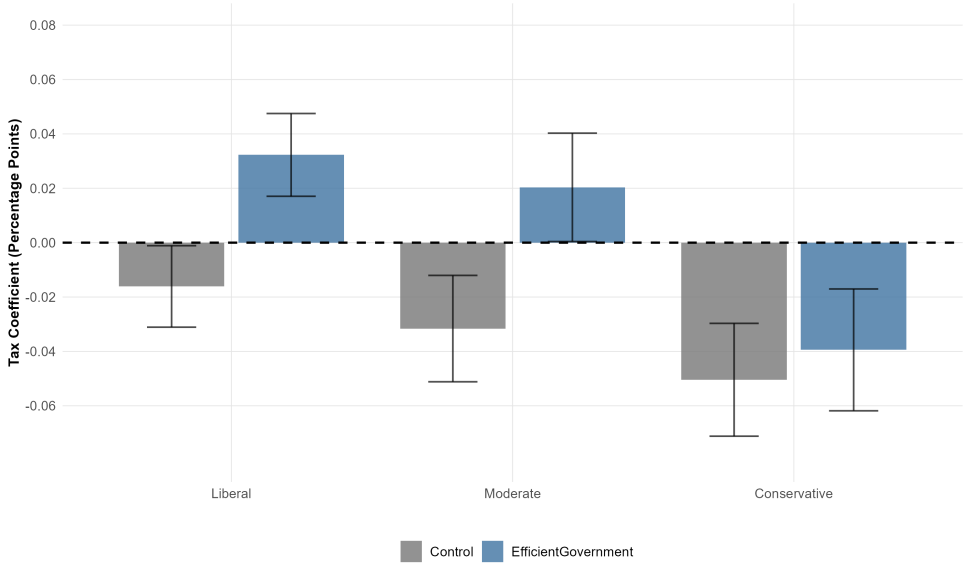
We plot the results for the case in which the government is perfectly efficient in Figure 13. For liberals (blue line), the metric reveals that government revenue has substantial but income-dependent value. At income levels below $250k, liberals are roughly indifferent between $1 going directly to the household versus to the government. As recipient income increases, the equivalent-transfer value rises exponentially. For a household earning $1 million, liberals would require approximately $86 to go directly to that household to feel indifferent to $1 going to the government. For conservatives (red line), the metric remains essentially flat near zero across the entire income distribution. Even for wealthy recipients, conservatives derive virtually no utility from tax revenue going to the government. This metric provides a portable, intuitive summary of government aversion that can be applied in future research on optimal taxation and political economy.

Overall, Wave 4 shows that government aversion is not a monolithic phenomenon. Its sources vary systematically across individuals, and this heterogeneity is central for understanding why strong redistributive preferences in private choice settings do not translate into broad political support for progressive taxation.

Figure 11: Valuation of Tax Revenue



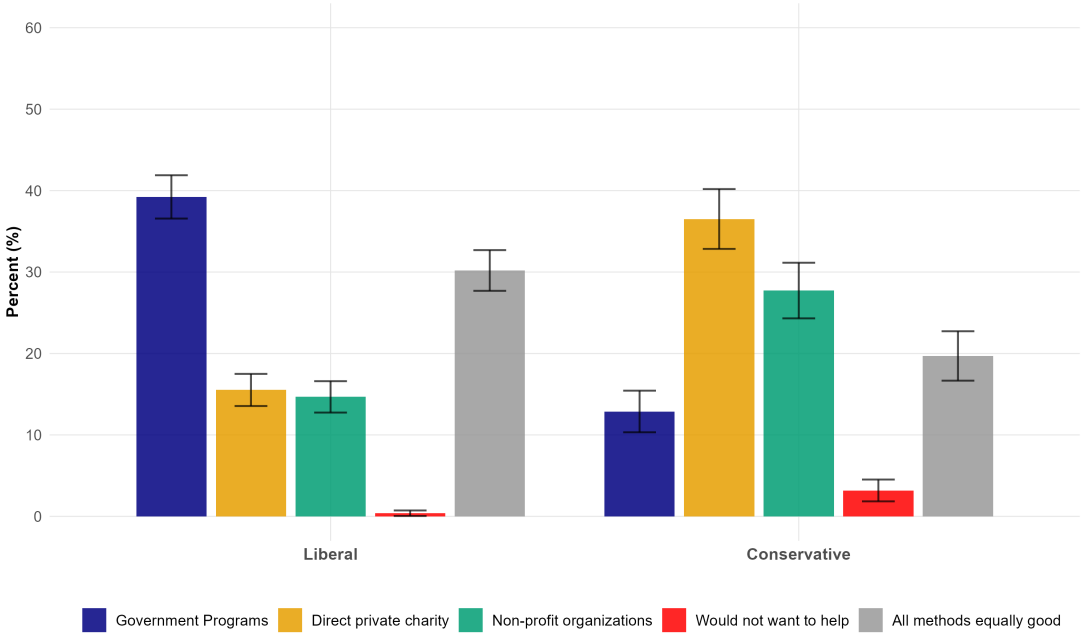
(a) By treatment



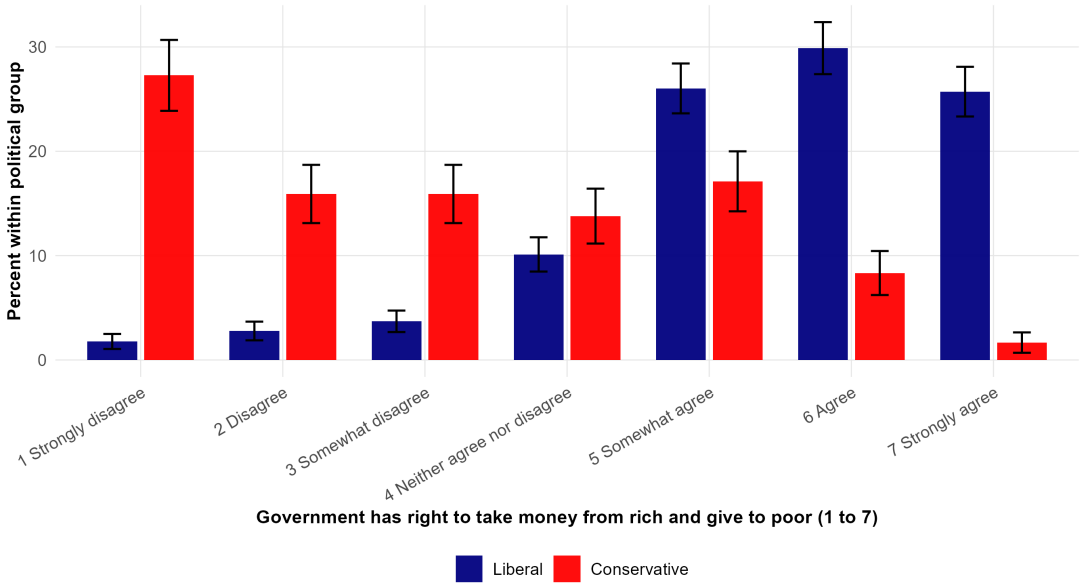
(b) Across political spectrum, by government efficiency treatment

Notes: This figure plots the marginal valuation of tax revenue and direct transfers from Wave 4. Panel A compares coefficients on tax revenue ($\tau_{iq1} - \tau_{iq2}$, blue bars) and transfer size ($T_{iq1} - T_{iq2}$, red bars) for the control group and the Efficient Government treatment. Panel B shows how the tax coefficient varies across the political spectrum under control (blue) and Efficient Government (red) conditions. Error bars represent 95% confidence intervals.

Figure 12: Attitudes toward Government Redistribution



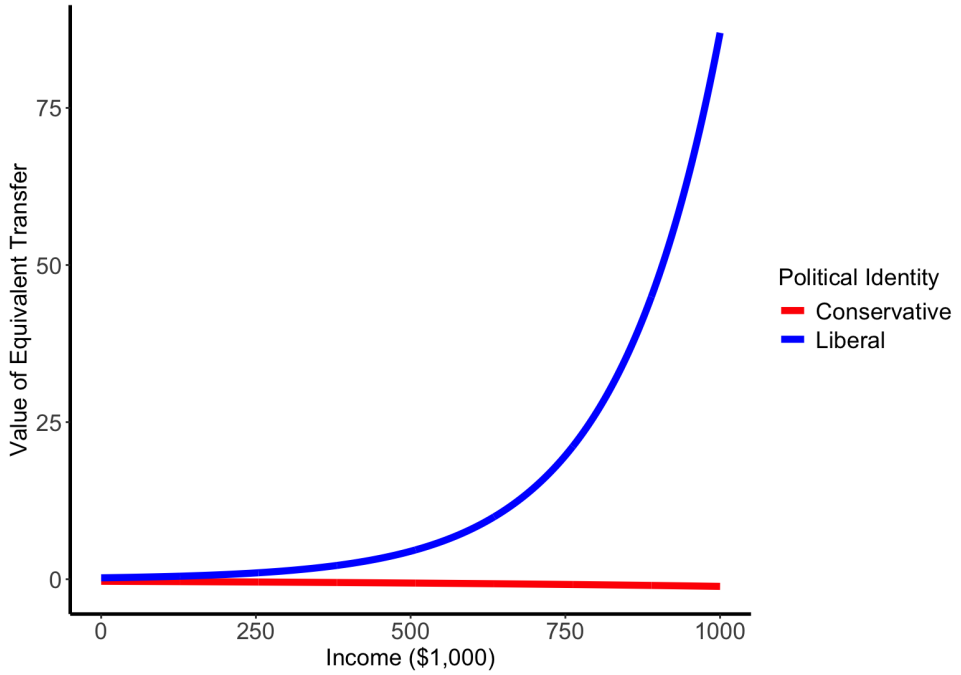
(a) Preferred Redistribution Method



(b) Principled Preference for Redistribution through the Government

Notes: Panel A displays respondents' stated preferences for how to help low-income families, separated by political ideology. Panel B plots responses to the question of whether the government has the right to take money from those who earn more to give to those who earn less, separated by political ideology. Error bars represent 95% confidence intervals.

Figure 13: The Equivalent-Transfer Metric of Government Funds



Notes: This figure plots the equivalent-transfer metric of government revenue, defined as the dollar amount that would need to go to a household at a given income level to make a voter indifferent to $1 going directly to the US Treasury. The blue line shows this metric for liberals, while the red line shows it for conservatives.

7 Conclusion

This paper studies a puzzle for redistributive policy. When individuals make incentivized choices that trade off the welfare of poor and rich households, they reveal strong and broadly shared concern for the poor. When these revealed preferences are mapped into the standard framework of optimal income taxation, they imply far more progressive tax policy than what the same individuals support politically. The gap between distributive preferences and political support is large, systematic, and present across the political spectrum.

Our main contribution is to quantify this gap and show that it is largely explained by aversion to government as an institution. Opposition to high taxes does not primarily reflect weak concern for low-income households. Instead, it reflects resistance to government-mediated redistribution. By experimentally varying whether redistribution is implemented privately or through actual payments to the U.S. Treasury, we show that government aversion accounts for over 80% of the divergence between distributive values and support for redistribution. Once institutional resistance is incorporated into the optimal tax problem, tax rates implied by transfer choices move substantially closer to the tax policies people endorse politically.

We further show that the sources of government aversion differ sharply across the political

spectrum. For liberals, aversion is largely instrumental. Concerns about inefficiency or misuse of public funds dampen support for taxation, but disappear when government is assumed to be perfectly efficient. For conservatives, aversion persists even under such assumptions and reflects principled objections tied to property rights, coercion, and the legitimacy of taxation itself. These differences help explain why disagreement over redistribution remains deep even when there is widespread willingness to help the poor.

Methodologically, the paper contributes a revealed preference approach to measuring social marginal welfare weights, a unified Bayesian elasticity function across the income distribution, and a portable metric that quantifies the value of public funds relative to private transfers. Together, these tools clarify the distinction between redistributive policy and society's distributive values. Substantively, the results suggest that much of the political conflict over taxation is not about whether to help low-income households, but about who should do so and through which institutions.

More broadly, our findings point to a limitation of standard normative frameworks that abstract from institutional legitimacy. Even when distributive preferences are identified and behavioral responses are well measured, they do not directly translate into politically feasible tax policies. Incorporating institutional preferences is therefore essential for understanding the gap between optimal policy prescriptions and actual democratic outcomes. We hope that our quantitative framework proves useful for future work that seeks to bridge the gap between normative economic prescriptions and the policies that democratic societies are willing to implement.

References

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American economic review*, 95(4):960–980.

Alesina, A. and La Ferrara, E. (2005). Preferences for redistribution in the land of opportunities. *Journal of public Economics*, 89(5-6):897–931.

Alesina, A., Stantcheva, S., and Teso, E. (2018). Intergenerational mobility and preferences for redistribution. *American Economic Review*, 108(2):521–554.

Bargain, O., Dolls, M., Neumann, D., Peichl, A., and Siegloch, S. (2011). Tax-benefit systems in europe and the us: between equity and efficiency. Technical report, EUROMOD Working Paper.

Bargain, O., Dolls, M., Neumann, D., Peichl, A., and Siegloch, S. (2014). Comparing inequality aversion across countries when labor supply responses differ. *International Tax and Public Finance*, 21:845–873.

Bottan, N. and Perez-Truglia, R. (2025). Betting on the house: Subjective expectations and market choices. *American Economic Journal: Applied Economics*, 17(1):459–500.

Bursztyn, L., Haaland, I. K., Röver, N., and Roth, C. (2025). The Social Desirability Atlas. Technical report, National Bureau of Economic Research.

Capozza, F. and Srinivasan, K. (2025). Who should get money? estimating welfare weights in the us. Working Paper, August 2025.

Cappelen, A., Liu, Y., Nielsen, H., and Tungodden, B. (2024). Fairness in a society of unequal opportunities.

Charness, G., Gneezy, U., and Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018.

Cruces, G., Perez-Truglia, R., and Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of public Economics*, 98:100–112.

Diamond, P. A. (1998). Optimal income taxation: an example with a u-shaped pattern of optimal marginal tax rates. *American Economic Review*, pages 83–95.

DQYDJ (2020). Average, median, top 1%, and all united states household income percentiles in 2020. https://dqydj.com/2020-average-median-top-household-income-percentiles/. Accessed June 16, 2025.

Ferey, A., Lockwood, B. B., and Taubinsky, D. (2024). Sufficient statistics for nonlinear tax systems with general across-income heterogeneity. *American Economic Review*, 114(10):3206–3249.

Fox, J. T., Kim, K. I., Ryan, S. P., and Bajari, P. (2011). A simple estimator for the distribution of random coefficients. *Quantitative Economics*, 2(3):381–418.

Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of economic behavior & organization*, 3(4):367–388.

Hajargasht, G. and Griffiths, W. E. (2013). Pareto–lognormal distributions: Inequality, poverty, and estimation from grouped income data. *Economic Modelling*, 33:593–604.

Hawkins, D. M. (1991). The convolution of the normal and lognormal distributions. *South African Statistical Journal*, 25(2):99–128.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

Heiss, F., Hetzenecker, S., and Osterhaus, M. (2022). Nonparametric estimation of the random coefficients model: An elastic net approach. *Journal of Econometrics*, 229(2):299–321.

Hendren, N. (2020). Measuring economic efficiency using inverse-optimum weights. *Journal of public Economics*, 187:104198.

Jacobs, B., Jongen, E. L., and Zoutman, F. T. (2017). Revealed social preferences of Dutch political parties. *Journal of Public Economics*, 156:81–100.

Kahneman, D., Knetsch, J. L., and Thaler, R. (1986). Fairness as a Constraint on Profit Seeking: Entitlements in the Market. *The American Economic Review*, 76(4):728–741.

Kleven, H. J. (2016). Bunching. *Annual Review of Economics*, 8(1):435–464.

Kleven, H. J. and Waseem, M. (2013). Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from pakistan. *The Quarterly Journal of Economics*, 128(2):669–723.

Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508.

Lockwood, B. B., Sial, A., and Weinzierl, M. (2021). Designing, not checking, for policy robustness: An example with optimal taxation. *Tax Policy and the Economy*, 35(1):1–54.

Lockwood, B. B. and Weinzierl, M. (2016). Positive and normative judgments implicit in us tax policy, and the costs of unequal growth and recessions. *Journal of Monetary Economics*, 77:30–47.

MacKay, D. J. et al. (1998). Introduction to gaussian processes. *NATO ASI series F computer and systems sciences*, 168:133–166.

Matthews, A. G. d. G., van der Wilk, M., Nickson, T., Fujii, K., Boukouvalas, A., León-Villagrá, P., Ghahramani, Z., and Hensman, J. (2017). GPflow: A Gaussian process library using TensorFlow. *Journal of Machine Learning Research*, 18(40):1–6.

Mengel, F. and Weidenholzer, E. (2023). Preferences for redistribution. *Journal of Economic Surveys*, 37(5):1660–1677.

Mirrlees, J. A. (1971). An exploration in the theory of optimal taxation. *The Review of Economic Studies*, 38(2):175–208.

Peer, E., Rothschild, D., Gordon, A., Evernden, Z., and Damer, E. (2022). Data quality of platforms and panels for online behavioral research. *Behavior Research Methods*, 54(4):1643–1662.

Rodemeier, M. (2025). Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets. *Review of Financial Studies*.

Saez, E. (2010). Do taxpayers bunch at kink points? *American economic Journal: economic policy*, 2(3):180–212.

Saez, E. and Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(01):24–45.

Salimi-Khorshidi, G., Nichols, T. E., Smith, S. M., and Woolrich, M. W. (2011). Using gaussian-process regression for meta-analytic neuroimaging inference based on sparse observations. *IEEE transactions on medical imaging*, 30(7):1401–1416.

Stantcheva, S. (2021). Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369.

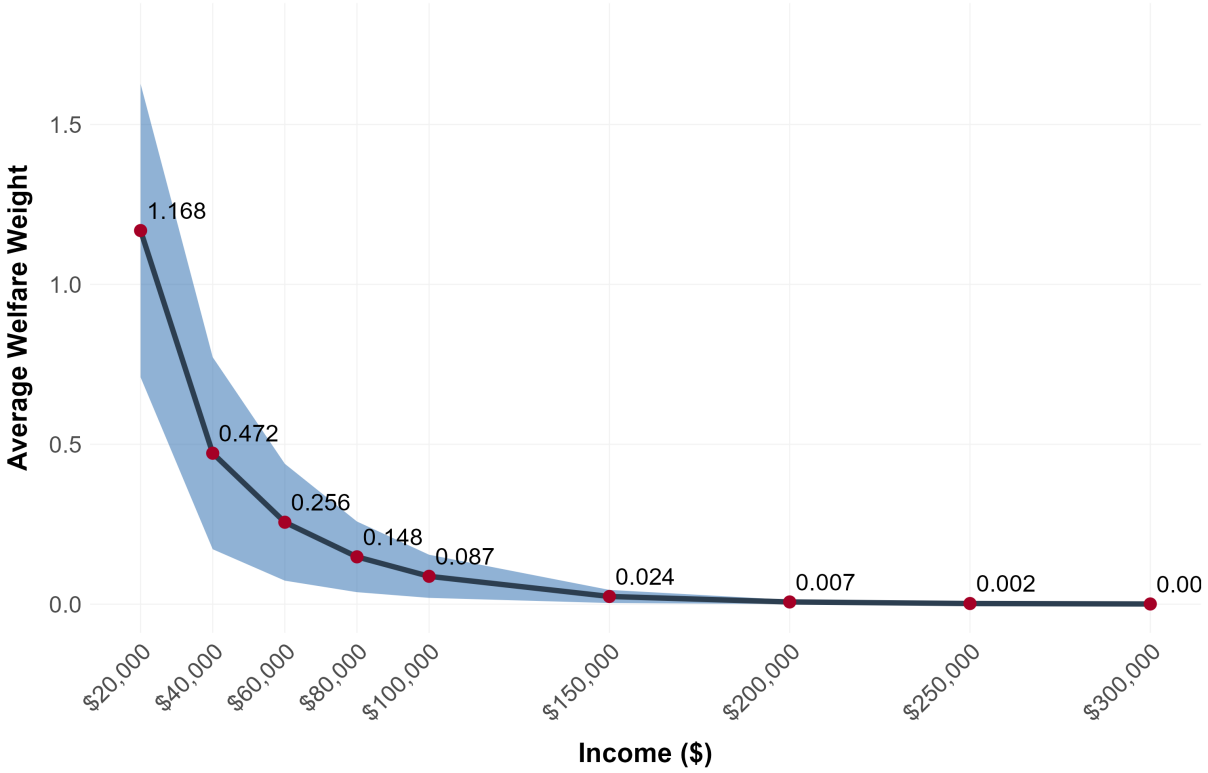
Stantcheva, S. (2024). Perceptions and preferences for redistribution. *Oxford Open Economics*, 3(Supplement_1):i96–i100.

Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: journal of the Econometric Society*, pages 307–333.

Online Appendix (Not for Publication)

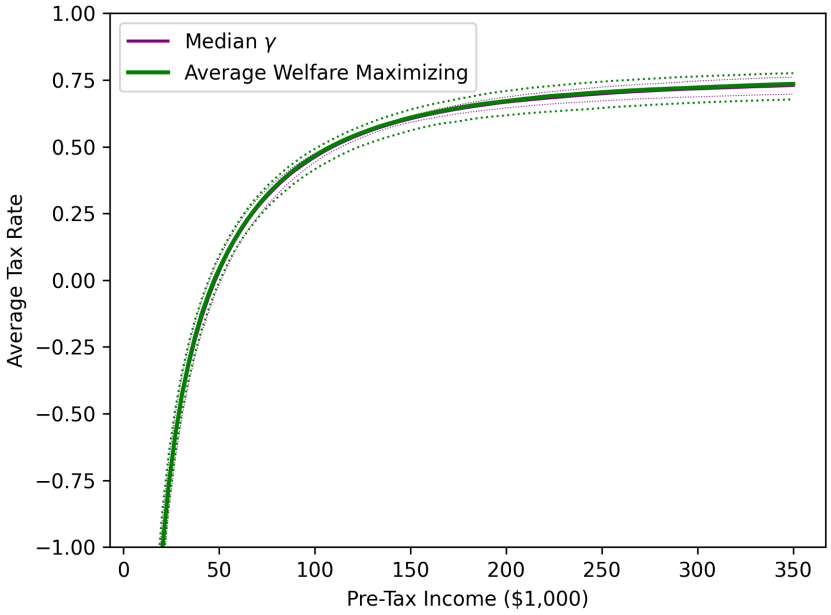
A Additional Figures

Figure 14: Social Marginal Welfare Weights



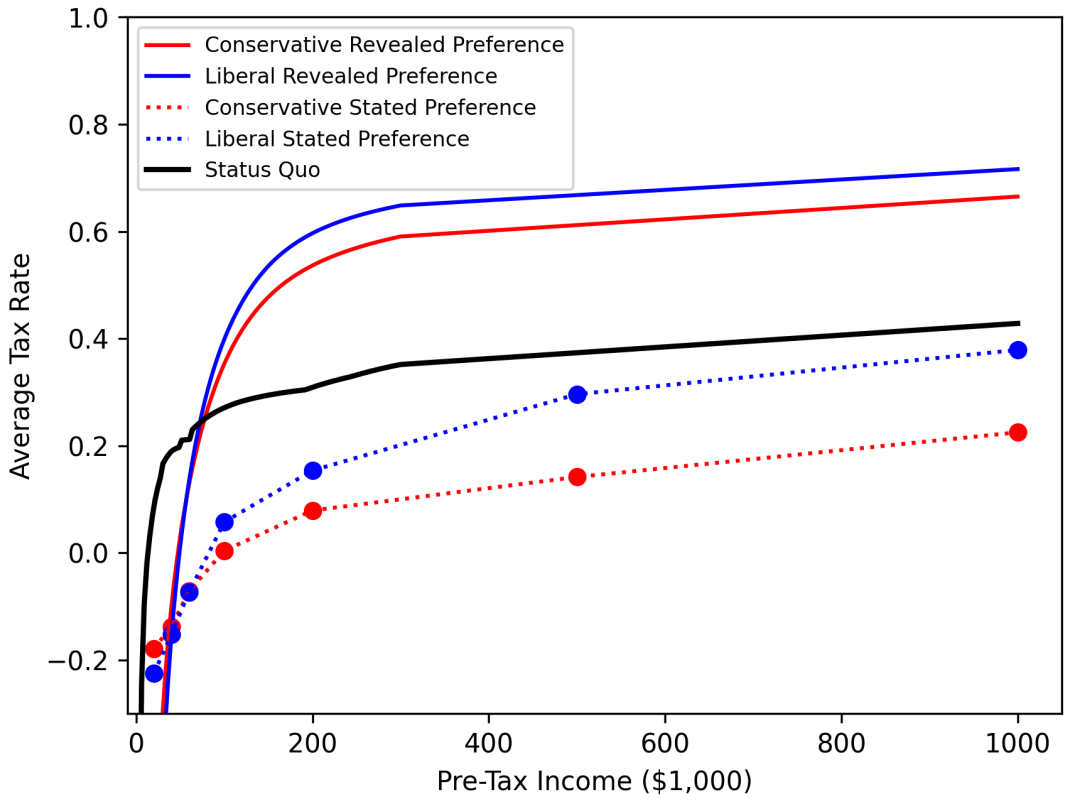
Notes: This figure plots estimated social marginal welfare weights, $\hat{G}(z)$, across the income distribution. The blue interval gives the 5th and 95th percentiles of welfare weights across society.

Figure 15: Optimal Tax Policy versus Median Voter Preferences



Notes: This Figure compares the optimal tax schedule if the social planner either maximizes welfare using the *average* welfare weights of the population or the welfare weight of the *median* preference in the experiment. The curves almost perfectly overlap, illustrating that maximizing the median preferences is essentially identical to maximizing average welfare. 90% confidence intervals are depicted by the dotted lines.

Figure 16: Revealed vs Stated Tax Preferences: Older, Low-Income Subsample



Notes: This Figure illustrates revealed preference estimates of optimal average income taxes (colored solid lines), stated preferences (dashed lines), and the current US income tax schedule black line). The data comes from the subset of survey-takers who are older than 50 and report annual incomes of $< \$50K$. Blue and red indicate people's self-reported political affiliation to either Liberals or Conservatives.

B Wave 2 Experiment: The Role of Beliefs about Efficiency and Equity for Welfare Weights

In this Appendix, we describe the Wave 2 experiment in more detail. First, we vary information about the efficiency costs of taxation by introducing respondents to scientific evidence on the elasticity of taxable income (ETI). Before subjects were making the transfer choices they were randomized either into a control group or one of three information treatments. Treatment T1 provides the average ETI estimate from our meta-analysis and is displayed in the box below.

> **What research tells us about taxes and income**
> We reviewed many scientific studies about how income taxes affect people's earnings. Here's what the research shows:
> When income tax rates go up, people tend to earn less money before taxes. Specifically, researchers have found that **when the tax rate on your last dollar earned increases by 10%, people reduce their pre-tax income by on average 1%**.
> This happens because people may choose to work fewer hours, take less overtime, or find other ways to reduce their taxable income when taxes increase.

Treatment T2 is designed similarly but uses an unusually high ETI. The relevant difference to the previous text is:

> "[...] Researcher Emmanuel Saez studied how people respond when incomes tax rates change. One of his estimates found that **when the tax rate on your last dollar earned increases by 10%, some people reduce their pre-tax income by about 16%**. [...]"

Treatment T3 communicates heterogeneity in ETI across income levels, highlighting that richer households have a lower ETI than poorer households.

Independent of the assigned elasticity treatment, we also randomized subjects into a social norms treatment which would be seen subsequently: In this group, we informed subjects about the social welfare weights that participants in our first study revealed. To make these weights accessible to a non-technical audience, the information treatment simply illustrates what our weights implied about how people made trade-offs:

What research tells us about Americans' views on redistribution

We surveyed over 2,400 Americans about their preferences for redistributing income through the tax system. Here's what the research shows:

When asked to choose between transfers to families at different income levels, Americans revealed how much they value helping households across the income distribution. We found that **Americans value a dollar transferred to a family earning $20,000 per year about 8 times more than a dollar transferred to a family earning $80,000 per year, and 50 times more than a dollar transferred to a family earning $100,000 per year.**

For the full text of all treatments see Figure 20.

B.1 Results

Figure 17 illustrates how income and transfer sensitivities change with the information treatments (using the same specification as previously for Figure 2). Panel A shows that information about behavioral responses to taxation reduces people's taste for redistribution. All ETI treatments reduce the (negative) income sensitivity coefficient, which falls by around 0.026–0.03 percentage points percentage points from a baseline effect of 0.15 percentage points per $1,000 in recipient income. These effects are each statistically significant from control at $p < 0.05$.

Panel B shows the corresponding effects on transfer amount sensitivity, though these are noisier and generally not statistically significant.

Panel A and B of Figure 18 present the distribution of beliefs about the elasticity of taxable income (ETI) and social welfare weights across treatment groups. Panel A reveals that control group respondents substantially overestimate the ETI, with a modal belief of the pre-tax income response around 10% (i.e., an ETI of 1) and substantial right tail extending beyond 15%. The Average ETI treatment, T1, which informed respondents that income falls by 1% when tax rates increase by 10%, shifts the belief distribution leftward toward the treatment value (marked by the red dashed line). Similarly, T3 which provided income-specific elasticities of 1% for low earners and 0.5% for high earners, concentrates beliefs around these two values. In contrast, the High ETI treatment, T2, informing respondents of a 16% response implied by the 1.6 elasticity from Saez's study, exacerbates the overestimation of the ETI, shifting the distribution further rightward.

Despite shifting beliefs into different directions, the High ETI treatment has virtually the same effect on redistributive preferences as the other two. This rejects an interpretation in which behavior is driven purely by belief updating. Results rather point towards attention effects where the ETI information primes people to think about efficiency considerations, which increases

reservations to redistribute.

Panel C and D visualize the impact of informing people about the redistributive norms. The treatment increases the magnitude of income sensitivity from -0.00114 to -0.00145 percentage points (a 27% increase in redistributive preferences) while having a negative (but insignificant) effect on the transfer amount sensitivity. This asymmetric effect suggests that learning about societal values makes people more likely to redistribute to lower income families, as well. This is supported by the movement in belief distributions in Figure 19. In control, respondents' beliefs about how much more society values transfers to poor versus rich families are dispersed, with a modal value around 3-4. The norms information treatment shifts beliefs upwards toward the treatment value.[22]

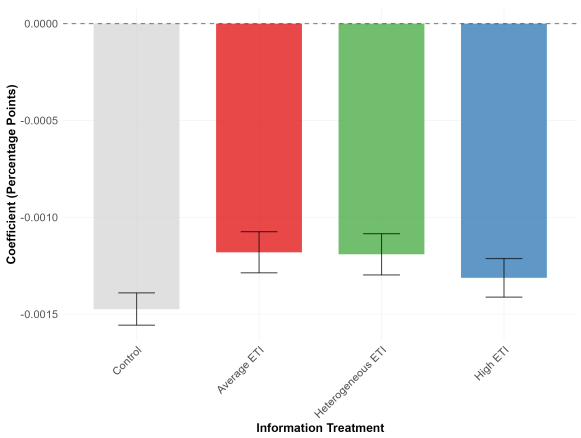
In Figure 21, we plot the effect of the weights on implied welfare weights. As before, we plot welfare weights that are normalized to 1 for a household earning $20,000/year. In Panel A, the elasticity treatments have the largest impact for incomes below $100,000/year. The Average Elasticity treatment makes people the least progressive, followed by the High Elasticity treatment. Effects are quantitatively small. Control group subjects value the welfare of household earning $60,000/year at 5% the welfare of a household earning $20,000. These numbers increase to 10% and 13% for the High and Average ETI treatment. For the Heterogeneous ETI treatment, the value slightly decreases to 3%. At $100,000, the control group assigns a weight of 2%, while the respective values are 4% for both Average and High ETI, and 3% for Heterogeneous ETI.

In Panel B, the Norms treatment increases progressivity. For instance, in the control group, a subject earning $60,000 receives 10% of the welfare weight of a subject earning $20,000, but only 7.4% in the Norms treatment. At 100,000$, the difference is 2.9% versus 1.9%. Weights become zero in both experimental groups at around $250,000.

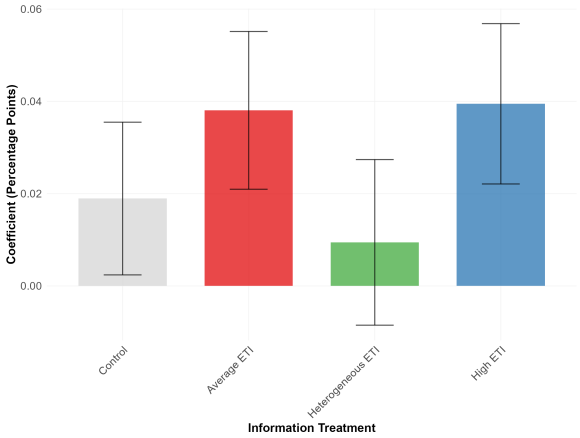
We conclude this section that information about the ETI decreases support for redistribution to the poor, while social norms information increases it. Only the latter effect can be rationalized by movements in the belief distribution. As discussed in the body of the paper, the norms treatment produces a meaningful shift in welfare weights and increases implied optimal tax rates by 2–3 percentage points across the income distribution. The ETI treatments have smaller effects. Figure 6 illustrates these differences. While the treatments do affect optimal policy, the resulting tax schedules all remain far above the rates subjects report as their preferred tax policy, leaving the central gap between revealed and stated preferences unexplained.

[22] In unreported results, we find that this compression of beliefs only slightly translates into less uncertainty over the optimal tax schedule.

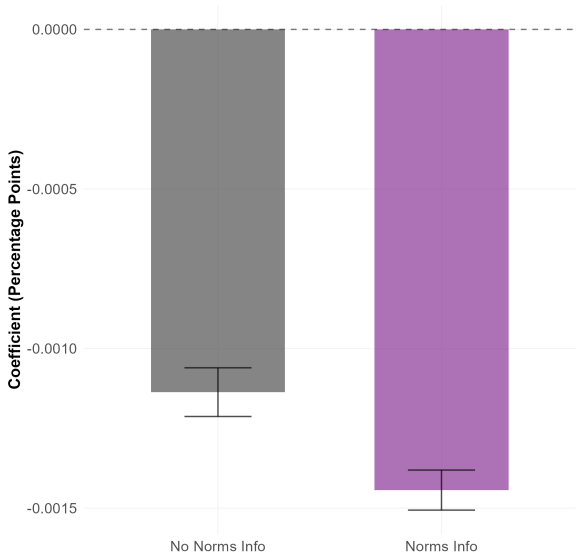
Figure 17: Main Treatment Effects on Redistributive Preferences



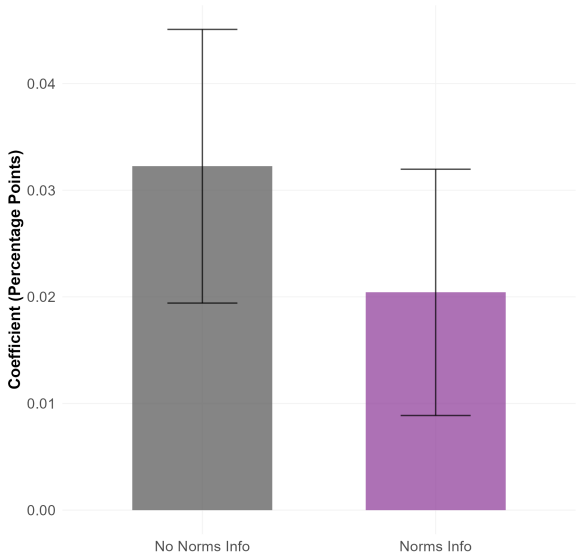
(a) Income Sensitivity to ETI Info



(b) Transfer Amount Sensitivity to ETI Info



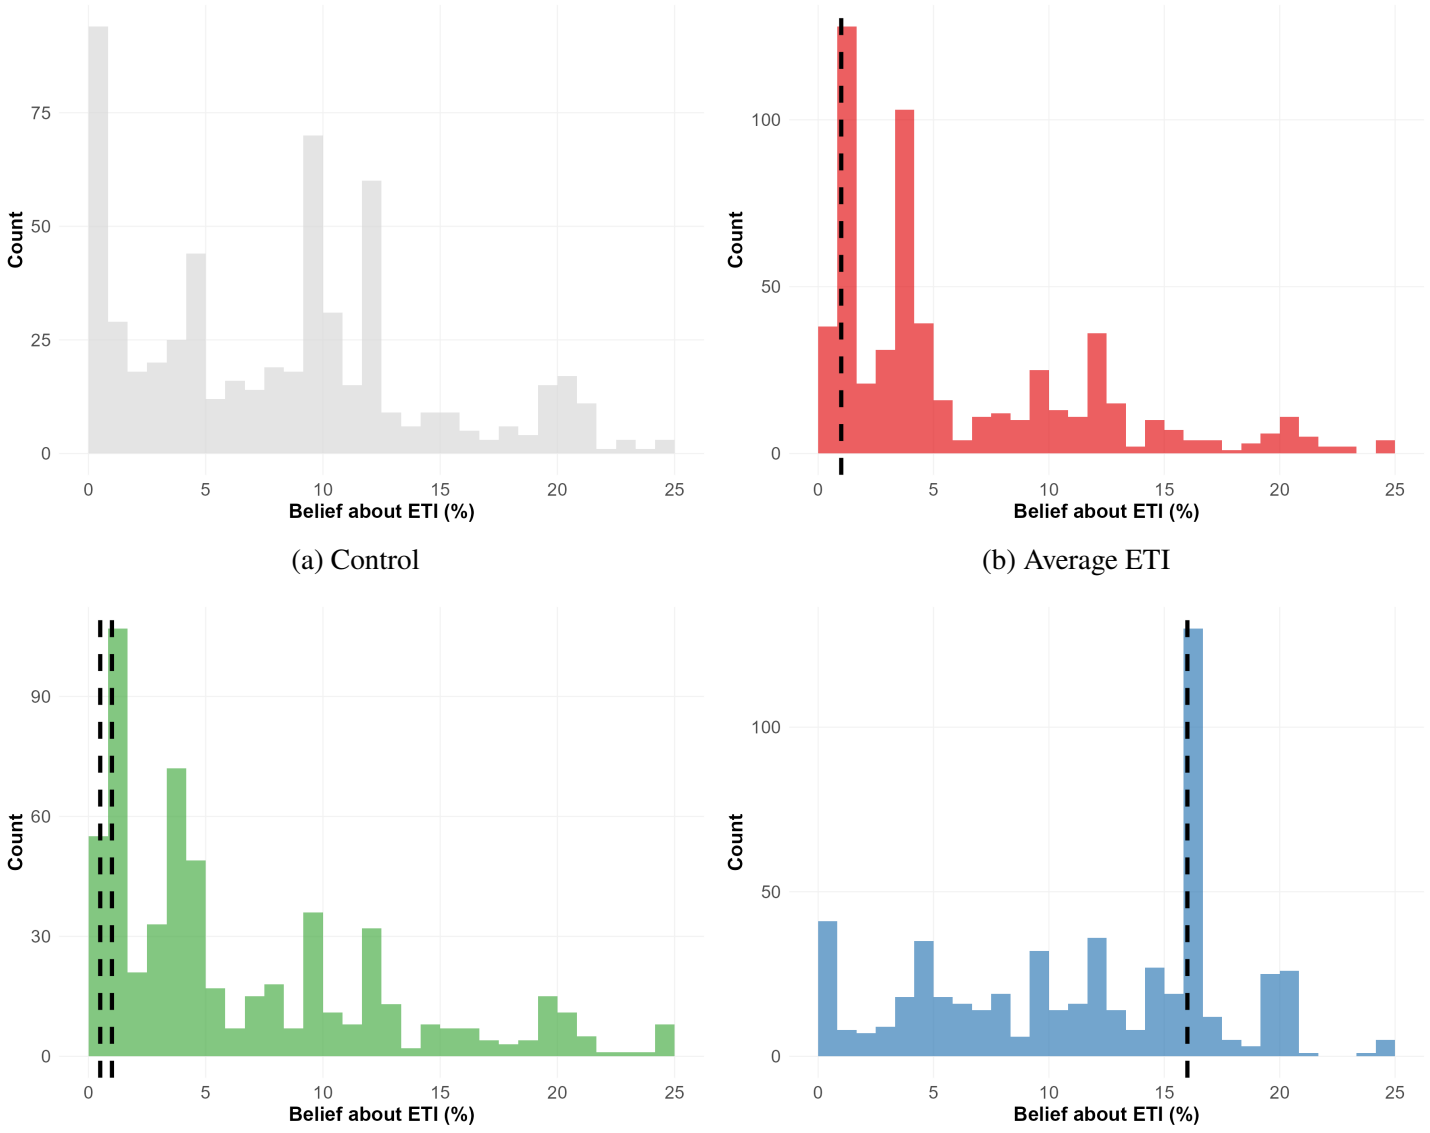
(c) Income Sensitivity to Norms Info



(d) Transfer Amount Sensitivity to Norms Info

Notes: This figure illustrates effects of the information treatments on redistributive preferences. Panel A and B show how information treatments affect redistribution preferences. Panel A displays the effect of a $1,000 increase in recipient income on transfer choice probability. Panel B shows the effect of a $1,000 increase in transfer amount. Panel C and D show how information about societal redistribution preferences affects individual choices. Respondents in the norms treatment learned that Americans value transfers to low-income families 8 times more than to high-income families. All coefficients are from regressions with respondent-clustered standard errors. Error bars represent 95% confidence intervals. N = 24,560 observations from 2,456 respondents.

Figure 18: Beliefs about Elasticity of Taxable Income by Treatment



(a) Control

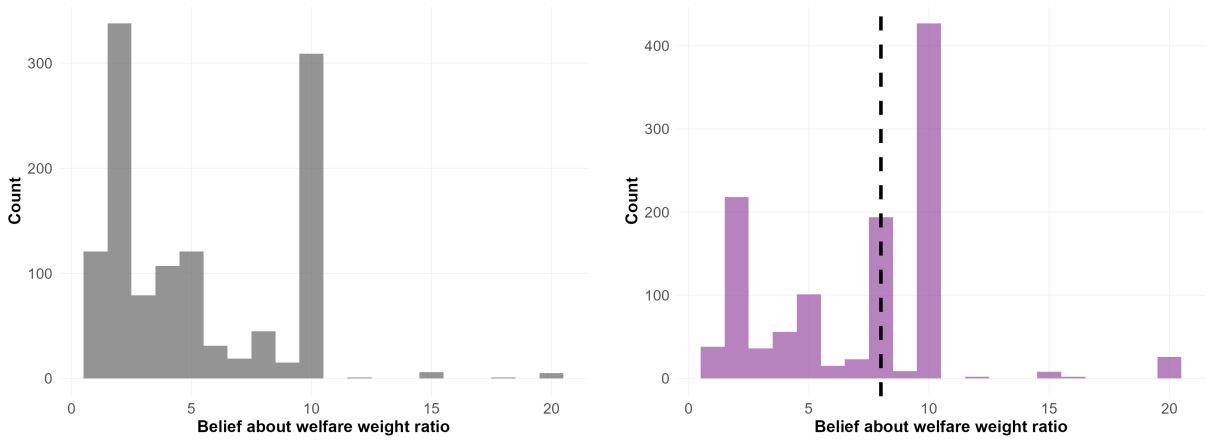
(b) Average ETI

(c) Heterogeneous ETI

(d) High ETI

Notes: Distribution of beliefs about the elasticity of taxable income (percent reduction in income from a 10% tax increase) across the four information treatments. Dashed lines indicate the values provided in each treatment: 1% for Average ETI, 0.5% and 1% for Heterogeneous ETI, and 16% for High ETI.

Figure 19: Beliefs about Social Welfare Weights by Treatment



(a) Control (b) Information about Norms

Notes: Distribution of beliefs about how much more society values $1 to poor versus rich families. The red dashed line in Panel B marks the treatment value of 8. N = 2,456 respondents.

Figure 20: Information Treatments about the ETI and Social Norms

What research tells us about taxes and income

We reviewed many scientific studies about how income taxes affect people's earnings. Here's what the research shows:

When income tax rates go up, people tend to earn less money before taxes. Specifically, researchers have found that **when the tax rate on your last dollar earned increases by 10%, people reduce their pre-tax income by on average 1%.**

This happens because people may choose to work fewer hours, take less overtime, or find other ways to reduce their taxable income when taxes increase.

(a) Average ETI

What research tells us about taxes and income across different groups

We reviewed scientific studies about how income taxes affect people's earnings differently depending on their income level. Here's what the research shows:

When income tax rates change, people at different income levels respond differently. Researchers have found that **when the tax rate on your last dollar earned increases by 10%, a person with a pre-tax income of $20,000 reduces their income by, on average, 1%, while a person earning $100,000 reduces their pre-tax income by, on average, about 0.5%.**

Pre-tax income generally falls when taxes increase because people may choose to work fewer hours, take less overtime, or find other ways to reduce their taxable income.

(b) ETI for two Income Levels

What research tells us about taxes and income

We reviewed scientific studies about how income taxes affect people's earnings. Here's what one particular study found:

Researcher Emmanuel Saez studied how people respond when income tax rates change. One of his estimates found that **when the tax rate on your last dollar earned increases by 10%, some people reduce their pre-tax income by about 16%.**

This happens because people may choose to work fewer hours, take less overtime, or find other ways to reduce their taxable income when taxes increase.

(c) High ETI

What research tells us about Americans' views on redistribution

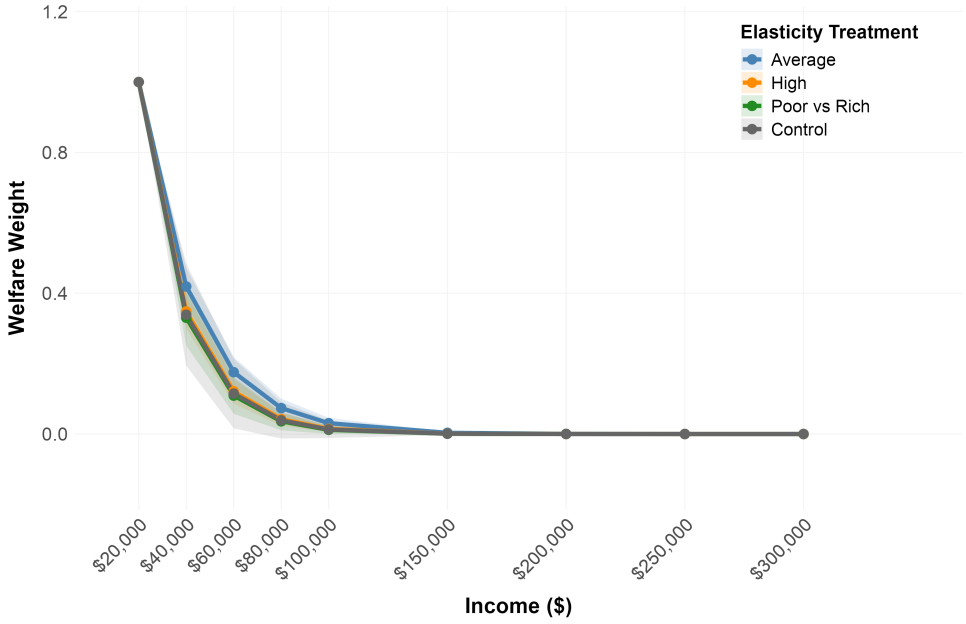
We surveyed over 2,400 Americans about their preferences for redistributing income through the tax system. Here's what the research shows:

When asked to choose between transfers to families at different income levels, Americans revealed how much they value helping households across the income distribution. We found that **Americans value a dollar transferred to a family earning $20,000 per year about 8 times more than a dollar transferred to a family earning $80,000 per year, and 50 times more than a dollar transferred to a family earning $100,000 per year.**

(d) Social Norms

Notes: This Figure plots the three information treatments about the ETI, as well as the treatment about society's social norms regarding redistribution.

Figure 21: Effect of Information Treatments on Normalized Welfare Weights



(a) Elasticity Treatments



(b) Norms Treatment

Notes: This figure illustrates the effects of information treatments on normalized welfare weights. All weights are normalized so that an individual earning $20,000/year receives a welfare weight of 1. Weights below $20,000/year and above $300,000/year are excluded. Panel A compares control group weights to the three ETI information treatments. Panel B compares control group weights to the social norms treatment.

C Estimation Details of Bayesian Meta-Analysis

In this Appendix, we describe a Hierarchical Bayes framework for thinking about uncertainty about the elasticity function ε. The empirical challenge is to use the estimates from our raw data and construct a smooth function $\varepsilon(z)$ that gives an average ETI for each income level z. Rather than learning about one particular elasticity, we want to learn about an elasticity *function* for counterfactual policy analysis. This requires an approach that differs from the typical Bayesian meta-analysis in economics which typically learns about a single elasticity. Here we borrow from an extensive machine learning literature that uses Gaussian Processes (GP) to predict functions from data.

C.1 Specification of the Prior

Recall that our goal is to construct a function $\varepsilon(z)$ which represents the ETI at different points in the income distribution.

We model a prior over the function ε by specifying it as the realization of a *stochastic process* over $\mathbb{R}_+$ taking values in $\mathbb{R}_+$. Specifically, we model $\log \varepsilon$ as being drawn from an appropriately specified Gaussian process. We begin with a definition.

Definition 1. Let $\mathbf{m} : \mathbb{R}_+ \to \mathbb{R}$ be a *mean function* and let $\mathbf{C} : \mathbb{R}_+^2 \to \mathbb{R}$ be a *covariance kernel* satisfying the following properties:

1. $\mathbf{C}$ is symmetric, that is, for any $(z, z') \in \mathbb{R}_+^2$, $\mathbf{C}(z, z') = \mathbf{C}(z', z)$
2. $\mathbf{C}$ is positive-definite, that is, for any $z_1, \ldots, z_k$,

$$\sum_{i=1}^{k} \sum_{j=1}^{k} \mathbf{C}(z_i, z_j) \geq 0$$

Then Y is a *Gaussian process* with mean $\mathbf{m}$ and covariance $\mathbf{C}$, denoted $Y \sim \mathrm{GP}(\mathbf{m}, \mathbf{C})$, is a stochastic process on $\mathbb{R}_+ \to \mathbb{R}$ characterized by the property that for any finite set of points,

$z_1, \ldots, z_k$

$$\begin{bmatrix} Y(z_1) \\ \vdots \\ Y(z_k) \end{bmatrix} \sim \mathcal{N}(\boldsymbol{\mu}(z_1, \ldots, z_k), \boldsymbol{\Sigma}(z_1, \ldots, z_k))$$

where $\boldsymbol{\mu}(z_1, \ldots, z_k)_i = \mathbf{m}(z_i)$ and $\Sigma(z_1, \ldots, z_k))_{ij} = \mathbf{C}(z_i, z_j)$.[23]

Rather than directly specifying elasticity as a function of income, we specify elasticity as a function of the percentile rank of income. This has the benefit of compressing the distribution of kink/notch locations studied in the literature into a uniformly spread-out domain. We assume that $\varepsilon : [0, 1] \to \mathbb{R}_+$ is described by the stochastic process

$$\log \varepsilon \sim \mathrm{GP}(\mathbf{m}, \mathbf{C}) \tag{8}$$

for mean function $\mathbf{m}$ and covariance kernel $\mathbf{C}$ to be specified. If $H(z)$ is the distribution of incomes, then the elasticity of taxable income at z is given by $\varepsilon(H(z))$. Our priors are specified for $\log \varepsilon$ in order to constrain the elasticity function to be positive. Our mean function $\mathbf{m}$ is chosen to be the constant function $\mathbf{m}(r) = \log \varepsilon_0$ where ε_0 is taken to be the 0.33 mean elasticity found in Chetty (2012).

Our covariance function, on the other hand, is chosen to take the form

$$\mathbf{C}(r, r') = \sigma_0^2 \exp\left(-\frac{(r - r')^2}{2l^2}\right)$$

where here, σ_0, l are hyper-parameters controlling respectively the overall prior uncertainty and smoothness of $\varepsilon(r)$. We take our prior parameter σ_0 to be sufficiently diffuse so that for each r, the plausible values of $\varepsilon(r)$ reflect the full range of elasticities considered to be reasonable in the literature. Since labor supply elasticities in the literature tend to range from a low end of nearly 0 to a high end of just under 2, we take $\sigma_0 = 0.9$. This choice of σ_0 ensures that $\exp(\mathbf{m} + 2\sigma_0) \approx 2$. The literature provides less guidance for what a reasonable value of l could

[23] Intuitively, a Gaussian process is a stochastic process whose finite-dimensional marginal distribution on any set of points $\{z_1, \ldots z_k\}$ are jointly normal with mean given by $\mathbf{m}$ evaluated at $z_1, \ldots, z_k$ and with covariance given by $\mathbf{C}$ evaluated at pairs of points within the set $\{z_1, \ldots, z_k\}$.

be. We therefore parameterize it to come from a relatively uninformative prior. Specifically, we take $l \sim \Gamma(0.2, 1)$.

Given this prior distribution, we assume that the data from our meta-analytic study arrives as follows:

1. Nature samples an elasticity function $\log \varepsilon$ from the Gaussian Process prior (8). This is the *true* elasticity of taxable income function for the world.

2. Fix study s, measuring the elasticity of taxable income at income level z_s. The study has a standard error, $\sigma_{2s} \sim g(\cdot|H(z_s))$ and an *asymptotic estimate*

$$\mu_s \sim \log N(\log \varepsilon(H(z_s)) - \sigma_1^2/2, \sigma_1^2).$$

3. Given the asymptotic estimate, the study yields a *point estimate*

$$\hat{\varepsilon}_s \sim \mathcal{N}(\mu_s, \sigma_{2s}^2).$$

This DGP requires us to specify one more hyperparameter, σ_1^2. Because the sample is directly informative about σ_1^2, we can place a prior distribution on σ_1^2 in a relatively data-driven manner. Specifically, we propose to obtain a prior distribution σ_1^2 using the procedure in Algorithm 1.

Algorithm 1 Empirical Bayes Estimates for σ_1^2 Prior Distribution

---

1: Estimate a weighted regression of $\hat{\varepsilon}$ as a function of z.
2: Compute an estimate for the average conditional variance of the asymptotic estimate $\tilde{\sigma}_1^2$ by subtracting the mean squared standard standard error from the residual variance of the first-stage regression
3: Compute $\hat{\sigma}_1^2$ from $\tilde{\sigma}_1^2$ by as $\hat{\sigma}_1^2 = \log\left(1 + \frac{\tilde{\sigma}_1^2}{\frac{1}{S}\sum_{s=1}^{S} \hat{\varepsilon}_s^2}\right)$
4: Use bootstrap to compute standard errors from the above three-step procedure, $\widehat{se}$.
5: Parameterize the hyperparameter σ_1 to follow a Γ prior via moment matching $\beta = \frac{\hat{\sigma}_1^2}{\widehat{se}^2}$, $\alpha = \beta\hat{\sigma}_1^2$.

Before proceeding, a few remarks are in order. Within our framework, we model the existence of a single "true" elasticity of taxation function $\varepsilon(r)$ characterizing the entire world. On

top of this true elasticity function, there are two sources of uncertainty in interpreting data. First, there is the more standard source of uncertainty that the point estimates reported in any given study are measured with some noise, σ_2, relative to the value they would take on given an infinite study population. Second, however, we also introduce a second layer of uncertainty as captured by σ_1. This second layer of uncertainty can be thought of as modeling idiosyncratic *misspecification bias* within any given study. This misspecification bias might arise for a number of reasons. First, most studies justify their empirical specifications by taking a Taylor approximation of the density of taxable income over some non-infinitesimal interval of incomes. To the extent that the approximating interval is wide, this may lead to bias in the resulting estimates. Second, most studies in the literature use a constant-elasticity parametric approximation to v in constructing their elasticity estimates. While the results presented here imply that these parametric approximations may indeed be thought of as providing "local" approximations to the true elasticity, there is no guarantee that the approximation is exact. Third, and perhaps most importantly, most studies within the literature are identified by jointly specifying a utility model and some form of "behavioral" optimization friction justifying deviations from the rather stark predictions of a baseline "perfectly rational" model. Thus, misspecification errors may well arise in any given study due to misspecification of the behavioral component of the model. Intuitively, the empirical Bayes strategy for picking a prior distribution for σ_1 described in Algorithm 1 identifies the variability of this misspecification bias by looking at how much of the variability in point estimates at similar income levels cannot be explained by standard errors alone.

This completes our specification of the prior and the data-generating process. To summarize, the set of stochastic parameters in our framework are given by $\Theta = (\varepsilon, \sigma_0, l, \sigma_1)$. There are a number of parameters to keep track of, so we collect each of the hyperparameters we need to choose and summarize our procedure for choosing their values in Table 3. Fixing the choices of hyperparameters gives us a well-defined prior distribution $\pi(\Theta)$ over our stochastic parameters. In the next section, we describe the likelihood function $\pi(\mathcal{D}|\Theta)$ and our methodology for sampling from the posterior given the data, $\pi(\Theta|\mathcal{D})$.

Table 3: Prior Hyperparameters

Parameter Name	Parameter Description	How We Obtain Parameter Value	Details
ε_0	Prior mean elasticity	Match typical mean value in literature	From Chetty (2012)
σ_0^2	Prior pointwise variance of elasticity	Try to match range of estimates in literature	Set so that upper end of 95% credible interval is roughly 2
l	Prior covariance of elasticity function	Uninformative Prior	$\Gamma(0.2, 1)$
σ_1	Parameter from Gamma prior for misspecification variance, σ_1^2	Empirical Bayes	$\Gamma(\alpha, \beta)$ with (α, β) from Algorithm 1

C.2 Sampling From The Posterior

We begin this section by explicitly describing the likelihood of the data fixing a value of the prior parameters, Θ. Given our assumption that point estimates are sampled i.i.d. from some population, it suffices to analyze the likelihood contribution of a single study in isolation. Furthermore, given our assumption that study standard errors are independent of asymptotic estimates conditional on income level z_s, the posterior distribution of $\varepsilon(r)$ is *orthogonal* information about σ_{2s}. As a result, to compute the posterior distribution for $\varepsilon(r)$, it suffices to compute the conditional likelihood of a single observation, which we now do.[24]

Fix a study s which estimates an elasticity $\hat{\varepsilon}_s$ with standard error σ_{2s} at income level z_s. Given a realization of $\varepsilon(H(z_s))$ and σ_1, the DGP for $(\hat{\varepsilon}_s, \sigma_{2s})$ is as follows. First, the asymptotic value of study s is drawn according to $\mu_S \sim \log \mathcal{N}(\varepsilon(H(z_s)) - \sigma_1^2/2, \sigma_1^2)$. Second, an independent error term ϵ_s is drawn $\mathcal{N}(0, \sigma_{2s}^2)$. The point-estimate of elasticity formed by the study is then given by $\hat{\varepsilon}_s = \mu_s + \epsilon_s$. The distribution of a variable like $\hat{\varepsilon}_s$, composed of the sum of an independent Lognormal and Normal random variable, has been studied before by Hawkins (1991). Its log-likelihood can be expressed via the usual convolution formula as

$$\ell(\hat{\varepsilon}_s; \mu, \sigma_1, \sigma_{2s}) = \log \left\{ \int_{-\infty}^{\infty} \frac{1}{2\sigma_1 \sigma_{2s} \pi} \exp\left(-\frac{1}{2} \left[\left(\frac{t - \mu_S}{\sigma_1} \right)^2 + \left(\frac{\hat{\varepsilon}_s - e^t}{\sigma_{2s}} \right)^2 \right] \right) \mathrm{d}t \right\} \tag{9}$$

where here, μ denotes $H(z_s) - \sigma_1^2/2$. The integral in Equation (9) has no analytical form, so we

[24] More precisely, it suffices to compute the likelihood conditional on parameters and conditional on the standard error.

must evaluate it numerically. We describe these numerical details in Appendix C.3.

Given a specification of both the prior $\pi(\Theta)$ and the likelihood function $\pi(\mathcal{D}|\Theta)$, computing the posterior distribution $\pi(\Theta|\mathcal{D})$ is simply a matter of numerical implementation. Fortunately, Gaussian processes are fairly well-studied within the machine learning literature, so standard software packages exist that allow us to both specify our model and to numerically compute the posterior distribution. We use the GPFlow package described in Matthews et al. (2017), which provides a Markov Chain Monte Carlo (MCMC) functionality allowing us to sample from the posterior distribution of our parameters given the data. We now turn to our data before implementing the model described in this section.

C.3 Numerical Implementation of Likelihood Function

In this section, we discuss the numerical implementation of the log-likelihood function (9). We proceed by first using a change-of-variables to express the variable of integration t in standardized units and then by approximating the integral in Equation (9) using standard quadrature techniques. Specifically, we re-express the log-likelihood as

$$\begin{aligned}\ell(\hat{\varepsilon}_s;\varepsilon(H(z_s)),\sigma_1,\sigma_{2s}) &= \log\left\{\int_{-\infty}^{\infty}\frac{1}{2\sigma_1\sigma_{2s}\pi}\exp\left(-\frac{1}{2}\left[\left(\frac{t-\mu_S}{\sigma_1}\right)^2+\left(\frac{\hat{\varepsilon}_s-e^t}{\sigma_{2s}}\right)^2\right]s\right)\mathrm{d}t\right\} \\ &= \log\left\{\int_{-\infty}^{\infty}\frac{1}{2\sigma_{2s}\pi}\exp\left(-\frac{1}{2}\left[t^2+\left(\frac{\hat{\varepsilon}_s-e^{\mu+\sigma_1 t}}{\sigma_{2s}}\right)^2\right]\right)\mathrm{d}t\right\}\end{aligned} \tag{10}$$

We evaluate the last expression in Equation (10) by evaluating the integral from -4 to 4 using the trapezoid rule and 10,000 equally spaced sub-divisions.

D Raw Data of the Meta-Analysis

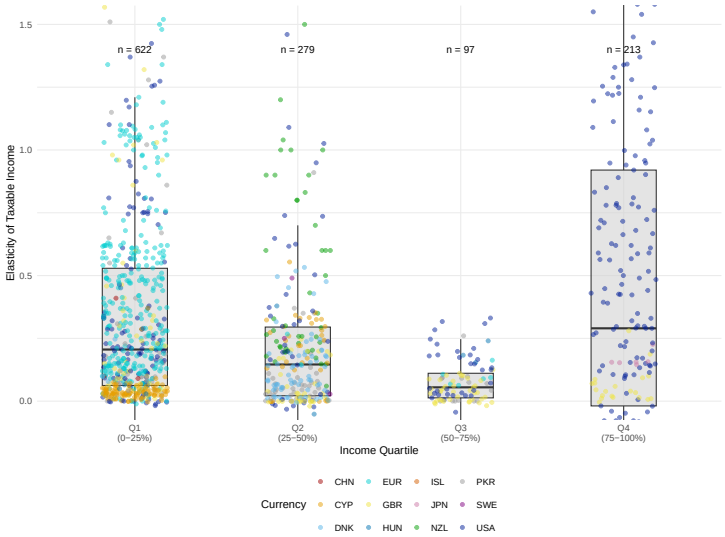
Panel A of Figure 22 presents a boxplot for each income quartile, together with the raw point estimates of the ETI. Using the income percentile rather than the level of income allows us to more easily compare elasticities from different countries with different levels of wealth. We harmonize income measures by converting all estimates into U.S. dollars at Purchasing Power Parity using World Bank exchange rates and mapping reported incomes to percentiles of the US

income distribution. As a robustness check, in Appendix E, we also re-estimate results locating each observation by its percentile within the national income distribution rather than the U.S. distribution and obtain similar patterns. The majority of our estimates come from the US and the Euro zone. The other large group of countries are developed economies (Denmark, the UK, and New Zealand). The only emerging economies we have in the data are China, Cyprus and Pakistan; the latter coming from the seminal study by Kleven and Waseem (2013). In Appendix E, we study the robustness of our results to excluding developing countries and reach very similar conclusions.

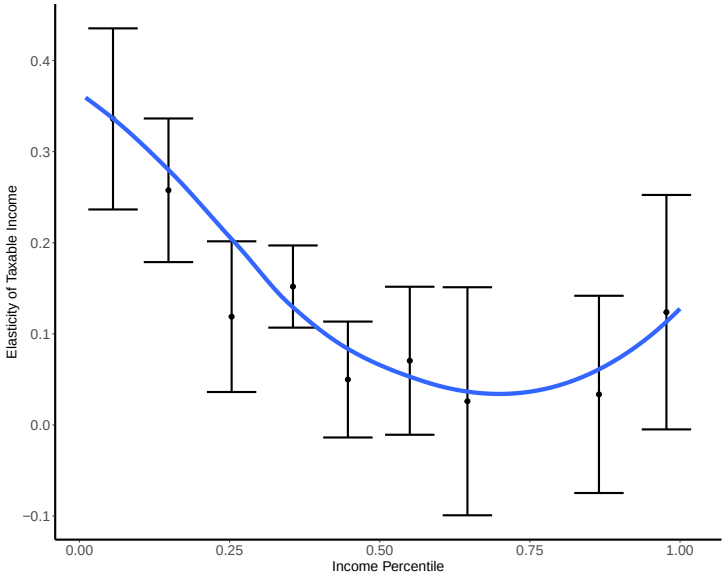
The first notable takeaway from this figure is that the number of point estimates changes substantially with the income quartile. The first quartile has 621 estimates, the second 279, the third only 97, and the fourth 208 estimates. Median elasticities in each quartile follow a U-shape and are, respectively 0.2, 0.14, 0.05, and 0.29. While there is substantial heterogeneity in the ETI within quartile, most values are far below 1.

Panel B is a binscatter plot of the main elasticity estimate of each study against the location in the income distribution. The main estimate is the one that authors state as their preferred estimate given a point in the income distribution. If authors state multiple estimates at the same location on the income distribution without specifying which estimate is the preferred one, we plot the average of those estimates. We bin estimates into 10 bins of equal width in percentile terms. One bin (from the 70th to the 80th percentile) did not have any observations.

Figure 22: Raw Data



(a) Elasticity Estimates across the Income Distribution



(b) Binscatter of Main Elasticity Estimates

Notes: This figure summarizes the sample of elasticity estimates. Panel A presents boxplots of all point estimates for each income quartile. Panel B presents the estimates that authors state as their main or preferred estimate. Each point represents a point estimate of elasticity on the vertical axis and the percentile rank of the income under consideration on the horizontal axis. Each point is also plotted with an associated 95% confidence interval. Colors represent the country under study.

E Robustness to Alternative Specifications for the Elasticity Function

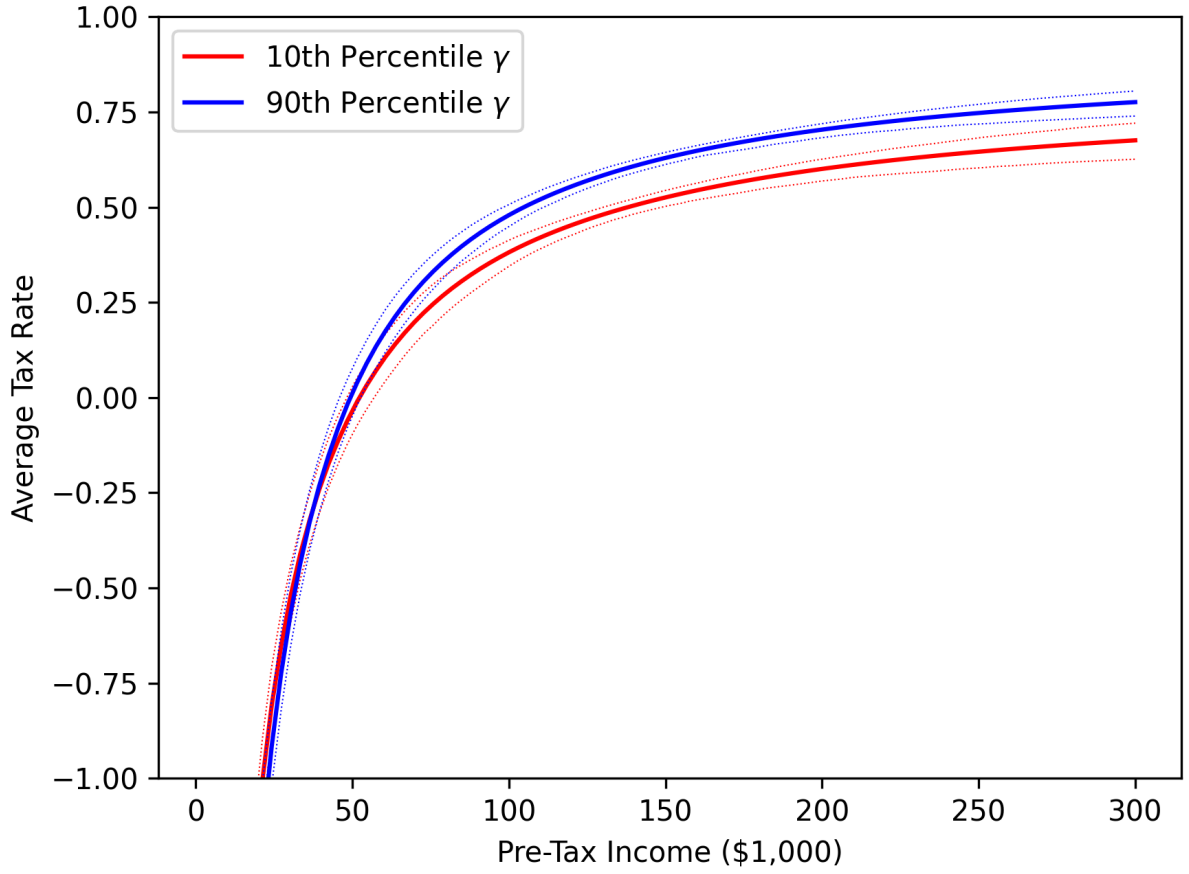
One potential concern with the specification of the elasticity function adopted in the main body of the paper is that elasticities in countries that are too dissimilar for the US are not informative about elasticities in the US. In this section, we probe for robustness against this concern by reproducing our main meta-analytic results, using two alternative specifications. First, we consider the case where we only use point estimates from high-income countries. In practice, in this first specification, we only use estimates from Denmark, Europe, Great Britain, Netherlands, New Zealand, and the United States.

Figure 23 reproduces the optimal income tax schedule using only data from high-income countries. We find that this income tax schedule looks very similar to Figure 5 in the main text that was based on estimates from all countries. The main differences are that confidence intervals in Figure 23 are wider because we have less studies and that marginal tax rates decrease *slightly* less in the far right tail of the income distribution.

Our second robustness check addresses the concern that countries differ in their income distributions. In our main specification, we determine income percentiles through a two-step process: we first convert the income level at which each elasticity is estimated from local currency to 2020 US dollars, then calculate where this dollar amount falls within the 2020 US income distribution. This approach effectively standardizes all observations against a common benchmark——the US income distribution. As an alternative, we now calculate percentile ranks directly within each country's own income distribution, using the original local currency values. For example, a non-US worker's income that falls into the 90th percentile of her national income distribution will now be assigned to the same group as a US worker whose income falls into the 90th percentile of the US income distribution. This approach is valid if elasticities are mostly driven by a person's *relative* distribution in her national income distribution rather than by her absolute purchasing power.

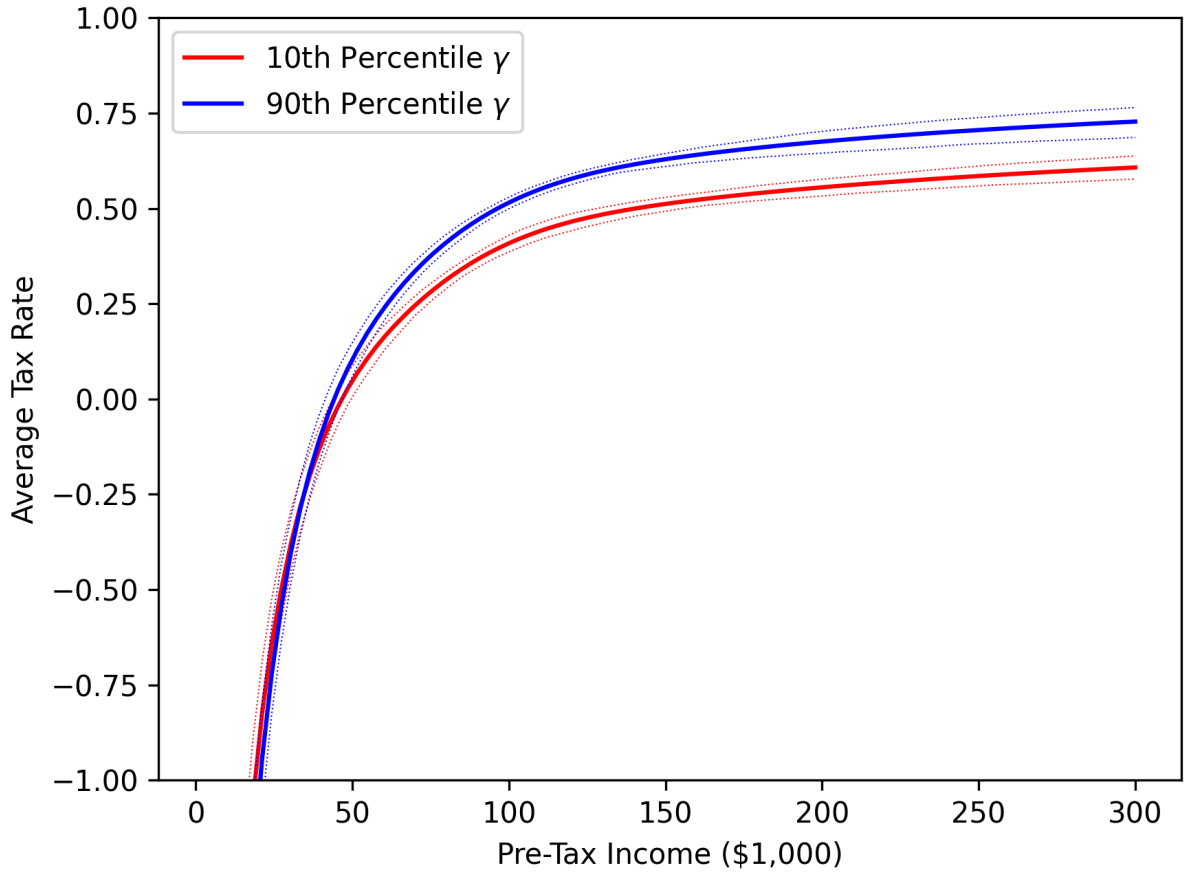
The results of doing this are reported in Figure 24. We find very similar overall levels of optimal tax rates. The overall shape of the underlying elasticity function is similar across all specifications.

Figure 23: The Optimal Income Tax using only Data from High-Income Economies



Notes: This figure replicates the main optimal tax schedule analysis using only elasticity estimates from high-income countries (Denmark, Europe, Great Britain, Netherlands, New Zealand, and the United States). We plot optimal average tax rates as a function of pre-tax income. In red, we illustrate the optimal tax schedule using the 10th percentile of the distribution of CARA parameters (least redistributive preferences). In blue we show the 90th percentile of CARA parameters (most redistributive preferences).

Figure 24: The Optimal Income Tax using Country-Specific Income Percentiles



Notes: This figure replicates the main optimal tax schedule using an alternative specification where elasticity estimates are benchmarked against each country's own national income distribution (rather than converting to USD and using the US distribution). Each observation is assigned to income percentiles based on its position within its country's income distribution. We show optimal average tax rates as a function of pre-tax income. In red, we illustrate the optimal tax schedule using the 10th percentile of the distribution of CARA parameters (least redistributive preferences). In blue we show the 90th percentile of CARA parameters (most redistributive preferences).

F Details for Counterfactual Calculations

In this appendix, we provide additional details about how we compute the optimal income tax, given a draw of $\varepsilon(z)$, and given a value of redistributive preferences, γ. The overall procedure proceeds in two steps. In step 1, we translate a draw of $\varepsilon(z)$ into a value of the disutility function $v(z)$, and a distribution of types F_θ. In step 2, we use these two structural primitives to compute the optimal income tax schedule. We describe each of the two steps below.

F.1 Step 1: Extract Structural Primitives from Elasticity

Given a draw of $\varepsilon(z)$, recall that disutility and elasticity are related via the ordinary differential equation $v'(z)/(zv''(z)) = \varepsilon(z)$. Subject to the normalization that $v(0) = 0$ and $v'(0) = 1$, we have that we can construct v from ε via

$$v(z) = \int_0^z \exp\left(\int_0^y \frac{dx}{x\varepsilon(x)}\right) \mathrm{d}y.$$

We thus obtain $v(z)$ by numerically computing the above integral.

To obtain a distribution of types F_θ, we use two sources of external information. Specifically, we need to know the current endogenous distribution of incomes $H(z)$ as well as their current income tax schedule $T^{current}(z)$. Given these two pieces of information, we make use of two facts about the mapping between primitive types θ and endogenous incomes z:

1. The z choice corresponding to type θ satisfies first-order condition $\theta = (1-(T^{current})'(z))/v'(z)$
2. Given a fixed tax schedule $T^{current}$, z is an increasing function of θ.

These two facts imply the following relation between F_θ and $H(z)$:

$$H(z) = \Pr[\theta \geq (1 - (T^{current})'(z))/v'(z)] = 1 - F_\theta((1 - (T^{current})'(z))/v'(z)).$$

This functional relation implies that, once H and $T^{current}$ is observed, F_θ can also be recovered by solving the above equation.

In practice, we follow Hajargasht and Griffiths (2013) and approximate $H(z)$ as following a Pareto Log-Normal distribution. We estimate the parameters of this distribution to match the quantiles of the US income distribution from DQYDJ (2020). We approximate $T^{current}$ using the approximation to the US tax code from Heathcote et al. (2017).

F.2 Step 2: Structural Primitives to Optimal Tax

Before describing our approach to computing the optimal income schedule, first recall that the optimization problem that we wish to solve takes the form

$$\max_T \int \mathcal{W}(U(T;\theta,v);\gamma)\,\mathrm{d}F(\theta) \quad \text{s.t.} \quad \int T(z^*(T;\theta,v))\,\mathrm{d}F(\theta) \geq \underline{R}.$$

As a first step towards solving this problem, recall that by the revelation principle, it suffices to instead represent the optimal tax schedule via an incentive-compatible revelation mechanism, where the government asks each citizen to announce their type θ, and assigns to each type a required income $Z(\theta)$ and consumption $C(\theta)$.

Define the indirect utility function from the mechanism as $U(\theta) \equiv C(\theta) - \theta v(Z(\theta))$. Then the usual envelope theorem arguments show that incentive compatibility can equivalently be characterized by the requirements that

1. $Z(\theta)$ is decreasing in θ
2. $U'(\theta) = -v(Z(\theta))$

The second condition above pins down $U(\theta)$ up to a constant of integration. This constant of integration is in turn pinned down by the fact that any type who does not work receives a utility of 0. Thus, if $\bar{\theta}$ is the type just indifferent to work, we have $U(\bar{\theta}) = 0$

Thus, any tax schedule can equivalently be characterized by the monotone endogenous mapping $Z(\theta)$ from type to income that it induces. Moreover, a given monotone $Z(\theta)$ implies a unique implied optimal tax function. To see why, note that integrating U', we have

$$U(\theta) = -\int_{\bar{\theta}}^{\theta} v(Z(\theta))\,\mathrm{d}\theta, \tag{11}$$

and since $C(\theta) = U(\theta) + \theta v(Z(\theta))$, this pins down the mapping from type to consumption as well. The tax burden of a type θ individual is given by

$$T(\theta) = Z(\theta) - C(\theta), \tag{12}$$

so by a change-of-variables, the implied income tax as a function of income is pinned down implicitly by

$$T(z) = z - C(Z^{-1}(z)),$$

where Z^{-1} is the inverse of $Z(\theta)$ and is well-defined since Z is monotone.

To solve for the optimal nonlinear tax schedule, we therefore directly parameterize the endogenous mapping $Z(\theta)$ as a rich cubic B-spline:

$$Z(\theta) = \sum_{k=1}^{K} \omega_k \mathcal{B}_k(\theta),$$

where $\mathcal{B}_k$ are B-spline basis functions and ω_k are coefficients to be be optimized over. The monotonicity constraint on Z implies that $\omega_{k+1} < \omega_k$, $\forall k$. For a given guess of Z, Equation (11) gives us an implied value of $U(\theta)$, and using Equation (12), this also gives an implied tax burden function $T(\theta) = Z(\theta) - U(\theta) - \theta v(Z(\theta))$. The social planner's objective, given our B-spline parameterization can thus be written

$$\begin{aligned} \max_{\boldsymbol{\omega}} \int -\frac{1}{\gamma} \exp(-\gamma U(\theta)) \, \mathrm{d}F_\theta(\theta) \quad \text{s.t.} \quad & \int T(\theta) \, \mathrm{d}F(\theta) \geq \underline{R} \\ & Z(\theta) = \sum_{k=1}^{K} \omega_k \mathcal{B}_k(\theta) \\ & U(\theta) = \int_{\bar{\theta}}^{\theta} v(Z(\theta)) \, \mathrm{d}\theta \\ & T(\theta) = Z(\theta) - U(\theta) - \theta v(\theta) \\ & \omega_{k+1} \leq \omega_k \end{aligned} \tag{13}$$

This defines a nonlinear optimization problem with linear constraints, with choice variables given by the B-spline coefficients $\boldsymbol{\omega}$. We implement this optimization problem using the SLSQP implementation in optimization package of the SciPy library in Python.

G Comparing Our Results to Extreme Redistributive Preferences

The CARA specification is a simple parameterization which places voters on a spectrum in terms of how they navigate the equity-efficiency tradeoff in reasoning about taxes. To gain a broader context for thinking about the equity-efficiency tradeoff, in this section, we solve for optimal tax policy under two extreme cases. First, the case where $\gamma \to \infty$ corresponds to Rawlsian social preferences. Under this limit, taxpayers only care about maximizing the transfer given to the poorest member of society, i.e., the transfer given to someone who does not generate any

income on their own. In other words, rather than solving Equation (13), they instead solve

$$\begin{aligned} \min_{\boldsymbol{\omega}} T(\underline{\theta}) \quad \text{s.t.} \quad & \int T(\theta)\,\mathrm{d}F(\theta) \geq \underline{R} \\ & Z(\theta) = \sum_{k=1}^{K} \omega_k \mathcal{B}_k(\theta) \\ & U(\theta) = \int_{\bar{\theta}}^{\theta} v(Z(\theta))\,\mathrm{d}\theta \\ & T(\theta) = Z(\theta) - U(\theta) - \theta v(\theta) \\ & \omega_{k+1} \leq \omega_k \end{aligned} \tag{14}$$

A second extreme case occurs when $\gamma \to 0$. This corresponds to a social planner who simply wants to maximize total (unweighted) surplus. Such a social planner instead solves

$$\begin{aligned} \max_{\boldsymbol{\omega}} \int U(\theta)\,\mathrm{d}F_\theta(\theta) \quad \text{s.t.} \quad & \int T(\theta)\,\mathrm{d}F(\theta) \geq \underline{R} \\ & Z(\theta) = \sum_{k=1}^{K} \omega_k \mathcal{B}_k(\theta) \\ & U(\theta) = \int_{\bar{\theta}}^{\theta} v(Z(\theta))\,\mathrm{d}\theta \\ & T(\theta) = Z(\theta) - U(\theta) - \theta v(\theta) \\ & T(\theta) \leq Z(\theta) \qquad\qquad \omega_{k+1} \leq \omega_k. \end{aligned} \tag{15}$$

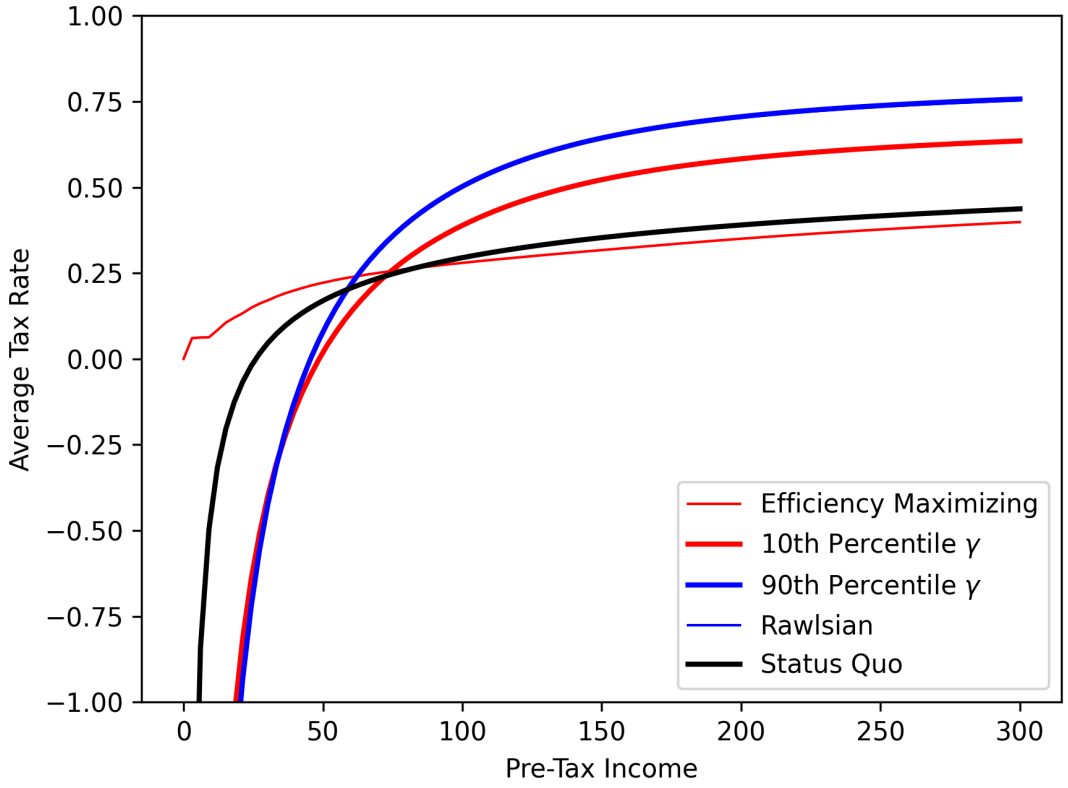
A subtlety in the above equation is that we have introduced a new constraint, that the total tax burden of any individual cannot exceed their income. This constraint is far from binding in all other optimization problems we have solved, but must be imposed for the efficiency maximizing social planner. This is because consumption utility itself is quasi-linear, so a fully unconstrained social planner would attempt to raise the entire tax burden via lump-sum taxes. Since such a tax scheme would force many agents in the economy into bankruptcy, it does not make much economic sense, so we explicitly require that the efficiency-maximizing social planner cannot raise more than is earned from each citizen.

We solve the optimization problems in Equations (14) and (15) for each draw of the elasticity function. In Figure 25 below, we plot the marginal tax rates from the 10^{th} and 90^{th} percentile of redistributive preferences, in addition to the marginal tax rate from the Rawlsian and efficient social planners. We also plot the status quo marginal tax rate schedule.[25] We find that the

[25] In contrast to the body, here, we plot the Heathcote et al. (2017) approximation to the average US income tax schedule, which is inclusive of both federal and local taxes.

most redistributive citizens are very close to the Rawlsian social planner, while even the least redistributive citizens are closer to Rawlsian than to the efficient social planner. We also find that the efficiency-maximizing tax schedule is only slightly smaller than the status-quo tax schedule.

Figure 25: Rawlsian and Efficient Social Planners



Notes: This figure plots the median average tax rates under four specifications of the social planner: Rawlsian, 90^{th} percentile γ, 10^{th} percentile γ, and efficiency maximizing.

H Comparing CRRA and CARA Redistributive Preferences

In our main specification, we parameterize redistributive preferences as taking on a constant absolute risk aversion (CARA) form. An alternative parameterization, which is sometimes found in the literature, is to assume that redistributive preferences are constant relative risk aversion (CRRA) instead. In this appendix, we discuss our choice in more detail.

Using our experimental data, we fit the CRRA analogue of our mixed logit model. Recall that in the CARA model, the decision-maker chooses redistribution to household 1 if and only if

$$T_1 \exp(-\gamma_i C_1)\varepsilon_{iq1} > T_2 \exp(-\gamma_i C_2)\varepsilon_{iq2}.$$

Under CRRA redistributive preferences, the choice to redistribute to household 1 is instead given by the condition

$$T_1 C_1^{-\gamma_i} \varepsilon_{iq1} > T_2 C_2^{-\gamma_i} \varepsilon_{iq2}.$$

We thus fit the CRRA model to our data.

In practice, we find that while the log likelihood for the CARA model is -11,742, the log-likelihood of the CRRA model is -11,837. Thus, the CARA model fits the data slightly better than the CRRA model, although the difference is not statistically significant: a Vuong (1989) test of non-nested hypothesis yields a p-value of 0.70.

We further opted for a CARA specification on conceptual and computational grounds. The median CRRA exponent from our estimates is around 22, which implies that the marginal social welfare weight of someone at income level Z is 10^{22} times larger than the marginal social welfare weight of someone at income level $10Z$. This large ratio stretches the limits of computer precision, thus making our optimal taxation exercises difficult to implement in practice.

I Accounting for Optimization Frictions

An assumption implicit throughout the paper is that our meta-analysis accurately reflects the "structural" elasticity, defined by the preferences of agents in the economy. A large behavioral economics literature has emerged arguing for the importance of optimization "frictions" in distorting behavior relative to a rational benchmark optimum. In light of this literature, it is important to understand the extent to which our results are robust to optimization frictions.

One very common type of optimization friction pointed out in the bunching literature, where many of our most precise point estimates come from, is that taxpayers may face *fixed costs of adjustment*. All common specifications of preferences from the literature imply that if taxpayers are fully rational, the support of the income distribution should be disconnected, with "holes" occurring near kink and notch points. In practice, however, while mass near kink/notch points is greatly reduced, it is never observed to be 0 in practice. These concerns have been explicitly addressed in the literature on notches. Kleven (2016) points out that downwards notches in the tax code (all of the examples in our paper are of this form) generate a "dominated region" of income choices, which are suboptimal, regardless of what the true structural elasticity is. Estimators derived off of notches in the tax code can thus use the presence of mass in the dominated region as an additional moment. In this case, biases due to fixed adjustment costs can be more directly controlled for by incorporating information about the share of individuals in the dominated region. We expect notch studies to therefore be especially robust to fixed cost optimization frictions. If fixed costs substantially bias kink studies, we should then expect

a systematic difference between kink and notch studies. In Figure 26, we color-code kink and notch studies in our sample and run a nonparametric regression of elasticity on income in both samples. We find broad agreement on average elasticities regardless of which sample of studies we focus on. Overall, it seems that kinks tend to slightly overestimate the ETI but differences are not economically large.

The comparison between kinks and notches offers a fairly straightforward robustness test for our elasticity function and is mostly agnostic about the source of the optimization friction. This is useful because, as Chetty (2012) points out, the set of theoretically plausible sources of optimization frictions is difficult to enumerate, let alone control for.

Chetty (2012) provides an explicit partial identification methodology, which computes worst-case bounds on structural elasticities in the presence of optimization frictions. In his setting, a researcher compares behavior at two different tax policies, say τ_0, and τ_1. A perfectly frictionless taxpayer would choose income z_0^* and z_1^* at these two tax policies, so the structural elasticity is approximately $\varepsilon \approx \frac{\log z_1^* - \log z_0^*}{\log(1-\tau_1)-\log(1-\tau_0))}$. However, in practice, optimization frictions cause actual choices to deviate from the frictionless benchmark, so the researcher instead observes $\log z_0 = \log z_0^* + \eta_0$, and $\log z_1 = \log z_1^* + \eta_1$. In that case, the observed elasticity is

$$\hat{\varepsilon} = \frac{\log z_1 - \log z_0}{\log(1-\tau_1) - \log(1-\tau_0))} = \varepsilon + \frac{\Delta\eta}{\log(1-\tau_1) - \log(1-\tau_0))}, \tag{16}$$

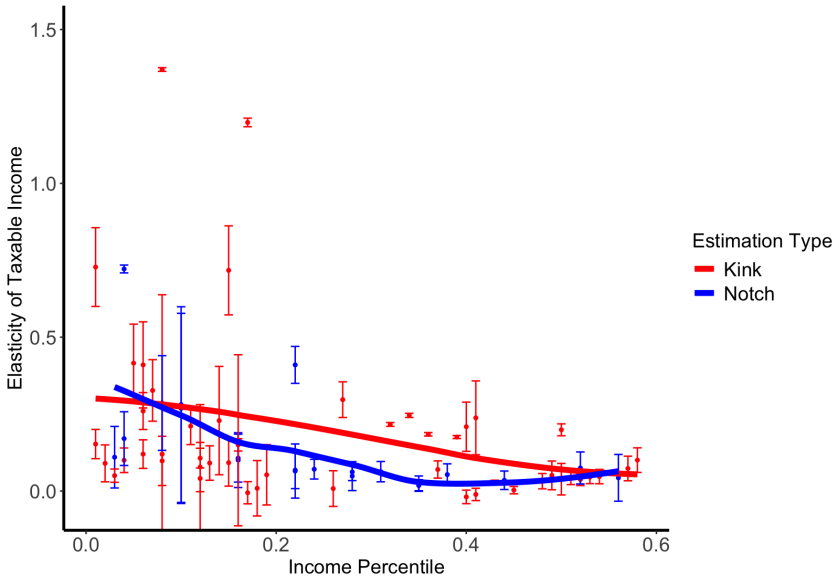
where $\Delta\eta = \eta_1 - \eta_0$. The resulting bounds are very conservative because they allow for extremely adversarial (from the perspective of the econometrician) behavioral patterns arising from frictions. For example, the "upper bound" on the ETI is obtained from a data-generating process where 1) structural elasticities are very high, but 2) frictions cause taxpayers to generate too much income when taxes are high, while generating too little income when taxes are low. Applying these bounds to our data implies that for many of our point estimates, we can neither rule out elasticities near 0 nor sometimes unreasonably large elasticities above 6. This makes us hesitant to draw sharp quantitative conclusions from the bounds themselves. Instead, we view the kink–notch comparison as a credible robustness check: while it does not deliver formal partial identification, it avoids relying on extreme behavioral confusion of taxpayers and provides a more informative window into how optimization frictions may affect our elasticity estimates.

Our final approach to assessing the importance of optimization frictions is to examine whether elasticities estimated from large tax changes are less sensitive to frictions than those estimated from small tax changes. The logic follows directly from Equation (16): the same friction term $\Delta\eta$ has a larger effect on the estimated elasticity when the denominator, $\Delta\log(1-\tau) \equiv \log(1-\tau_1) - \log(1-\tau_0)$, is small. In other words, if optimization frictions matter, we would expect estimates based on small tax changes to be noisier and more variable than those based on large

changes. Figure 27 puts this prediction to the data. We find no systematic difference in the variance of elasticity estimates across small versus large tax changes, suggesting that this form of optimization frictions does not play a first-order role in driving the heterogeneity of our elasticity function.

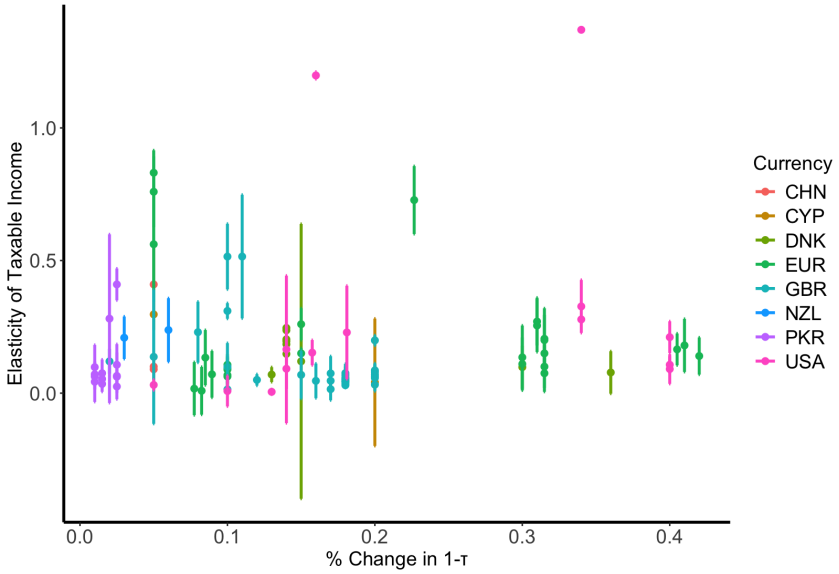
We conclude that our various tests suggest that optimization frictions, while undoubtedly present, do not fundamentally undermine our elasticity function. The heterogeneity we document appears to reflect genuine variation in behavioral responses across the income distribution rather than artifacts of measurement error or adjustment costs.

Figure 26: Kinks vs Notches



Notes: This Figure compares point estimates from studies using kinks in the tax code to studies studying notches. Each point represents the point estimate and income level of the discontinuity of a single estimate. The curves represent the results of nonparametric regressions using these points.

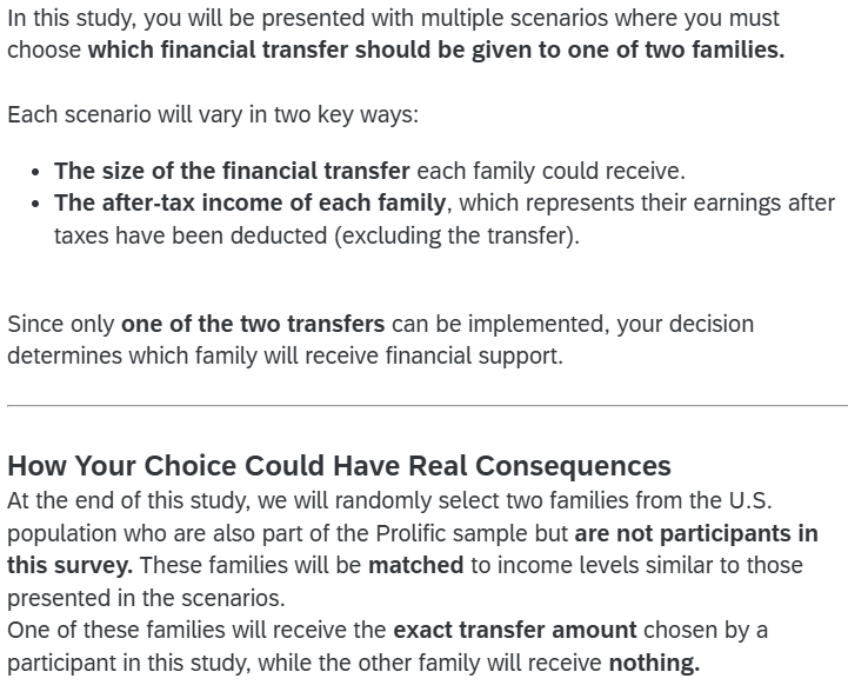
Figure 27: Estimates by $\%\Delta(1-\tau)$



Notes: This Figure plots estimates of elasticity of taxable income by the percent change in one minus the marginal tax rate.

J Experiment Screenshots

Figure 28: Experimental Instructions and Incentive Compatibility



In this study, you will be presented with multiple scenarios where you must choose **which financial transfer should be given to one of two families.**

Each scenario will vary in two key ways:

- **The size of the financial transfer** each family could receive.
- **The after-tax income of each family**, which represents their earnings after taxes have been deducted (excluding the transfer).

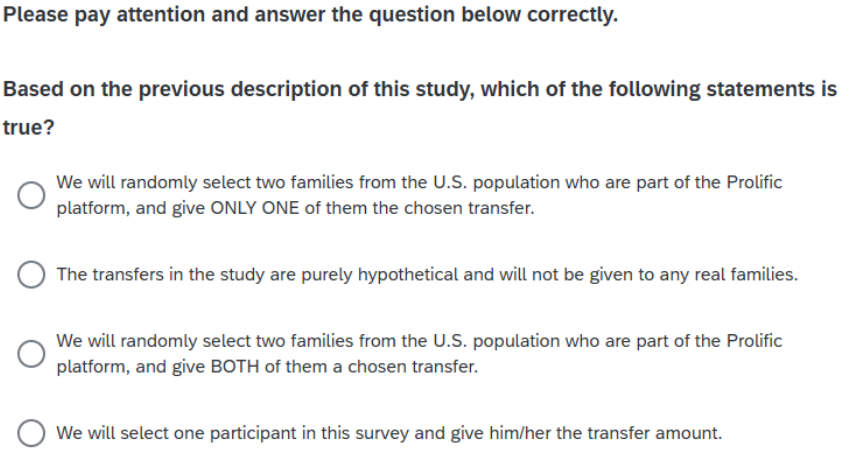
Since only **one of the two transfers** can be implemented, your decision determines which family will receive financial support.

How Your Choice Could Have Real Consequences

At the end of this study, we will randomly select two families from the U.S. population who are also part of the Prolific sample but **are not participants in this survey.** These families will be **matched** to income levels similar to those presented in the scenarios.
One of these families will receive the **exact transfer amount** chosen by a participant in this study, while the other family will receive **nothing.**

This means your decisions are not just hypothetical—your choices could directly impact how much money a real family will receive from us.

Figure 29: Comprehension Check



Please pay attention and answer the question below correctly.

Based on the previous description of this study, which of the following statements is true?

- ○ We will randomly select two families from the U.S. population who are part of the Prolific platform, and give ONLY ONE of them the chosen transfer.
- ○ The transfers in the study are purely hypothetical and will not be given to any real families.
- ○ We will randomly select two families from the U.S. population who are part of the Prolific platform, and give BOTH of them a chosen transfer.
- ○ We will select one participant in this survey and give him/her the transfer amount.

K Full Survey

K.1 Wave 1

We are a non-partisan group of researchers from Washington University in St. Louis, USA, and Bocconi University in Milan, Italy.

Our goal is to understand the diversity of views on the desirability of income redistribution. This question has important implications for government policy, and your participation in this survey contributes to our knowledge as a society. It is important that you **answer all questions honestly**. Anytime you do not know an answer, it is ok to give your best guess. It is also very important for the success of our research project that you **complete the survey until the end**, once you have started. The survey should take (on average) about **8 minutes** to complete.

YOU MUST BE A RESIDENT OF THE **UNITED STATES OF AMERICA** TO PARTICIPATE IN THIS SURVEY

If you have any questions about this study, you may contact us at redistributionpreferences@wustl.edu.

*Please enter your unique Prolific ID.

Next page ›

*The fruit test you are about to take part in is very simple. When asked to enter a fruit, you must select "Apple". This is an attention check.

Based on the text you read above, which fruit have you been asked to enter?

- ◯ Cherry
- ◯ Orange
- ◯ Banana
- ◯ Apple
- ◯ Mango

Next page ›

Powered by Qualtrics

Overview

In this study, you will be presented with multiple scenarios where you must choose **which financial transfer should be given to one of two families.**

Each scenario will vary in two key ways:

- **The size of the financial transfer** each family could receive.
- **The after-tax income of each family**, which represents their earnings after taxes have been deducted (excluding the transfer).

Since only **one of the two transfers** can be implemented, your decision determines which family will receive financial support.

How Your Choice Could Have Real Consequences

At the end of this study, we will randomly select two families from the U.S. population who are also part of the Prolific sample but **are not participants in this survey.** These families will be **matched** to income levels similar to those presented in the scenarios.
One of these families will receive the **exact transfer amount** chosen by a participant in this study, while the other family will receive **nothing.**

This means your decisions are not just hypothetical—your choices could directly impact how much money a real family will receive from us.

Next page >

Powered by Qualtrics

*Please pay attention and answer the question below correctly.

Based on the previous description of this study, which of the following statements is true?

- ◯ We will randomly select two families from the U.S. population who are part of the Prolific platform, and give BOTH of them a chosen transfer.
- ◯ We will randomly select two families from the U.S. population who are part of the Prolific platform, and give ONLY ONE of them the chosen transfer.
- ◯ We will select one participant in this survey and give him/her the transfer amount.
- ◯ The transfers in the study are purely hypothetical and will not be given to any real families.

Next page ›

Powered by Qualtrics

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1700	$1600
Current After-Tax Income	$400,000/year	$80,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1700	$700
Current After-Tax Income	$320,000/year	$265,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$500	$900

Current After-Tax Income	$75,000/year	$310,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1700	$1800
Current After-Tax Income	$175,000/year	$400,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1000	$200
Current After-Tax Income	$190,000/year	$210,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$300	$1300
Current After-Tax Income	$65,000/year	$10,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1500	$1000
Current After-Tax Income	$140,000/year	$75,000/year

Based on the information above, which of the two transfers do you prefer us to make?

○ Transfer 1 ○ Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1500	$400
Current After-Tax Income	$145,000/year	$370,000/year

Based on the information above, which of the two transfers do you prefer us to make?

Transfer 1 Transfer 2

*Consider the following of the two transfers.

	Transfer 1	Transfer 2
Transfer Amount	$500	$1700
Current After-Tax Income	$230,000/year	$60,000/year

Based on the information above, which of the two transfers do you prefer us to make?

Transfer 1 Transfer 2

*Consider the following transfers.

	Transfer 1	Transfer 2
Transfer Amount	$1600	$500
Current After-Tax Income	$125,000/year	$105,000/year

Based on the information above, which of the two transfers do you prefer us to make?

Transfer 1 Transfer 2

Next page >

Powered by Qualtrics

*Which factors were most important in guiding your previous transfer decisions? Please respond in full sentences.

Next page

Powered by Qualtrics

*The chocolate test you are about to take part in is very simple. When asked to enter a type of chocolate, you must select "Bittersweet Chocolate". This is an attention check.

Based on the text you read above, which type of chocolate have you been asked to enter?

- ◯ White Chocolate
- ◯ Dark Chocolate
- ◯ Milk Chocolate
- ◯ Bittersweet Chocolate

Next page ›

Powered by Qualtrics

What is your gender?

- ◯ Male
- ◯ Female
- ◯ Non-binary / third gender
- ◯ Prefer not to say

What is your age?

Please indicate your marital status

- ◯ Single
- ◯ Married

Do you have children living with you?

- ◯ Yes
- ◯ No

How would you describe your ethnicity/race?

- ◯ European American/White

African American/Black

Hispanic/Latino

Asian/Asian American

Other

Were you born in the United States?

Yes

No

In which state do you live?

Select one

Next page

Powered by Qualtrics

Which category best describes your highest level of education?

- ◯ Eighth Grade or Less
- ◯ Some High School
- ◯ High School Degree/GED
- ◯ Some College
- ◯ 2-year College Degree
- ◯ 4-year College Degree
- ◯ Master's Degree
- ◯ Doctoral Degree
- ◯ Professional Degree (JD, MD, MBA)

What is your current employment status?

- ◯ Full-time Employee
- ◯ Part-time Employee
- ◯ Self-employed or Small Business Owner
- ◯ Unemployed and Looking for Work
- ◯ Student
- ◯ Not in Labor Force (For Example: Retired, or Full-time Parent)

What was your TOTAL household income, before taxes, last year (2024)?

- ◯ <$10,000/year
- ◯ $10,000-$14,999/year
- ◯ $15,000-19,999/year
- ◯ $20,000-$29,999/year
- ◯ $30,000-$39,999/year
- ◯ $40,000-$49,999/year
- ◯ $50,000-$74,999/year
- ◯ 75,000-$99,999/year
- ◯ $100,000-$149,999/year
- ◯ $150,000-$199,999/year
- ◯ $200,000+/year

On economic policy matters, where do you see yourself on the liberal/conservative spectrum?

- ◯ Very Conservative
- ◯ Conservative
- ◯ Moderate
- ◯ Liberal

Very Liberal

Next page

Powered by Qualtrics

K.2 Belief Elicitation in Wave 2

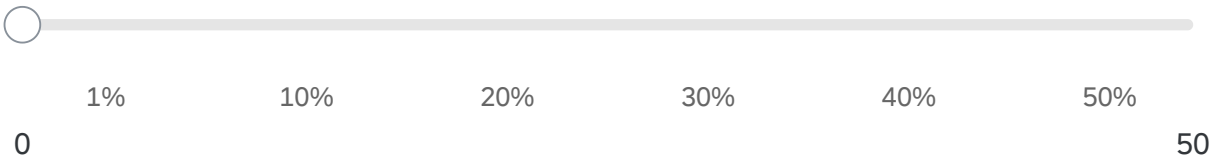
*Your estimate

What do you think happens to the income of an average US citizen when income taxes change?

Examples:
• Enter 0 if you think people don't change their income at all
• Enter 1 if you think people reduce their income by 1%
• Enter 5 if you think people reduce their income by 5%
• Enter 10 if you think people reduce their income by 10%

If the tax rate on the person's last dollar earned increases by **10%**, by what percentage do you think the person reduces the pre-tax income?

% change in pre-tax income



Next page ›

*Your estimate

How much more do you think the average American values helping lower-income families compared to higher-income families?

Specifically, compared to a dollar transferred to a family earning $80,000 per year, how many times more do you think the average American values a dollar transferred to a family earning $20,000 per year?

Examples:
- Enter 1 if you think the average American values both transfers equally
- Enter 2 if you think the average American values the transfer to the $20,000 family twice as much
- Enter 10 if you think the average American values the transfer to the $20,000 family 10 times as much

Please enter a number:

Next page ›

Powered by Qualtrics

K.3 Transfer Choices with Government Taxation in Wave 3

Overview

In this study, you will be presented with multiple scenarios where you must choose which financial transfer should be given to one of two families. In each scenario you will choose which of two families should receive a financial transfer. Only one option can be implemented per scenario.

What you will see for each option

- **After-tax household income (before your transfer decision):** the family's current annual income *after regular taxes*, ignoring the transfer on this screen.
- **Total transfer amount:** the dollar amount our study will send to the family.
- **Proportion of transfer taxed:** the share of that transfer that will be **taxed away**.
- **Transfer received:** the dollars that actually reach the family after tax: received = transfer × (1 − tax rate).

Example: If the transfer is $800 and the tax rate is 30%, the family receives $560 and $240 is taxed away. The other family receives nothing in that scenario.

Real-world consequence

At the end of the study, we will randomly select two U.S. families (from the Prolific pool, not survey participants) whose incomes match those shown in the scenarios. One family will receive the **Transfer received** amount from a randomly chosen participant's decision. **The money that is taxed-away will be donated by us to the U.S. government via the U.S. Treasury's Pay.gov portal** (https://fiscal.treasury.gov/pay.gov/).

Next page ›

*

	Option 1	Option 2
After-tax household income (before this decision)	$95,000/year	$100,000/year
Total Transfer Amount	$2400	$3600
Proportion of Transfer Taxed	**0%**	**50%**
Transfer Received	$2400	$1800

Based on the information above, which option do you prefer us to implement?

○ Transfer 1 ○ Transfer 2

*

	Option 1	Option 2
After-tax household income (before this decision)	$255,000/year	$35,000/year
Total Transfer Amount	$800	$3600
Proportion of Transfer Taxed	**50%**	**0%**
Transfer Received	$400	$3600

Based on the information above, which option do you prefer us to implement?

Transfer 1 Transfer 2

*

	Option 1	**Option 2**
After-tax household income (before this decision)	$75,000/year	$175,000/year
Total Transfer Amount	$2400	$2800
Proportion of Transfer Taxed	**0%**	**25%**
Transfer Received	$2400	$2100

Based on the information above, which option do you prefer us to implement?

Transfer 1 Transfer 2

*

	Option 1	**Option 2**
After-tax household income (before this decision)	$230,000/year	$445,000/year
Total Transfer Amount	$1200	$2400
Proportion of Transfer Taxed	**75%**	**0%**
Transfer Received	$300	$2400

Based on the information above, which option do you prefer us to implement?

Transfer 1 Transfer 2

*

	Option 1	**Option 2**
After-tax household income (before this decision)	$365,000/year	$140,000/year
Total Transfer Amount	$1200	$2800
Proportion of Transfer Taxed	**75%**	**0%**
Transfer Received	$300	$2800

Based on the information above, which option do you prefer us to implement?

○ ○

Transfer 1 Transfer 2

*

	Option 1	**Option 2**
After-tax household income (before this decision)	$185,000/year	$40,000/year
Total Transfer Amount	$2400	$3600
Proportion of Transfer Taxed	**75%**	**0%**
Transfer Received	$600	$3600

Based on the information above, which option do you prefer us to implement?

○ ○

Transfer 1 Transfer 2

*

	Option 1	Option 2
After-tax household income (before this decision)	$40,000/year	$335,000/year
Total Transfer Amount	$1200	$3600
Proportion of Transfer Taxed	**0%**	**25%**
Transfer Received	$1200	$2700

Based on the information above, which option do you prefer us to implement?

○ Transfer 1 ○ Transfer 2

*

	Option 1	Option 2
After-tax household income (before this decision)	$340,000/year	$480,000/year
Total Transfer Amount	$400	$1200
Proportion of Transfer Taxed	**25%**	**50%**
Transfer Received	$300	$600

Based on the information above, which option do you prefer us to implement?

○ Transfer 1 ○ Transfer 2

*

	Option 1	Option 2
After-tax household income (before this decision)	$450,000/year	$100,000/year
Total Transfer Amount	$2800	$3200
Proportion of Transfer Taxed	**0%**	**25%**
Transfer Received	$2800	$2400

Based on the information above, which option do you prefer us to implement?

○ Transfer 1 ○ Transfer 2

*

	Option 1	Option 2
After-tax household income (before this decision)	$290,000/year	$20,000/year
Total Transfer Amount	$3600	$3200
Proportion of Transfer Taxed	**0%**	**75%**
Transfer Received	$3600	$800

Based on the information above, which option do you prefer us to implement?

○ Transfer 1 ○ Transfer 2

*Imagine you could decide how much the U.S. government pays in transfers and charges in taxes to people with different annual incomes. For each income level below, choose the household's **average tax rate**, which is the share of income paid in net taxes (taxes minus transfers). Negative values mean a net cash benefit.

For example, if the average tax rate of a household making $20,000/year is -50%, then the government transfers $10,000 to that household, so after taxes and transfers, they take home $30,000/year. If the average tax rate of the household making $20,000/year is 50%, then the household owes the government $10,000/year in taxes, and they will take home $10,000/year.

Please choose the tax rates that you would like the U.S. government to implement.

Average tax rate in % (negative values mean transfers)

$20,000

-50 100

$40,000

-50 100

$60,000

-50 100

$100,000

-50 100

$200,000

-50 100

$500,000

-50 100

$1 million

-50 100

Next page

Powered by Qualtrics

K.4 Elicitation of Government Attitudes in Wave 4

***Now we would like to know your actual beliefs about the U.S. government.**

What do you personally believe about how the government actually operates.

- ◯ 1 - Very inefficiently - Most tax revenue is wasted
- ◯ 2 - Inefficiently
- ◯ 3 - Somewhat inefficiently
- ◯ 4 - Neither efficient nor inefficient
- ◯ 5 - Somewhat efficiently
- ◯ 6 - Efficiently
- ◯ 7 - Very efficiently - Tax revenue is used very effectively

*Out of every dollar the federal government collects in taxes, approximately how many percent do you believe are wasted (due to excessive bureaucracy, inefficiency, corruption, or misallocation)?

% of tax revenue wasted

0 100

*How much do you trust the federal government to do what is right?

- ◯ Not at all
- ◯ A little
- ◯ A moderate amount
- ◯ A lot
- ◯ A great deal

*If you wanted to help low-income families, which method would you prefer?

- ◯ Direct private charity (e.g., donating to a family directly)
- ◯ Non-profit organizations (e.g., United Way, food banks)
- ◯ Government redistribution through taxes and transfers
- ◯ All methods are equally good
- ◯ I would not want to help low-income families through any method

*When the government spends tax revenue, how much of it do you believe goes toward programs you consider valuable?

- ◯ None at all
- ◯ A little
- ◯ A moderate amount
- ◯ A lot
- ◯ A great deal

*To what extent do you agree with the following statement: "The government has the right to take money from those who earn more to give to those who earn less."

- ◯ 1 - Strongly disagree
- ◯ 2 - Disagree
- ◯ 3 - Somewhat disagree
- ◯ 4 - Neither agree nor disagree
- ◯ 5 - Somewhat agree
- ◯ 6 - Agree
- ◯ 7 - Strongly agree

Next page >