

Mirrlees Vs The Masses: Optimal Income Redistribution Under a Popular Opinion Constraint

Gregory Sun*

October 23, 2026

[Click here for the latest version](#)

Abstract

Economics is an important input into public policy, but its influence is often constrained by political feasibility concerns that derive from the views of members of the general public. I study this constraint and measure its importance by combining text data from naturally occurring online policy debates with an online survey experiment to quantify how the general public evaluates the desirability of different income redistribution policy reforms. Public opinion is diverse and not well described by the standard optimal income taxation model used in economics. Combining estimated policy preferences with structural estimates from a standard income taxation model, I find that the social planner's preferred tax reform is deeply unpopular. However, modest tweaks can achieve majority support while preserving most of the welfare gains.

*I thank Brent Hickman, Bart Hamilton, Brad Larsen, Ryan McDevitt, Martin García-Vásquez, and Adriana Lleras-Muney for helpful comments that have improved this paper.

1 Introduction

Economics provides a set of tools to identify and quantify the key tradeoffs involved when designing public policy. Yet economic analysis often has limited influence in practice, partly because the reforms economists recommend can fail to garner popular support. For example, broad tariffs have regained political traction in the US despite little support among economists, and rent control remains popular despite evidence that it is counterproductive (Diamond et al., 2019). Such disagreements reflect deeper differences in how economists and the public reason (Autor et al., 2025; Elmendorf et al., 2025). This paper studies what these disconnects imply for the advice economists give within a fundamental area of public policy: taxation and redistribution.

The seminal model of Mirrlees (1971) has long established a standard framework within which economists debate the costs and benefits of prioritizing different tax-and-transfer policy imperatives (see survey by Piketty and Saez, 2013). Within this framework, heterogeneous individuals choose endogenous labor supply given the tax-and-transfer schedule, which the social planner designs to maximize social welfare, subject to a government budget constraint. Optimal marginal tax rates then depend on the social marginal welfare weights attached to individuals at different incomes, the shape of the income distribution, and the elasticity of taxable income with respect to taxes. Applied work estimates these objects and computes the policies they imply, but economists’ model-based recommendations often diverge sharply from policies the public supports. For instance, even the low end of optimal top marginal tax rate estimates is over 10 percentage points higher than the rates preferred by a typical liberal member of the public (Rodemeier and Sun, 2026; Capozza and Srinivasan, 2024). Because reforms lacking broad public support are unlikely to be politically feasible, the optimal policies economists derive may have little practical relevance.

To study the tension between economic analysis and political feasibility, I estimate how the optimal policy of the Mirrleesian planner changes when she must also satisfy a popular support constraint, which requires that a sufficiently large share of the public prefers any

proposed reform to the status quo. Solving this problem requires two empirical inputs. First, the distribution of the public’s preferences over tax-and-transfer policy determines how people evaluate the consequences of a reform. Second, the structural primitives of the Mirrlees model determine what those consequences are, as earnings endogenously respond to tax policies that shape the marginal price of labor. With both in hand, I ask (i) the extent to which the popular support constraint binds, (ii) how the planner should modify her policies to satisfy it, and (iii) what the social welfare costs of doing so are.

Estimating the first input requires specifying the set of considerations a model of public policy preferences should include, but existing evidence offers limited guidance. A large literature documents that fairness concerns shape attitudes toward redistribution (Konow, 2000; Alesina and Angeletos, 2005; Fleurbaey and Maniquet, 2006; Cappelen et al., 2007, 2013, 2018; Almås et al., 2020; Hvidberg et al., 2023), but also that the way in which these concerns translate into behavior depends on subtle details of the institutional context (Fehr and Schmidt, 1999; List, 2007; Fershtman et al., 2012). Saez and Stantcheva (2016) provide a quantitative framework to measure how fairness concerns translate into citizens’ policy preferences. This approach requires the researcher to pre-specify a set of policy-relevant tradeoffs, but, as previous studies suggest, fairness notions most relevant to the broader public are difficult to anticipate. To overcome this challenge, I develop a novel approach that allows contemporary public discourse data to determine which considerations to model.

I first gather a large corpus of text data from online policy discussions about proposals to reform redistribution policy in the US. I then apply topic modeling methods, in the form of a Gaussian mixture model estimated on neural network embeddings, in order to empirically identify the most salient tradeoffs that the public raises when evaluating redistribution policy. Two findings emerge. First, the considerations people raise are diverse, and many of them lie outside the standard Mirrleesian framework, such as whether transfers are made in cash or in kind, their fiscal cost, and concerns about automation. Second, even considerations that economists share with the public are often understood differently. While economists care

about reductions in work only because they make taxation and redistribution less efficient, many commenters articulate direct normative concerns about these responses. In summary, I find that the public differs from economists both in how it evaluates familiar policy tradeoffs, and also in the broader range of considerations it views as relevant to redistribution policy.

To translate these findings into actionable policy preference estimates, I adopt the generalized social marginal welfare weight framework of Saez and Stantcheva (2016), in which each citizen evaluates a reform by weighting its effects on different households according to their characteristics. The topics identified in the empirical text analysis determine which household characteristics enter citizens' welfare weight functions. I quantify the importance of each consideration using a survey experiment in which respondents choose between pairs of hypothetical government transfers that vary along relevant dimensions identified within the text analysis. I implement this survey design both on a sample of US respondents and on a separate sample of European respondents.

Within the US sample, respondents exhibit strong preferences for redistribution toward poor households, but they also are highly averse to negative labor-supply responses induced by transfers. These preferences are shared by American liberals and conservatives alike, though weighted differently across political identities. Respondents also generally dislike the fiscal costs of redistribution, but preferences on this and other dimensions vary substantially across respondents. Results from the European sample are similar in many ways, with a key difference being that only conservative Europeans show any normative concern for labor-supply responses, with the conservative European aversion being roughly at the midpoint between American conservatives and liberals. I estimate the corresponding distribution of welfare weights separately for the European and US data using a random coefficients discrete choice model. Because these welfare weights determine how each citizen evaluates any reform, the estimates provide a quantitative characterization of political feasibility constraints.

To determine the social planner's objectives, I also need to predict how households respond to changes in taxes and transfers, which is governed by the structural primitives of the

Mirrlees model. Standard empirical specifications assume away income effects, ruling out, by construction, the widely held concern that even lump-sum transfers discourage work. For my present purposes, this standard restriction is problematic for the goal of estimating counterfactuals where both levels and slopes of income change. I therefore extend the standard Mirrleesian empirical specification by allowing both substitution *and* income effects to vary flexibly across the income distribution. Moreover, I am able to nonparametrically identify the structural primitives, including household utility curvature and the distribution of labor productivity. This is accomplished by matching elasticities of taxable income and marginal propensities to earn across the income distribution, using a meta-analysis of recent experimental and quasi-experimental data.

Combining these inputs, I study how the popular support constraint shapes optimal redistribution policy. A standard Mirrleesian planner raises taxes on high earners (relative to the status quo) to fund relatively generous transfers to low-income households, withdrawing these transfers rapidly as earnings rise in order to target transfers at the poorest households. This rapid withdrawal raises marginal tax rates at the bottom of the income distribution and induces substantial reductions in work. These consequences of the Mirrleesian planner's policy clash with popular preference estimates because the public's aversion to labor-supply reductions outweighs its taste for the policy's redistributive effects. Empirically, only about a third of the American public supports the reform, implying that a political feasibility constraint requiring majority support would bind. Relaxing the government's budget constraint to fund more generous transfers does not help, because the public also dislikes government spending. Interestingly, due to Europeans' lower normative aversion to transfer-induced labor responses, the majority-support constraint would not bind within the European sample.

In a second set of counterfactuals, I solve for a set of alternative redistribution policies that take the preferences of the American public into account. These policies maximize Mirrleesian social welfare, subject to an additional constraint that a sufficiently large fraction of the US public supports the reform. This implies a continuum of potential policy

recommendations, but I report two in particular: one with a majority support constraint, which provides a bound on the set of politically feasible policies, and another requiring super-majority support. Relative to the Mirrleesian optimum, the constrained policies give smaller transfers at the bottom of the income distribution but phase them out more gradually. This lowers marginal tax rates at the bottom while still eventually phasing out transfers, so that the reform remains budget neutral, but induces smaller labor-supply responses among transfer recipients. Modifications of this kind achieve majority support while retaining over 90% of the welfare gains from the Mirrleesian planner’s preferred reform, so satisfying the empirical popular support constraint is relatively inexpensive from the perspective of traditional economic analysis. These results survive a battery of robustness checks.

I additionally validate my findings directly in a second, pre-registered survey experiment. I directly represent the tradeoffs implied by the Mirrleesian and constrained optimal policies in simple terms to respondents. I include interactive elements to help respondents visualize policy implications of the proposed reforms, including how they impact tax burdens across the income distribution and labor-supply responses by transfer recipients. I then ask respondents whether they would prefer each reform to the status quo. As my model estimates predict, the baseline Mirrleesian planner’s preferred reform falls short of majority support among American respondents. On the other hand, the simple-majority constrained optimal reform attains majority support of 57.5% (95% confidence interval of [54.1%, 60.9%]). This validation exercise provides independent evidence that social policies favored by traditional economic analysis can be altered for greater relevance to practical policy-making.

Taken together, my findings show that satisfying the popular support constraint requires only modest departures from the planner’s preferred policy, provided that those departures are guided by what the public cares about. Identifying these considerations is possible here because the public’s own views about redistribution are estimable within online text data. Analogous text data may be available in other settings where economists and the public disagree, including the trade and housing policy examples discussed above. Aside

from contributing empirical insights to the tax-and-transfer debate, this paper proposes a portable methodology to combine text analysis, survey measurement, and structural policy design to study political feasibility constraints within broader economic policy settings.

The remainder of the paper is organized as follows. Section 2 develops the framework for describing the public’s policy preferences and estimates them. Section 3 estimates the structural primitives of the Mirrlees model. Section 4 combines these estimates to characterize optimal policy under the popular support constraint. Section 5 concludes.

1.1 Related Literature

This paper relates to a long tradition in political economics that incorporates political constraints into models of income taxation (Roberts, 1977; Meltzer and Richard, 1981; Snyder and Kramer, 1988; Roemer, 1999; Benabou and Ok, 2001; Acemoglu et al., 2010; Bierbrauer and Boyer, 2016; Berliant and Gouveia, 2025). Much of this literature characterizes the political equilibrium that arises under specific models of voter preferences. I complement this work by studying a more local question: which small reforms to the status quo are politically feasible. Focusing on local reforms allows voter preferences to be specified flexibly and estimated empirically, an approach well suited to the heterogeneous preferences I document.

A large body of work documents systematic differences in how economists and the public view redistribution. The public often holds inaccurate beliefs about the extent of inequality (Norton and Ariely, 2011; Cruces et al., 2013; Kuziemko et al., 2015), evaluates tax-and-transfer policy using normative criteria that differ from those of the standard framework (Kuziemko et al., 2015; Saez and Stantcheva, 2016; Fisman et al., 2020; Stantcheva, 2021), and reaches different conclusions about optimal policy (Rodemeier and Sun, 2026; Capozza and Srinivasan, 2024). I contribute to this literature in two ways. First, I use text data to determine empirically which of the many documented considerations are most relevant in debates over redistribution. Second, I incorporate the resulting estimates into an explicit policy design problem, translating documented differences in views into their implications

for optimal policy.

Most closely, this paper contributes to an emerging literature on policy design under political feasibility constraints. Dreyfuss (2025) and Jiang et al. (2026) study settings in which the public’s *positive* beliefs about how the economy works differ from those of economists, while Shapiro et al. (2026) measure demand for firearm policy among those directly affected by it. I add to this discussion by focusing on *normative* differences between economists and the public, measured among third-party observers of redistribution.

Finally, the paper builds on evidence on how households respond to taxes and transfers, including the meta-analysis of elasticities of taxable income in Rodemeier and Sun (2026) and experimental and quasi-experimental estimates of the labor-supply response to unearned income (Imbens et al., 2001; Cesarini et al., 2017; Golosov et al., 2024; Vivalt et al., 2024). Relative to this work and to a growing literature on the broader effects of unconditional cash transfers (Bartik et al., 2024; Miller et al., 2024; Jaroszewicz et al., 2024), I synthesize existing estimates into the structural primitives of an optimal taxation model, which allows them to inform the design of politically feasible policy.

2 How Do Non-Economists Evaluate Policies?

The goal of this section is to quantify the most important dimensions along which economists and the public reason differently in thinking about redistribution. My point of departure is the most common framework in the economics literature for thinking about redistributive taxation: the optimal income taxation model of Mirrlees (1971). I therefore first briefly review key concepts from the standard theory.

Consider a population of households characterized by an unobservably heterogeneous type $\theta \sim F_\theta$. A type θ household has preferences over consumption c and pre-tax income z , which are described by a utility function $U(c, z; \theta)$.¹ The government can set a tax-and-

¹This differs slightly from the original setup of Mirrlees (1971), where preferences over consumption and labor supply ℓ are homogeneous and described by utility $U(c, \ell)$, and θ represents unobserved *productivity*

transfer policy $T(z)$, which maps a given pre-tax income level z into a net tax burden. Negative values $T(z) < 0$ imply that a household with income z receives assistance from the government on net. Denote $C(z) \equiv z - T(z)$ as the post-tax income enjoyed by someone who generates pre-tax income z . Given the tax schedule, households choose pre-tax income to maximize utility: $z(T, \theta) \equiv \arg\max_z U(z - T(z), z, \theta)$. The indirect utility function for a type- θ household facing tax schedule T is denoted $V(T, \theta) \equiv \max_z U(z - T(z), z, \theta)$, and the social planner's preferences over redistributive outcomes are described by social welfare function

$$\int_{\theta} \mathcal{W}(V(T, \theta), \theta) dF_{\theta}(\theta). \quad (1)$$

This framework is flexible enough to encode fairly general redistributive concerns, through the specification of $\mathcal{W}$, which quantifies how much weight the planner places on different members of society. A quantity which plays an important role for analysis is the marginal welfare weight

$$\Gamma_T(z) \equiv \frac{\mathbb{E}[\mathcal{W}_1(V(T, \theta), \theta) \cdot U_1(z(T, \theta) - T(z(T, \theta)), z(T, \theta), \theta) | z(T, \theta) = z]}{\mathbb{E}[\mathcal{W}_1(V(T, \theta), \theta) \cdot U_1(z(T, \theta) - T(z(T, \theta)), z(T, \theta), \theta)]},$$

where here and in the rest of the exposition, I use subscripts under the functions $\mathcal{W}$ and U to denote derivatives with respect to the argument indexed by the subscript. The numerator of the above expression denotes the marginal utility of giving an extra dollar of income to a randomly chosen household that, under the current tax system T , earns income z . The denominator is a normalizing factor that ensures that welfare weights integrate to 1. The social planner values redistribution if she places more weight on the poor versus the rich, which is summarized by the extent to which $\Gamma'_T(z)$ is negative.

The planner wishes to maximize social welfare subject to two constraints. First, given a choice of tax policy T , households endogenously respond and choose their income level

heterogeneity. This induces a set of preferences over consumption and pre-tax income via the identity, $U(c, z, \theta) = U(c, z/\theta)$. Most empirical applications work with data where pre-tax income is the observed choice, so formulating preferences in terms of income here is more convenient.

$z(T, \theta)$, which implicitly encodes an *incentive constraint*. Second, the government wishes to raise a sufficient level of revenue $\underline{R}$ from the taxes that it sets. Combining these constraints with the social welfare function from Equation (1), the social planner's problem is

$$\max_T \int_{\theta} \mathcal{W}(V(T, \theta), \theta) dF_{\theta}(\theta) \quad \text{s.t.} \quad \int_{\theta} T(z(T, \theta)) dF_{\theta}(\theta) \geq \underline{R}. \quad (2)$$

While the resulting optimization problem is complicated, the conditions for optimality that arise from it are fairly parsimonious. To describe them, it is necessary to introduce a number of (endogenous) quantities associated with a given status quo level of taxes T . Let $G_T(z) \equiv \Pr[z(T, \theta) \leq z]$ denote the distribution of pre-tax incomes induced by a given tax policy T , and let g_T denote its density. For a given $G_T(z)$, the “local Pareto” parameter, $\alpha_T(z) \equiv \frac{g_T(z)z}{1-G_T(z)}$, measures how dispersed the income distribution is around a given level of income z . For a type- θ household, denote the Hicks elasticity of taxable income by $\zeta_T(\theta) \equiv \frac{\partial z(T, \theta)}{\partial (1-T'(z(T, \theta)))} \frac{1-T'(z)}{z(T, \theta)}$ and define the income effect as $\eta_T(\theta) \equiv \frac{\partial z(T, \theta)}{\partial T(z(T, \theta))} (1 - T'(z))$.² Define the average Hicks elasticity and income effect respectively by $\zeta_T(z) \equiv \mathbb{E}[\zeta_T(\theta) | z(T, \theta) = z]$ and $\eta_T(z) \equiv \mathbb{E}[\eta_T(\theta) | z(T, \theta) = z]$.

Under a broad set of circumstances (see, for instance, Ferey et al., 2024), optimal marginal tax rates are characterized by the formula

$$\frac{T'(z)}{1 - T'(z)} = \mathbb{E} \left[1 - \Gamma_T(z(T, \theta)) + \frac{T'(z(T, \theta))}{1 - T'(z(T, \theta))} \eta_T(z(T, \theta)) \middle| z(T, \theta) \geq z \right] \times \frac{1}{\alpha_T(z)} \times \frac{1}{\zeta_T(z)}. \quad (3)$$

This formula roughly decomposes the social planner's first-order condition as depending on only three factors: *redistributive preference* of the social planner $\Gamma_T(z)$, *income inequality* in society summarized locally by $\alpha_T(z)$, and *behavioral responses* to tax policy described by $\zeta_T(z)$ and $\eta_T(z)$.

²In the definition of the Hicks elasticity, the derivative with respect to $1 - T'(z(T, \theta))$ should be understood to mean with respect to a small tax reform that changes the marginal tax rate at income level $z(T, \theta)$ while keeping the overall tax burden $T(z(T, \theta))$ fixed. Conversely, the derivative with respect to $T(z(T, \theta))$ should be understood as holding the marginal tax rate $T'(z(T, \theta))$ fixed.

The structure of Equation (3) thus suggests a straightforward empirical test of whether people who are discussing redistribution policy reason similarly to the economists who uses this standard framework. If the general public’s conversation topics largely center around either (i) how much to prioritize helping the poor over the rich, (ii) the level of inequality in society, or (iii) behavioral responses to taxation, then we cannot rule out that they reason like economists. Conversely, if other concerns dominate, then the standard economic framework does not capture the the general public’s most salient concerns. Section 2.1 formalizes and implements the above test using text analysis tools. Qualitatively, people bring up many concerns outside of the three mentioned here, so in Section 2.2, I implement and analyze a survey experiment to estimate the quantitative tradeoff between the most salient concerns identified in the text analysis.

2.1 Qualitative Analysis: Topic Modeling

This subsection asks which topics naturally arise when members of the general public debate redistribution. To answer this question, I use standard methods from the topic modeling literature in machine learning (Vayansky and Kumar, 2020). These methods are used organize large collection of text data into *topics*, such that two pieces of text that belong to the same topic are semantically similar to one another while two pieces of text that belong to different topics are dissimilar from one other.

I use data from the social media platform Reddit.com, which is organized into discussion boards (“subreddits”) devoted to common themes. Within a subreddit, an original poster (OP) starts a thread on a given topic, and other users respond with comments. A comment can reply either directly to the OP, in which case I refer to it as a *top-level* comment, or to another comment. While Reddit users are not a representative sample of the population, they comprise a fairly broad swath of it. Moreover, because the goal of this section is to identify *which* topics arise, a fully representative sample is not essential, provided that the most common topics on Reddit coincide with the most common topics in the broader

population.

To identify a large set of discussions that are clearly about redistribution policy, I focus on online debates surrounding universal basic income (UBI). Debates during the 2020 presidential election popularized UBI as a potential way to reform redistribution policy in the US, and concerns about AI and automation have since sustained public interest in the idea. The nature of the resulting debates is particularly convenient for my empirical analysis for two reasons. First, much of the discussion surrounding UBI has occurred on internet forums, which makes it visible to the econometrician. Second, UBI serves as a useful focal point and provides a convenient keyword for finding discussions which are clearly about redistribution policy. Together, these features make the UBI debate a convenient window into the considerations that ordinary people bring to bear when evaluating redistribution policy more generally, which are precisely the inputs needed to characterize the public support constraint.

I construct my main analysis sample by collecting every top-level comment on threads whose title contains “universal basic income”, “UBI”, or “basic income” in three popular subreddits, “Politics”, “NoStupidQuestions”, and “AskReddit”, as of July 31, 2025.³ “Politics” is among the most frequented subreddits on the site and focuses on political topics, while the other two are popular question-and-answer subreddits in which threads are often structured as prompts for debate. The resulting sample contains 528 threads with 11,183 top-level comments. For each comment, I also record its number of “likes” as a measure of how well received it is.

Before analyzing these comments, I apply two preprocessing steps. First, not all comments are relevant for understanding judgments about the desirability of UBI as a policy. For example, one thread asked users how they would spend a hypothetical \$1,000 monthly UBI transfer,⁴ and the responses describe what commenters would do with the money, which is uninformative about the merits of UBI from a policymaker’s point of view. To filter such

³The search does not restrict posting dates, so all qualifying threads posted before this date are included.

⁴This amount was proposed by Andrew Yang in his 2020 presidential campaign.

comments at scale, I use the Anthropic API (Anthropic, 2025) to prompt a large language model (the March 7, 2024 version of Claude 3 Haiku) to classify whether each comment makes an argument about the desirability of UBI as a policy. Second, to prevent stylistic and tonal differences across commenters from driving the analysis, I prompt the same model to summarize the main substantive argument in each relevant comment. The resulting dataset consists of 4,835 relevant comments and their summaries. Details of both steps are given in Appendix A.

Clustering comments into topics requires first overcoming the fact that, unlike the numerical variables typically used in economic analysis, raw text is unstructured. I address this using a byproduct of the large language models (LLMs) underlying recent advances in artificial intelligence. LLMs are neural networks, which map an input, such as a piece of text, into an output through a sequence of parameterized transformations whose parameters are fit to perform a prediction task. The output of the next-to-last transformation, known as an *embedding*, is often useful for tasks other than the one the network was originally trained to perform (Buchanan et al., 2025). In the language context, embeddings are useful because they map semantically similar pieces of text to nearby vectors (Mikolov et al., 2013; Reimers and Gurevych, 2019), so that Euclidean distance between embeddings provides a meaningful measure of semantic similarity.

I compute embeddings of all 4,835 comment summaries using a popular open-source embedding model.⁵ Because these embeddings are high-dimensional, I further map them into $\mathbb{R}^3$ using Uniform Manifold Approximation and Projection (UMAP), a common nonlinear dimensionality-reduction technique (McInnes et al., 2018).⁶ The result is a set of 4,835 three-dimensional vectors in which nearby vectors correspond to semantically similar comments.

I cluster these vectors by modeling them as draws from a Gaussian mixture model with

⁵I use the “all-MiniLM-L6-v2” model from the SentenceTransformers Python library.

⁶I use UMAP instead of a more familiar linear dimension technique such as principal components analysis because neural network embeddings are typically believed to lie on a low-dimensional nonlinear manifold, not in a low-dimensional linear subspace (Goodfellow et al., 2016).

20 components.⁷ Each component of the mixture model corresponds to a distinct cluster, and comments are sorted into the cluster whose mixture component mean is closest to that comment’s embedding.⁸ I read the comments closest to the center of each cluster (reported in Table 6 in Appendix A) and assign it an intuitive label to aid interpretation. For example, Cluster 9 captures concerns about the fiscal responsibility of UBI. When several clusters discuss different aspects of the same underlying topic, I group them into a single topic (e.g., Clusters 10 and 11 both do for the relationship between UBI and existing welfare programs).

Panel A of Figure 1 reports the resulting topics and their relative frequencies. Because the prominence of a topic may depend not only on how often it is raised but also on how well received it is, Panel B relates each topic’s frequency to the average number of likes its comments receive. Likes vary relatively little across topics, and most topics do not differ significantly from one another. Accordingly, the like-weighted topic shares in Panel C closely resemble the frequency shares in Panel A.


A number of insights emerge from Figure 1. First, the considerations people raise are diverse. Commenters discuss the distinction between cash and in-kind benefits, the fiscal costs to the government, the effects of redistribution on labor supply, automation, inflation, and redistribution itself, and even the most commonly discussed topic—cash versus in-kind—accounts for only about 30% of comments.⁹ Second, many of these topics are missed by the standard Mirrleesian framework, and even considerations that the framework does include are often understood differently. Comments about labor supply, for instance, frequently express direct normative judgments about whether recipients work, as exemplified by the following comment justifying why the commenter does not support UBI: “No why should lazy people get money for doing nothing...This is beyond stupid!” This contrasts with the Mirrleesian

⁷I also experimented with 15 and 25 components. The former missed topics identified with 20 components, while the latter produced redundant components.


⁸The first two steps of this pipeline follow Grootendorst (2022). I use a Gaussian mixture for the final clustering step because it specifies an explicit data-generating process for the embeddings.

⁹These concerns also differ in focus from the recent empirical economics literature on unconditional cash transfers in developed countries. While some papers focus on the labor-market impact of UBI (Vivalt et al., 2024; Golosov et al., 2024), a primary focus of this literature is on measuring the effect of cash on various aspects of recipient well-being (Bartik et al., 2024; Miller et al., 2024; Jaroszewicz et al., 2024).


Figure 1: Topic Analysis



(a) Topic Frequency



(b) Topic Frequency and Likes



(c) Like-Weighted Topic Frequency

Notes: Panel A displays the proportion of comments corresponding to each high-level category of topic. Panel B contains a scatterplot with frequency of comments on the horizontal axis and average number of likes per comment on the vertical axis. Panel C displays the proportion of comments corresponding to each high-level category of topic, weighted by likes.

framework, where labor-supply responses matter only insofar as they make taxation less efficient. These findings raise the possibility that even if the social planner is convinced that her framework is normatively correct, she may have difficulty persuading the public to support her preferred policies unless she engages with these considerations. To translate this qualitative insight into policy implications I next quantify how the public trades off the various concerns identified in the text analysis.

2.2 Quantitative Analysis: Survey Experiment

To model how the considerations identified in the previous subsection translate into policy preferences, I use the generalized social marginal welfare weight (GSMWW) framework introduced by Saez and Stantcheva (2016). This approach summarizes a citizen’s preferences over tax-and-transfer policy via a welfare weight function over household characteristics. I connect GSMWWs to my topic analysis by assuming that people tend to talk about what is important to them, which justifies treating the topics identified above as directly informative about which characteristics to include in a specification of citizens’ welfare weight functions.

Formally, I model a population of citizens with preferences over redistribution policies. Policies differ by their effect on recipient households, which are in turn characterized by a vector of characteristics, denoted X . Citizen i ’s policy preferences can be summarized via their GSMWW function, $\alpha_i(X)$, which depends on recipient characteristics. For example, if i is primarily interested in helping the poor, then one of the dimensions of X , say X_1 , would correspond to income, and $\alpha_i(X)$ would be decreasing in X_1 . GSMWWs are a convenient way to represent preferences over redistributive policies because they have a natural interpretation in terms of a discrete choice experiment. Specifically, $\alpha_i(X)$ is the value of α such that citizen i is indifferent between giving a household with characteristics X one dollar versus giving α dollars to a randomly chosen household, drawn from the rest of society. Thus, $\alpha_i(X) > 1$ indicates that citizen i places above-average priority on households with characteristics X , while $\alpha_i(X) < 1$ indicates below-average priority.

Preferences over transfers to individual households can then be used to construct preferences over redistributive policies in general. Consider a budget-neutral reform to the existing system of redistributive taxes, and let $\Delta T(X)$ denote the (net of tax) cash transfer to a household with characteristics X induced by the reform, prior to any behavioral responses on the part of that household. In the GSMWW framework, citizen i prefers this policy if and only if $\int \alpha_i(X) \Delta T(X) dF(X) \geq 0$. Motivated by my finding that the distinction between cash and in-kind transfers is salient to many people discussing UBI, I extend this framework slightly to allow for the possibility that people place different values on the reform depending on the nature of the transfer. I also allow for reforms that are not budget neutral. Thus, let a hypothetical redistribution policy be characterized both by the cash transfer ΔT and by a change in the in-kind benefit amount ΔK . If this policy costs the government ΔC to implement over the status quo, then citizen i supports the policy if and only if

$$\int \alpha_i(X) [\Delta T(X) + \beta_i^K \Delta K(X)] dF(X) - \beta_i^C \Delta C \geq 0, \quad (4)$$

where β_i^K is the relative weight that citizen i places on in-kind benefits relative to cash, and β_i^C is the weight on fiscal concerns. To streamline notation, I adopt the convention that $\beta_{i1} \equiv \beta_i^K$ while $\beta_{i2} \equiv \beta_i^C$. A standard approach to estimating the GSMWW function is to posit a specification for $\alpha_i(X)$, ask survey respondents to choose between hypothetical transfers, and connect their choices to the specification using standard discrete choice methods.¹⁰ Relative to the standard approach, the text analysis from the previous subsection disciplines the specification of $\alpha_i(X)$ in a more data-driven way.

I connect $\alpha_i(X)$ to the text analysis by assuming that people tend to talk about topics that are important to them. To formalize this, I parameterize the welfare weight function as linear in characteristics. Specifically, after normalizing each characteristic X_k to have unit

¹⁰For instance, Saez and Stantcheva (2016) argue that allowing welfare weights to depend on both pre-tax and post-tax incomes suffices to capture both utilitarianism and libertarianism and conduct a survey experiment varying both to estimate the extent to which people exhibit behavior consistent with one of these two common ethical benchmarks.

variance, I assume

$$\alpha_i(X) = \sum_{j=3}^J \beta_{ij} X_j. \quad (5)$$

The ratio β_{ij}/β_i^C can be interpreted as quantifying how much weight i places on household characteristic j compared to sending a dollar to the government, while β_i^K captures the relative value that citizen i places on in-kind benefits compared to cash.

I formalize the assumption that citizens talk about what matters most to them as follows.

Assumption 1. In text data, $\Pr[i \text{ discusses topic } j] = \frac{|\beta_{ij}|}{\sum_{j'} |\beta_{ij'}|}$.

A consequence of this assumption is that there is a one-to-one mapping between the topics that arise in my text analysis and components of the GSMWW function $\alpha_i(X)$. In particular, the most frequent topics should be exactly those that are most crucial to understand in order to approximate citizens' preferences over policy well. I thus design a survey experiment to elicit GSMWWs, based on the topics recovered in the Reddit text analysis.

2.2.1 Survey Design

I conducted an online survey with 1,070 respondents to estimate the preferences of the public with respect to different redistributive policies. After a number of basic demographic questions, the main body of the survey consists of 10 questions in which respondents choose between income transfers to two hypothetical recipients. Transfers vary in how much the recipient receives, the efficiency loss to the government from implementing the transfer, and whether the transfer is cash or in-kind. Recipients differ in their income level, how much they change their labor supply in response to the transfer, and whether their current jobs are likely to be automated by technology such as AI. These attributes directly correspond to the main topics from the text analysis. For each question, respondents are presented with a table describing two hypothetical transfers and are asked to indicate which transfer they believe the government should make. The values of the attributes are randomized across questions, allowing me to identify how each attribute affects respondents' choices. The specific wording

of the question, with a concrete example, is given in Figure 2. Details about randomization are given in Appendix B.¹¹

Figure 2: Survey Question Example

For the following questions, consider a government who wishes to spend additional money to finance one of the following two transfers. Each differs by the **amount** and **type of transfer**, and by the **characteristics of the recipient household**. Full details are provided in the tables in each question.

Based on the information contained in the table, for each question, **mark which of the two transfers you think the government should make, if you had to pick exactly one of the two.**

*

	Transfer 1	Transfer 2
Transfer Characteristics:		
Transfer Amount	\$13,000	\$16,000
Transfer Type	Housing Assistance	Cash
Recipient Characteristics:		
Recipient's after-tax household income	\$65,000/year	\$65,000/year
The transfer causes the recipient to work ___% fewer hours/week	6%	4%
Recipient may soon lose job due to AI automation?	YES	NO
Cost to government to make transfer	\$16,000	\$26,000

☐

The government should make Transfer 1

☐

The government should make Transfer 2

Notes: This figure shows an example of a hypothetical choice question used in the survey. The contents of each row in the table are randomized independently for each question.

The survey was implemented in Qualtrics and distributed through Prolific.com, a platform commonly used by economists and other social scientists to recruit participants for online survey research. The experiment was pre-registered on the AEA RCT registry as trial number 14351. To compare American preferences with those in a setting with more extensive redistribution, I also field the same survey, with minor adaptations to the European context,

¹¹Readers who are interested in the overall flow of the survey can take a version of it for themselves at the following link: <http://bit.ly/4pksVxu>.

to a sample of 709 respondents residing in Europe (see Appendix C for details).

Table 1 summarizes the demographic composition of the US sample. Respondents are evenly split by gender and have an average age of 46. The median respondent reports a household income between \$50K and \$75K, somewhat below the 2023 US median of \$80,610 (Guzman and Kollar, 2024). The sample is somewhat more educated than the US adult population, with 62% of respondents holding at least a four-year college degree. The sample is politically balanced: 36% of respondents identify as liberal, 38% as conservative, and the remaining 26% as moderate.

2.2.2 Results

I first provide descriptive evidence about how each attribute affects choices by estimating a linear probability model of how respondents’ choices respond to the attributes varied in the survey. Let Y_{iq} denote an indicator for whether respondent i chooses transfer 1 (versus transfer 2) in question q , and let ΔX_{iq}^j denote the difference in attribute j between transfer 1 and transfer 2. The linear probability model is

$$Y_{iq} = \delta_0 + \sum_{j=1}^J \delta_j \Delta X_{iq}^j + \varepsilon_{iq}. \quad (6)$$

The first column of Table 2 reports the results for the US sample. Because each of the 1,070 respondents answers 10 questions, all coefficients are precisely estimated, and all are statistically significant with intuitive signs. Respondents prefer transfers to lower-income recipients ($\hat{\delta}_{\text{income}} = -0.126$), dislike transfers that induce large labor-supply reductions ($\hat{\delta}_{\text{labor}} = -0.333$), prefer in-kind transfers over cash ($\hat{\delta}_{\text{in-kind}} = 0.065$), favor transfers to workers at risk of AI displacement ($\hat{\delta}_{\text{AI}} = 0.029$), and dislike transfers with high fiscal costs ($\hat{\delta}_{\text{cost}} = -0.094$).

The negative sign on the labor-supply coefficients indicates that, holding all else equal, respondents prefer transfers that induce smaller reductions in labor supply. Because this

Table 1: Survey Sample Demographics

	% of sample
Female	50.0
Age (mean, years)	45.9
Married	50.0
Has children/dependents	45.0
Born in US	95.0
<i>Race/ethnicity</i>	
White	65.6
Black	12.5
Hispanic	10.4
Asian / Pacific Islander	6.6
Other	4.9
<i>Education</i>	
High school or less	11.4
Some college / 2-year degree	26.1
4-year degree	38.3
Graduate degree	24.2
<i>Household income (2023)</i>	
Less than \$50K	33.8
\$50K–\$100K	33.8
\$100K or more	32.3
<i>Employment</i>	
Working	76.4
Student	2.9
Not working	20.7
<i>Economic ideology</i>	
Liberal	36.4
Moderate	25.7
Conservative	37.9
Respondents	1,070

Notes: This table reports descriptive statistics of the respondents in the survey experiment. Entries are expressed in percentages unless noted.

experiment holds the cost to the government fixed, results on the labor-supply coefficient imply that the public places direct weight on labor-supply outcomes beyond what standard efficiency concerns would imply. This finding echoes the qualitative evidence from the topic modeling exercise and has important implications for the counterfactual analysis in Section 4. This sensitivity, however, is to a large extent specific to the US context. The second column of Table 2 reports the same regression for the European sample. European respondents

weigh recipient income, fiscal costs, in-kind provision, and AI risk much as Americans do, but the coefficient on recipient labor-supply reductions is small and statistically insignificant ($\hat{\delta}_{\text{labor}} = -0.078$, with a standard error of 0.065).

Table 2: Linear Probability Model Results

	<i>Dependent variable: Chose Family 1</i>	
	United States	Europe
Income	−0.126*** (0.004)	−0.136*** (0.005)
Cost of Transfer	−0.094*** (0.015)	−0.124*** (0.019)
Decrease in Labor	−0.333*** (0.054)	−0.078 (0.065)
Transfer In-Kind	0.065*** (0.006)	0.083*** (0.008)
AI Risk	0.029*** (0.006)	0.034*** (0.008)
Observations	10,700	7,090

*p<0.1; **p<0.05; ***p<0.01

Notes: This table contains regression coefficients for a linear probability model regression, where the outcome is an indicator for choosing the first family. The regressors are constructed by taking the quantity for family 1 and subtracting the analogous quantity for family 2. The first column uses the US sample, and the second column uses the European sample. Standard errors are in parentheses.

I also estimate the linear probability model separately by political identity. Table 3 shows that while both liberals and conservatives care about all attributes in the expected directions, the magnitudes differ. In line with political stereotypes, liberals place substantially more weight on recipient income ($\hat{\delta}_{\text{income}} = -0.174$ versus -0.077), indicating stronger redistributive preferences, while conservatives are more sensitive to fiscal costs ($\hat{\delta}_{\text{cost}} = -0.129$ versus -0.056). Both coefficient differences are statistically significant at the 1% level. However, contrary to political stereotypes, the coefficients on both factors for both political ideologies are negative and different from zero at the 1% significance level. In other words, American

conservatives do care about recipient income while American liberals do care about fiscal costs of redistribution. Moreover, conservatives and liberals alike express a strong aversion to transfer recipient labor-supply responses ($\hat{\delta}_{\text{labor}} = -0.365$ versus -0.227). While the difference between the two coefficients is not statistically significant, both coefficients are again negative and different from zero at the 1% level. Again, this confirms one political stereotype while contradicting another: American conservatives unsurprisingly are averse to transfer recipient labor-supply responses, but the same is true of American liberals as well. Interestingly, only American liberals place significant positive weight on AI automation risk ($\hat{\delta}_{\text{AI}} = 0.057$ versus -0.005).

The ubiquity of distaste for labor-supply responses among Americans contrasts sharply with results from the European sample. As shown in the last two columns of Table 3, European liberals place no statistically significant weight on recipient labor-supply responses ($\hat{\delta}_{\text{labor}} = -0.019$), while the coefficient for European conservatives ($\hat{\delta}_{\text{labor}} = -0.298$) is roughly halfway between the coefficients for American liberals and conservatives. These estimates imply that only conservative Europeans are concerned about transfer-induced labor-supply responses.

The descriptive analysis thus far does not capture respondent-level heterogeneity beyond differences by political ideology, and it does not directly quantify the welfare weights needed for policy evaluation. I therefore next estimate a mixed logit model that allows for flexible preference heterogeneity. The utility that respondent i derives from transfer k in question q is specified as

$$U_{iqk} = \frac{1}{\sigma_i} \left\{ \frac{\exp(-\gamma_i Z_{iqk}) + \beta_i^K (K_{iqk} - 1) + \beta_i^A A_{iqk}}{\gamma_i} T_{iqk} + \beta_i^L L_{iqk} - \beta_i^C C_{iqk} \right\} + \varepsilon_{iqk}, \quad (7)$$

where Z_{iqk} is recipient pre-tax income, K_{iqk} is an indicator for in-kind transfers, A_{iqk} is an indicator for AI automation risk, T_{iqk} is the transfer amount, L_{iqk} is the labor-supply reduction, C_{iqk} is the fiscal cost, and ε_{iqk} is a Type-I extreme value error. Because $K_{iqk} - 1$ equals zero

Table 3: Results by Political Affiliation

	United States		Europe	
	Liberal	Conservative	Liberal	Conservative
Income	−0.174*** (0.006)	−0.077*** (0.007)	−0.158*** (0.007)	−0.118*** (0.011)
Cost of Transfer	−0.056** (0.024)	−0.129*** (0.026)	−0.126*** (0.027)	−0.173*** (0.044)
Decrease in Labor	−0.227*** (0.085)	−0.365*** (0.090)	−0.019 (0.096)	−0.298** (0.152)
Transfer In-Kind	0.062*** (0.010)	0.055*** (0.011)	0.086*** (0.011)	0.058*** (0.018)
AI Risk	0.057*** (0.010)	−0.005 (0.011)	0.047*** (0.011)	−0.013 (0.018)
Observations	3,890	4,060	3,160	1,400

*p<0.1; **p<0.05; ***p<0.01

Notes: This table contains regression coefficients for a linear probability model regression, where the outcome is an indicator for choosing the first family. The regressors are constructed by taking the quantity for family 1 and subtracting the analogous quantity for family 2. Columns subset each sample to self-identified liberals and conservatives. In the European survey, these are respondents who place themselves on the left and right of the political spectrum, respectively. Standard errors are in parentheses.

for in-kind transfers, the remaining terms describe preferences over redistribution delivered through existing in-kind programs, and β_i^K captures how the valuation of a transfer changes when it is made in cash. The term $\exp(-\gamma_i Z_{iqk})/\gamma_i$ implies a standard constant absolute risk aversion (CARA) form for redistributive preferences over incomes. Respondents with higher values of $\gamma_i > 0$ place higher weight on lower-income recipients. The parameter σ_i captures the precision of respondent i 's responses, with higher values indicating less noisy choices. The coefficient β_i^C measures the weight i places on the government's budget constraint.

I specify preference parameters as distributed using a multivariate normal distribution. I impose sign restrictions on β_i^C , γ_i , and $1/\sigma_i$ to ensure economically sensible preferences by assuming that $\log(\exp(\beta_i^C) - 1)$, $\log(\exp(\gamma_i) - 1)$, and $\log(\exp(1/\sigma_i) - 1)$ are jointly normal,

Table 4: Mixed Logit Estimates

	United States		Europe	
	Mean	Std. Dev.	Mean	Std. Dev.
$\log(\exp(1/\sigma_i) - 1)$	13.02 (1.57)	21.93 (3.25)	11.31 (6.96)	19.53 (5.46)
$\log(\exp(\gamma_i) - 1)$	3.31 (0.21)	0.81 (0.12)	2.67 (1.49)	0.17 (0.56)
β_i^L	-0.46 (0.07)	0.40 (0.07)	-0.30 (0.46)	0.51 (0.42)
β_i^K	0.18 (0.02)	0.20 (0.02)	0.29 (0.07)	0.21 (0.11)
β_i^A	0.07 (0.02)	0.21 (0.02)	0.09 (0.16)	0.14 (0.09)
$\log(\exp(\beta_i^C) - 1)$	-2.38 (0.15)	1.06 (0.12)	-2.06 (2.17)	0.96 (0.50)

Notes: This table reports coefficient estimates from the mixed logit model, estimated separately on the US and European samples. For each sample, the first column contains the estimates of the mean random coefficient, while the second column contains the standard deviation. Standard errors are in parentheses.

so β_i^C , γ_i , and $1/\sigma_i$ are positive.¹² Table 4 reports the estimated means and standard deviations of the preference distribution. The results confirm substantial preference heterogeneity across respondents. For instance, the standard deviation of the labor-supply coefficient β_i^L is nearly as large as its mean, indicating that while respondents on average dislike labor-supply reductions, there is substantial disagreement about how much weight to place on this consideration, with some respondents even placing *positive* weight on policies that reduce work among recipients. The European estimates, reported in the last two columns of Table 4, are less precise given the smaller sample, but they mirror the linear probability results: the mean of β_i^L is smaller in magnitude than in the US sample and is not statistically different from zero. The estimated distributions of random coefficients are used in Section 4 to evaluate the popularity of different policy proposals. Before doing so, I need to model the responses of households to changes in redistributive policy, which I turn to next.

¹²In analyzing results from a pilot experiment, I originally experimented with parameterizing the marginal distributions of these random coefficients as log-normal, but that specification was numerically unstable.

3 Behavioral Responses to Tax/Transfer Reform

In this section, I describe the model of income taxation used to simulate counterfactual responses to tax reforms. Before discussing my empirical specification in more detail, it is useful to contrast my approach with the more standard specification used in empirical optimal taxation applications. Specifically, in many papers (Saez, 2010; Best and Kleven, 2012; Chetty, 2012; Chetty and Saez, 2013; Kleven and Waseem, 2013; Lockwood, 2020), utility is assumed to take the form

$$U(c, z, \theta) = c - \frac{(z/\theta)^{1+\epsilon}}{1+\epsilon}. \quad (8)$$

This specification implies two important restrictions on the data-generating process. First, by quasi-linearity, it assumes that lump-sum transfers do not cause income effects. Second, the parametric form of the labor disutility function implies that elasticities of taxable income are constant throughout the income distribution. Recent work by Rodemeier and Sun (2026) provides evidence against the second assumption, while the first assumption rules out income effects, whose existence has been demonstrated in recent empirical work (Goloso et al., 2024; Vivalt et al., 2024). Motivated by these issues, I relax both of the above restrictions in my empirical specification:


$$U(c, z, \theta) = u(c) - v(z/\theta). \quad (9)$$

In this specification, curvature in the consumption utility index u generates income effects, while non-constant curvature in v allows the elasticity of taxable income to vary by income. The structural primitives to identify are the utility of consumption u , the disutility of labor v , and the distribution F_θ of unobserved productivity θ .

Sun (2026) establishes nonparametric identification of models of this form given two empirical ingredients. The first ingredient is the Hicks elasticity of taxable income and how it varies across the income distribution. Rodemeier and Sun (2026) estimate this quantity,

with full uncertainty quantification, in a Bayesian meta-analysis, and Figure 3 reproduces their findings. The second ingredient is the marginal propensity to earn labor income out of unearned income (MPE) and how it varies across income levels.¹³ I estimate this quantity by drawing on recent experimental and quasi-experimental studies.

Figure 3: Posterior Distribution of ETI



Notes: This figure replicates Figure 4 from Rodemeier and Sun (2026). Each green line represents a single draw from the posterior distribution of their estimated elasticity of taxable income functions. The thick line represents the posterior mean, while the dotted lines represent 90% Bayesian credible intervals.

I draw on estimates from four papers that use plausibly exogenous variation to estimate MPEs. I treat each reported estimate as the average MPE among households in the corresponding range of the 2020 US income distribution. Two of the papers report estimates separately by income quantile, so each estimate maps to the corresponding income range. The third, a natural field experiment, restricts eligibility to participants with incomes below \$37K per year, so its estimate maps to incomes below \$37K. The fourth reports a single estimate for its entire sample, which I treat as the average over the entire income distribution. For the three lottery studies, I use each paper’s preferred MPE estimate. The field experiment reports the effect of a \$10,000 transfer on labor income in thousands of dollars, so I divide its estimate by 10 to obtain an MPE. Details about the papers and the MPE estimates I extract from them are provided in Table 5.

¹³Formally, the MPE is the derivative $\frac{\partial z(T, \theta)}{\partial T(z(T, \theta))}$. Because tax burdens and incomes are both expressed in dollars, this quantity is scale-free, allowing for easy comparison across studies.

Table 5: Marginal Propensity to Earn Estimates from the Literature

Study	MPE	Std. Error	Income Range (\$)	Source
Vivalt et al. (2024)	−0.420	(0.100)	0–37,470	0.1 times first row of Table A.35
Golosov et al. (2023)	−0.308	(0.024)	0–34,301	First row of Table 4.1
Golosov et al. (2023)	−0.520	(0.020)	34,301–68,400	First row of Table 4.1
Golosov et al. (2023)	−0.589	(0.022)	68,400–123,580	First row of Table 4.1
Golosov et al. (2023)	−0.674	(0.039)	123,580–500,000	First row of Table 4.1
Imbens et al. (2001)	−0.124	(0.054)	0–500,000	First row of Table 5
Cesarini et al. (2017)	−0.580	(0.113)	0–44,488	Bottom left panel of Table A3
Cesarini et al. (2017)	−0.574	(0.117)	44,488–100,000	Bottom left panel of Table A3
Cesarini et al. (2017)	−1.230	(0.278)	100,000–500,000	Bottom left panel of Table A3

Notes: This table contains information about point estimates and standard errors of marginal propensities to earn out of unearned income (MPEs) for a number of studies. The last column describes the exact location in the respective paper each estimate comes from.

3.1 Estimation

Given the two empirical ingredients described above—Hicks elasticities from Rodemeier and Sun (2026) and marginal propensities to earn from the papers summarized in Table 5—I now describe how I estimate the structural primitives u , v , and F_θ . Because Sun (2026) proves nonparametric identification in these settings, I parameterize both u and v flexibly as fourth-degree B-splines, which makes both functions thrice continuously differentiable. This level of smoothness implies that standard comparative statics quantities (which depend on the second derivatives of u and v) are continuous as a function of income. I parameterize the consumption utility function as $\hat{u}(c; \phi_u) \equiv \sum_{k=1}^{K_u+4} \phi_{uk} \mathcal{B}_{uk}(c)$, and the labor disutility function as $\hat{v}(\ell; \phi_v) \equiv \sum_{k=1}^{K_v+4} \phi_{vk} \mathcal{B}_{vk}(\ell)$, where the basis functions $\mathcal{B}_u = \{\mathcal{B}_{u1}, \dots, \mathcal{B}_{u, K_u+4}\}$ and $\mathcal{B}_v = \{\mathcal{B}_{v1}, \dots, \mathcal{B}_{v, K_v+4}\}$ are determined by the Cox–de Boor recursion formula (de Boor, 2001) given knot vectors $\kappa_u = \{\kappa_{u1} < \dots < \kappa_{u, K_u+1}\}$ and $\kappa_v = \{\kappa_{v1} < \dots < \kappa_{v, K_v+1}\}$, which partition the support of consumption and labor supply respectively into K_u and K_v uniformly spaced sub-intervals. I set $K_u = 4$ and $K_v = 6$. I use fewer basis functions when parameterizing u because, as shown in Sun (2026), the utility function largely controls the size of the MPE, and the heterogeneity analysis in the MPE literature I draw on in this paper has limited granularity. Thus, using a richer specification would require extrapolating income effect heterogeneity beyond what the data can reasonably identify. Nonetheless, the specification of consumption utility is reasonably flexible.

Identification is based on rationalizing ETIs and MPEs, so my estimator chooses utility parameters to match both sets of quantities simultaneously. Each moment is an average among households in a given income range, where averages over income z are taken with respect to the 2020 US income distribution G described below. Because Rodemeier and Sun (2026) estimate the Hicks elasticity as a function of income, $\zeta(z)$, its average can be computed over any range. I split each of the K_v sub-intervals defined by the knots κ_v at its midpoint, which yields $2K_v$ intervals $I_1, \dots, I_{2K_v}$. For each $j = 1, \dots, 2K_v$, the empirical moment ζ_j^{data} is the mean of $\mathbb{E}[\zeta^{(s)}(z)|z \in I_j]$ across draws $\zeta^{(s)}$ from their posterior distribution, and its model counterpart $\zeta_j^{\text{model}}(\phi_u, \phi_v)$ is the average over I_j of the Hicks elasticity implied by $\hat{u}(\cdot; \phi_u)$ and $\hat{v}(\cdot; \phi_v)$. Similarly, for each $k = 1, \dots, 9$, the k th estimate in Table 5 gives an empirical MPE moment η_k^{data} , whose model counterpart $\eta_k^{\text{model}}(\phi_u, \phi_v)$ is the average MPE implied by $\hat{u}$ and $\hat{v}$ over the corresponding income range. Point estimates then solve

$$\begin{aligned} \hat{\phi}_u, \hat{\phi}_v = \underset{\phi_u, \phi_v}{\operatorname{argmin}} \quad & \sum_j w_j^\zeta (\zeta_j^{\text{model}}(\phi_u, \phi_v) - \zeta_j^{\text{data}})^2 + \sum_k w_k^\eta (\eta_k^{\text{model}}(\phi_u, \phi_v) - \eta_k^{\text{data}})^2 \\ \text{s.t.} \quad & \hat{u}(\cdot; \phi_u) \text{ increasing, concave, } \hat{v}(\cdot; \phi_v) \text{ increasing, convex} \\ & \hat{u}(0; \phi_u) = \hat{v}(0; \phi_v) = 0, \hat{u}'(0; \phi_u) = 1, \theta(z_{0.5}) = 1. \end{aligned} \tag{10}$$

The weights w_j^ζ and w_k^η are inversely proportional to the variance of the corresponding empirical moment (across posterior draws for the ETI moments, and the squared standard error in Table 5 for the MPE moments), so more precisely estimated moments receive greater weight in the objective function. The shape restrictions on u and v are implemented as linear inequality constraints on the B-spline coefficients.¹⁴ The last constraint that $\theta(z_{0.5}) = 1$ is a scale normalization which implies that the median productivity type has $\theta = 1$. Thus, θ represents how many times more productive a household is relative to the median.

Jointly fitting all structural parameters to both sets of moments is numerically unstable, so I adopt a two-step procedure. I first fit B-spline approximations to the observed rela-

¹⁴When using a B-spline parameterization, monotonicity and curvature constraints can be expressed as linear inequalities in the coefficients; see, for instance, Bodoh-Creed et al. (2026).

relationship between pre-tax income and MPE and Hicks elasticities, minimizing the weighted squared distance between the spline-implied and empirical moments. I then use these reduced-form fits as starting values for the structural estimation, which fits u and v to rationalize both sets of moments simultaneously.


Once the B-spline coefficients for u and v are estimated, the distribution of unobserved heterogeneity F_θ can be recovered by inverting the relationship between observables and the structural model. Specifically, I parameterize the distribution of income $G(z)$ as coming from a Pareto-Lognormal distribution (Hajargasht and Griffiths, 2013) and fit the parameters to match the percentiles of the 2020 US income distribution reported in the 2020 Annual Social and Economic Supplement of the US Census. Then, for each income level z in the support of the observed income distribution, I compute the implied type $\theta(z)$ by solving the first-order condition of the household's optimization problem, which states that the marginal rate of substitution between consumption and labor equals the net-of-tax wage rate. The distribution of unobserved types is then given by $F_\theta(\theta(z)) = G(z)$.

To quantify uncertainty in my estimates, I employ a parametric bootstrap procedure. I treat the point estimates and variance-covariance matrix of the ETI and MPE moments as defining a joint normal distribution over the empirical inputs. For each of 500 bootstrap iterations, I draw a new set of moments from this distribution and re-estimate the model. This yields a distribution of estimates for the structural primitives, from which I construct 95% confidence intervals.

Figure 4 displays the estimation results. Because utility functions are difficult to visually interpret, I report two structural primitives that summarize the estimates. The left panel plots the estimated distribution of log productivity, $\log(\theta)$, across quantiles of the income distribution, with the solid line showing the point estimate and the dotted lines indicating the 95% confidence interval. Sun (2026) shows that this quantity is identified even in a more general specification of utility $U(c, z/\theta)$, so this distribution is robustly estimated even if the additively separable form of utility is misspecified. Similarly, Sun (2026) shows that

if additive separability is assumed, but utility depends more flexibly on θ , the (inverse) curvature of the consumption utility function, $u'(c)/(c \cdot u''(c))$, is identified, even if θ is not interpreted as productivity. The inverse curvature of utility is informative about the ratio between substitution and income effects. The fact that this curvature is near 0 implies the existence of large income effects, which is unsurprising given the large MPEs observed in my meta-analysis.

Figure 4: Structural Estimates



Notes: This figure plots functions that summarize structural utility models. Panel A plots the quantile function of log productivity, while Panel B plots the inverse curvature, $u'(c)/(cu''(c))$, of the consumption utility function.

Because the B-spline approximations are defined only up to approximately the 99th percentile of the income distribution, while counterfactuals also depend on behavior in the upper tail, I extrapolate both utility functions beyond this range by assuming that they exhibit constant relative risk aversion above it, as described in Appendix D.1.

4 Counterfactuals

With estimates of preferences over redistributive policy from Section 2 and estimates of behavioral responses to tax reform from Section 3 in hand, I conduct counterfactuals to understand the consequences of these estimates for policymakers. To keep my approach as

comparable as possible to prior work on optimal income taxation, I begin by solving the canonical social planner’s problem. I then introduce an additional constraint that the policy chosen by the social planner must be agreeable to a sufficiently large fraction of citizens, and show how the social planner should optimally alter her policy to meet this constraint.¹⁵

Recall that in the policy preferences of respondents described by Equation (7), the random coefficient γ_i parameterizes standard CARA preferences for redistribution. I take the median value of γ_i from the random coefficient model, $\bar{\gamma}$, as the social planner’s preference. More precisely, let T_0 denote the status quo tax policy.¹⁶ The social planner’s objective is

$$\text{SWF}(T) = \int_{\theta} \omega(\theta) [u(z(T, \theta) - T(z(T, \theta))) - v(z(T, \theta)/\theta)] dF_{\theta}(\theta), \quad (11)$$

where the weights $\omega(\theta)$ are chosen to match the survey estimates. In the survey analysis, the value that the median respondent places on an extra dollar given to a household with pre-tax income z is proportional to $\exp(-\bar{\gamma}z)$. In the model, the social value of an extra dollar given to a type- θ household equals the planner’s weight on that household’s utility, $\omega(\theta)$, times the household’s marginal utility of consumption. Equating the two under the status quo yields

$$\omega(\theta) = \frac{\exp(-\bar{\gamma}z(T_0, \theta))}{u' [z(T_0, \theta) - T_0(z(T_0, \theta))]}.$$

The social planner aims to maximize this social welfare function, subject to a revenue constraint, which I assume that the status quo policy satisfies exactly: $\underline{R} = \int_{\theta} T_0(z(T_0, \theta)) dF_{\theta}(\theta)$. Rather than finding a globally optimal tax policy, I instead solve for an optimal local reform. I do this because my specification of redistributive preferences of the public using GSMWWs implicitly linearizes the preferences of the public. As a result, this framework requires more extrapolation the further the proposed policy is from the status quo, and the focus on small

¹⁵While this approach prioritizes the social planner’s normative judgments, an alternative framework might place citizens’ preferences on equal footing. Appendix E shows how my empirical estimates are rich enough to accommodate such a formulation as well.

¹⁶Details about how T_0 is constructed are given in Appendix D.3.

reforms limits the need to extrapolate. Putting everything together, the social planner solves

$$\begin{aligned} \max_T \text{SWF}(T) \quad \text{s.t.} \quad & \int_{\theta} T(z(T, \theta)) dF_{\theta}(\theta) = \int_{\theta} T_0(z(T_0, \theta)) dF_{\theta}(\theta), \\ & \frac{\int_{\theta} (T'(z(T_0, \theta)) - T'_0(z(T_0, \theta)))^2 z(T_0, \theta)^2 dF_{\theta}(\theta)}{\int_{\theta} z(T_0, \theta)^2 dF_{\theta}(\theta)} \leq 0.01. \end{aligned} \quad (12)$$

The second constraint above requires that the weighted root mean squared difference between the status quo and new marginal tax rates not exceed 10%. I parameterize the tax function T flexibly as a B-spline, as described in Appendix D.2, and numerically solve the constrained optimization problem (12).


In Figure 5, I plot the results of the optimization. Panel A shows consumption as a function of pre-tax income, while Panel B shows implied marginal tax rates. The result reveals that the Mirrleesian social planner desires far more redistribution than the current tax-and-transfer system implements. The optimal reform involves raising taxes on top earners to fund substantially larger transfers to households making less than \$60,000/year. In doing so, the social planner raises marginal tax rates across the board, most sharply at low and middle incomes. As a result, it should be unsurprising that this preferred policy causes a notable decrease in labor supply.

Given the public's aversion to labor-supply reductions documented in Section 2.2, the large decline in work induced by the Mirrleesian policy suggests that it may be unpopular. To evaluate the popularity of a given policy, I simulate a population of citizens by drawing preference parameters from the random coefficients estimated from my survey sample. Specifically, for each simulated citizen i , I draw a vector of preference parameters $(\gamma_i, \beta_i^L, \beta_i^K, \beta_i^A, \beta_i^C)$ from the multivariate normal distribution estimated in Section 2.2. Given these parameters, citizen i 's decision utility when evaluating tax reform T relative to the status quo T_0 is

$$\nu_i(T) = \int \omega_i(\theta) \Delta U(\theta) dF_{\theta}(\theta) - \beta_i^L \cdot \Delta Z - \frac{\beta_i^K}{\gamma_i} \cdot \kappa + \beta_i^C \cdot \Delta R, \quad (13)$$

where $\Delta U(\theta)$ denotes the utility change for a type- θ household, κ denotes the dollar amount

Figure 5: Social Planner's Proposed Policy



Notes: This figure compares the social planner's (blue) most preferred policy change to the status quo policy (black). Panel A displays post-tax income as a function of pre-tax income, while Panel B displays implied marginal tax rates.

of in-kind transfers converted to cash under the reform, and ΔR denotes the change in net government revenue. The welfare weight function $\omega_i(\theta)$ captures citizen i 's preferences over which types of recipients should benefit from redistribution,

$$\omega_i(\theta) = \frac{\exp(-\gamma_i z(T_0, \theta)) + \beta_i^A \cdot a(z(T_0, \theta))}{\gamma_i u'[z(T_0, \theta) - T_0(z(T_0, \theta))]} \quad (14)$$

In the numerator, the first term reflects standard CARA preferences for redistribution, shared in common with the social planner, while the second term captures the extent to which citizen i places additional weight on workers facing automation risk, as measured by the coefficient β_i^A times a probability of AI automation $a(z)$. This probability of AI automation function, in turn, is based on estimates derived in Hampole et al. (2025). Implementation details can be found in Appendix D.4. The aggregate labor-supply effect ΔZ is defined as the average percent change in pre-tax income induced by the reform,

$$\Delta Z \equiv \int \left(\frac{z(T, \theta) - z(T_0, \theta)}{z(T_0, \theta)} \right) dF_\theta(\theta). \quad (15)$$

A citizen supports the proposed policy if and only if $\nu_i(T) > 0$, so the popularity of a policy is measured as the fraction of simulated citizens who support it. Plugging in the social planner’s preferred policy, I find that *just under one third* of the general public supports the social planner’s favorite policy. By contrast, when I evaluate the same policy using the preference distribution estimated from the European sample (Table 4), 52% of European citizens support it, so a majority-support constraint would not bind. This contrast is consistent with the view that the near-ubiquitous aversion to labor-supply responses among Americans, which is far less widespread among Europeans, may help explain why European governments implement more redistribution than the US (Alesina and Angeletos, 2005).

The unpopularity of the social planner’s most preferred policy raises the following question: what tweaks would help make the policy look more desirable to the US public?¹⁷ To answer this, I consider two modifications. First, because the public places direct weight on labor supply, the planner could seek reforms that induce smaller disincentives to work. Second, the planner could relax the government budget constraint, which would allow transfers to be phased out more slowly, and thus discourage work less, without requiring higher taxes elsewhere. Both modifications correspond to prominent features of the UBI debate studied in Section 2.1. Advocates emphasize that UBI would induce smaller work disincentives than traditional redistribution, while the version of UBI proposed during the 2020 presidential campaign involved a substantial increase in net government spending.

Both modifications can be incorporated into my framework. The first, reducing labor-supply impacts, is implemented by adding a penalty term $-\lambda \cdot \Delta Z$ to the social planner’s objective function. Solving this penalized problem for different values of λ traces out the same frontier of policies as would be obtained by directly constraining labor-supply effects to lie below progressively tighter thresholds, and I vary λ from 0—which corresponds to pure welfare maximization—up to values large enough that the resulting policy induces

¹⁷In the counterfactual exercises described in the body of the paper, I let the social planner’s preferences take normative precedence over the views of the general public, which act as an undesirable constraint on the social planner. In Appendix E, I also compute the policy that maximizes the public’s preferences.


negligible changes in aggregate labor supply. The second, relaxing the budget constraint, is implemented by allowing for a specified level of deficit spending. However, I find that relaxing the budget constraint consistently reduces popularity, because the public’s aversion to increases in government spending, as captured by the budget sensitivity coefficient β_i^C , outweighs the benefits of slower phase-outs. As a result, I do not pursue this possibility further in the main analysis.

In Figure 6, I show how modifications to the social planner’s preferred policy affect both the planner’s welfare function and public support. The horizontal axes report the labor-supply penalty multiplier λ , while the two vertical axes report social welfare (Panel A) and the fraction of the public supporting the policy (Panel B). The results reveal a mild tradeoff between the social planner’s preferences and political feasibility. At $\lambda = 0$, corresponding to pure welfare maximization, only about one third of the public supports the policy despite its welfare benefits. As λ increases and the social planner places greater weight on minimizing labor-supply effects, popularity rises while social welfare falls. Popularity crosses the 50% threshold, indicating majority support, at approximately $\lambda = 3.4$. The policy achieves super-majority support, defined as approval from more than two-thirds of the population, at approximately $\lambda = 7.7$. The policy attaining majority support retains over 90% of the gains the social planner would have attained relative to the status quo from her optimal unconstrained policy, and even the super-majority policy retains roughly three-quarters of these gains, showing that these tweaks do not require the planner to give up very much on her objective.¹⁸

In Figure 7, I plot the versions of the social planner’s policy that attain majority and super-majority support. I find that the form of the changes is intuitive. The pure redis-

¹⁸As shown in Appendix E, a social planner who fully caters to the public ends up with a policy that looks like a more extreme version of the best policy that attains super-majority support. Transfers to the poorest members of society are phased out even more slowly, and transfers at the bottom slightly decrease. An average member of the public thus does not appear to prioritize improving the welfare of the poorest members of society and instead views mitigating work disincentives among welfare recipients to be a much larger priority. The social welfare function barely increases under the policy that maximizes the average welfare of the public.

Figure 6: Desirability and Popularity of Modified Social Planner's Policies




Notes: This figure examines the welfare and popular support implications of adding a labor-supply penalty to the social planner's optimization problem (horizontal axis). Panel A shows how the social planner's objective changes as the labor penalty increases. Panel B shows how popular support changes as a function of the labor penalty.

tributive policy, corresponding to $\lambda = 0$, features high marginal tax rates at low incomes due to the rapid phase-out of transfers to the poor. This phase-out is efficient in the pure Mirrleesian problem because it targets transfers to those with the highest marginal welfare weights. However, it creates substantial work disincentives that the public finds objectionable. The majority-supported policy, corresponding to $\lambda \approx 3.4$, reduces the generosity of transfers slightly at the bottom of the income distribution and phases them out more gradually, thereby lowering marginal tax rates in the phase-out region while keeping marginal tax rates at high incomes approximately unchanged. The super-majority policy, corresponding to $\lambda \approx 7.7$, further softens the phase-out, resulting in a flatter marginal tax rate schedule throughout the lower portion of the income distribution.

These modifications echo the central features of UBI proposals: slower phase-outs reduce work disincentives, which the public cares about directly and independently of standard efficiency concerns. Critically, however, unlike many UBI proposals that never fully phase out transfers, the policies I identify eventually eliminate benefits as income rises, thereby maintaining budget neutrality. This distinction may help explain why existing UBI proposals

have struggled to gain traction: by phasing out transfers too slowly or not at all, they incur large fiscal costs that the public also dislikes. UBI advocates thus appear to have correctly intuited which features of a reform appeal to the US public, while economic analysis remains useful for translating these intuitions into policies that are also fiscally responsible.

Figure 7: Modified Social Planner’s Policies



Notes: This figure plots various versions of the social planner’s policy compared to the status quo tax-and-transfer system. The dotted line represents the unconstrained social planner’s policy. The dashed line represents the policy required to obtain majority support, while the solid blue line represents the policy required to obtain super-majority (more than two-thirds) support. Panel A plots post-tax income as a function of pre-tax income while Panel B plots implied marginal tax rates.

The main takeaways of this section are robust across a variety of alternative specifications. In Appendix D.5, I show that my findings remain similar even under substantial bias in the income and substitution effect moments used for estimation. Two concerns are especially important. First, the Hicks elasticities entering the meta-analysis of Rodemeier and Sun (2026) are typically short-run, and longer-run elasticities may be larger. Second, much of the evidence on MPEs comes from lottery winners, who may be a selected sample of the US population. Even if Hicks elasticities are three times larger than those I use, the social planner still prefers more redistribution, this redistribution remains deeply unpopular, and the same modification still improves its popularity. Similarly, if I instead set MPEs equal across the income distribution to the only experimental estimate I use (Vivalt et al., 2024), I

obtain nearly the same findings. The same appendix shows that alternative specifications of the random coefficients model of GSMWWs do not substantially affect the main conclusions of this section.

4.1 Do Welfare Weights Predict Policy Support?

While the welfare weight framework used in this paper is standard and convenient for its tractability, it can be restrictive. On the margin, it implies that preferences for policies can be fully described by *who* those policies affect. Welfare weights therefore have difficulty capturing some plausible forms of voter behavior. For instance, a voter who exhibits status quo bias (Samuelson and Zeckhauser, 1988) and thus rejects *all* deviations from the status quo cannot easily be described purely within the welfare weights framework. Conversely, voters could exhibit action bias (Patt and Zeckhauser, 2000), which increases preferences over *any* policy that differs from the status quo. Motivated by these possibilities, I design an additional survey experiment to test how well my policy preference estimates predict behavior when survey respondents directly choose whether or not to support a given policy. In the experiment, 812 respondents are shown three tax-and-transfer reform descriptions.¹⁹

The first description, which I refer to as *Mirrlees*, is of the policy computed by maximizing the social planner’s objective function. The description includes text explaining that the reform is motivated by a “desire to provide further assistance to low and medium income Americans.” It explains that (i) the policy will be financed by increasing taxes on the top 25% of earners, (ii) the policy is forecast to be budget neutral, and (iii) the policy will discourage work, with quantitative estimates for how much less the average American will work based on the predictions of the model estimated in Section 3.

The second description, which I refer to as *Popular (no AI)*, is similar to the first but shows numbers based on the modified policy that is forecast by the welfare weights to obtain majority support. A final description, *Popular (AI)*, is similar to the second, but explic-

¹⁹The experiment was pre-registered on the AEA RCT registry as trial number 19220.

itly adds additional text explaining that the proposed redistribution is meant to “mitigate concerns that artificial intelligence and related automation technologies will displace a large number of workers in the United States”.

Figure 8: Direct Policy Choice Survey Questions

The below information describes a proposed reform to the existing system of income taxation and redistribution. The reform design is motivated by a **desire to provide further assistance to low and medium income Americans**. The reform will consist of three parts.

- Generosity of income support for low-income Americans will be expanded, with the poorest Americans receiving additional monetary income assistance.
- Taxes on the top 30% of earners in the US will increase.

The details of reforms 1 and 2 above can be explored in more detail by interacting with the below tool. Each choice of slider value represents a level of pre-tax income. The tool will then inform you about what the effect of the proposed reform will be on the net tax burden of an individual with that level of income.

Pre-tax income: \$50,000
Receive **\$3,500** more in government benefits annually

Economists helped inform the design of the above proposal by choosing details so that the overall effect of the above reform is likely to be **budget neutral for the government**. In addition, economists calculated the likely effects of the above reform on how much Americans work. On average, they estimate that this policy proposal will cause the **average American** to work **18% less each year**. The effect is especially large amongst **recipients of the additional cash assistance**, who on average would be forecasted will work **43% less each year** under this reform.

Based on this information, do you support the above reform?

☐ I would support this policy
☐ I would not support this policy

The below information describes a proposed reform to the existing system of income taxation and redistribution. The reform design is motivated by a **desire to provide further assistance to low and medium income Americans**. The reform will consist of three parts.

- 20% of all existing federal government assistance for low-income Americans (for example SNAP, medicare, TANF) will be replaced with a direct cash transfer.
- Generosity of income support for low-income Americans will be expanded, with the poorest Americans receiving additional monetary income assistance.
- Taxes on the top 25% of earners in the US will increase.

The details of reforms 2 and 3 above can be explored in more detail by interacting with the below tool. Each choice of slider value represents a level of pre-tax income. The tool will then inform you about what the effect of the proposed reform will be on the net tax burden of an individual with that level of income.

Pre-tax income: \$50,000
Receive **\$5,800** more in government benefits annually

Economists helped inform the design of the above proposal by choosing details so that the overall effect of the above reform is likely to be **budget neutral for the government**. In addition, economists calculated the likely effects of the above reform on how much Americans work. On average, they estimate that this policy proposal will cause the **average American** to work **9% less each year**. The effect is especially large amongst **recipients of the additional cash assistance**, who on average would be forecasted will work **19% less each year** under this reform.

Based on this information, do you support the above reform?

☐ I would support this policy
☐ I would not support this policy

The below information describes a proposed reform to the existing system of income taxation and redistribution. The reform design is motivated by a **desire to provide further assistance to low and medium income Americans and mitigate concerns that artificial intelligence and related automation technologies will displace a large number of workers in the United States**. The reform will consist of three parts.

- 20% of all existing federal government assistance for low-income Americans (for example SNAP, medicare, TANF) will be replaced with a direct cash transfer.
- Generosity of income support for low-income Americans will be expanded, with the poorest Americans receiving additional monetary income assistance.
- Taxes on the top 25% of earners in the US will increase.

The details of reforms 2 and 3 above can be explored in more detail by interacting with the below tool. Each choice of slider value represents a level of pre-tax income. The tool will then inform you about what the effect of the proposed reform will be on the net tax burden of an individual with that level of income.

Pre-tax income: \$127,000
Pay **\$12,400** more in taxes annually

Economists helped inform the design of the above proposal by choosing details so that the overall effect of the above reform is likely to be **budget neutral for the government**. In addition, economists calculated the likely effects of the above reform on how much Americans work. On average, they estimate that this policy proposal will cause the **average American** to work **9% less each year**. The effect is especially large amongst **recipients of the additional cash assistance**, who on average would be forecasted will work **19% less each year** under this reform.

Based on this information, do you support the above reform?

☐ I would support this policy
☐ I would not support this policy

(a) Mirrlees

(b) Popular (no AI)


(c) Popular (AI)

Notes: This figure provides screenshots of the survey questions directly asking about preferences for the policies computed in the counterfactual exercises. The slider is interactive and will dynamically adjust as the respondent sets its level.

To ensure that respondents fully understand the proposed reform, in addition to the text, each description embeds an interactive slider that shows, for any pre-tax income level selected by the respondent, how the proposed policy would change that household’s net tax burden. The numbers displayed by the slider exactly match the policies computed in this section. Figure 8 provides screenshots of the policy descriptions as well as their accompanying sliders.

Each respondent is shown all three policy descriptions, with the order of presentation

Figure 9: Support for Various Policies



Notes: This figure plots the proportion of respondents who choose to support each of the three policies. Error bars represent 95% confidence intervals.

following a cyclic Latin square design.²⁰ In Figure 9, I plot the proportion of respondents expressing support for each of the three policies.²¹ The findings are consistent with the predictions of the welfare weight model. The social planner’s policy is found to be unpopular, with only 42.8% of respondents (95% CI: [39.4%, 46.3%]) choosing to support it. By contrast, both of the modified policy descriptions are much more popular. The *Popular (AI)* description attracts 57.5% support (95% CI: [54.1%, 60.9%]), and the *Popular (no AI)* description attracts 56.3% (95% CI: [52.9%, 59.8%]). Importantly, the modified policy receives statistically significantly more than majority support at the 5% significance level, regardless of the AI framing. Meanwhile, the social planner’s policy attracts statistically significantly less than majority support. The AI framing, if anything, slightly increases popularity, but the point estimates are very close and the confidence intervals are too wide to reject equality. The differences in support are, however, smaller than the welfare weights predict, which

²⁰Specifically, respondents are shown policies in one of the following three orders: Mirrlees → Popular (no AI) → Popular (AI), Popular (no AI) → Popular (AI) → Mirrlees, or Popular (AI) → Mirrlees → Popular (no AI).

²¹I use all choices from the experiment to construct this figure. Appendix F checks that the findings reported here are robust to order effects.

could reflect either misspecification of my particular parameterization of the welfare weights or limitations of the welfare weight framework more generally. Nonetheless, the results show that the estimated weights are informative about which modifications the social planner should make to her policy to substantially improve political feasibility.

5 Conclusion

This paper studies how the public’s views constrain the redistribution policies that economists can feasibly recommend, measuring those views with a combination of text data from online policy debates and a survey experiment. When ordinary citizens discuss redistribution, they emphasize considerations that fall outside the scope of the standard Mirrleesian framework.

Estimated preferences imply that the Mirrleesian social planner’s optimal policy is deeply unpopular among Americans, garnering support from just under one third of the general public. However, modest adjustments achieve majority or even super-majority support while still accomplishing much of the social planner’s redistributive objective. These politically feasible policies address the public’s aversion to the negative labor-supply responses induced by redistribution by giving smaller transfers to the lowest-income households but phasing these transfers out more slowly than the Mirrleesian optimum in order to reduce marginal tax rates. They do so while maintaining budget neutrality, which is important because the public is also concerned about the fiscal responsibility of the government. I directly validate the popularity of these modified reforms in a pre-registered experiment.

My findings have several implications for redistributive policy design. First, they suggest that economists seeking to influence redistribution policy would benefit from engaging with the public’s concerns about work incentives, even when those concerns go beyond what standard efficiency analysis justifies. Second, the results empirically show that accommodating the public’s concerns need not be expensive: the planner can satisfy the popular support constraint while retaining the vast majority of the welfare gains from her preferred reform.

More broadly, understanding how the public reasons about policy is essential for designing reforms that improve welfare while remaining politically viable. This paper establishes that in the case of redistribution policy, the gap between economists and the public is consequential but not insurmountable. By taking the considerations that matter to ordinary citizens seriously, it is possible to identify reforms that economists favor and that a majority of the public also supports. Similar tensions arise in other domains, including trade, housing, and climate change, and the methodology developed in this paper can help identify analogous opportunities for policy design in these other contexts as well.

References

- Acemoglu, D., Golosov, M., and Tsyvinski, A. (2010). Dynamic mirrlees taxation under political economy constraints. *The Review of Economic Studies*, 77(3):841–881.
- Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American economic review*, 95(4):960–980.
- Almås, I., Cappelen, A. W., and Tungodden, B. (2020). Cutthroat capitalism versus cuddly socialism: Are americans more meritocratic and efficiency-seeking than scandinavians? *Journal of Political Economy*, 128(5):1753–1788.
- Anthropic (2025). Claude API. Large language model API.
- Autor, D., Dorn, D., and Hanson, G. (2025). How the china trade shock impacted us manufacturing workers and labor markets, and the consequences for us politics. *Washington Center for Equitable Growth*, May.
- Azevedo-McCaffrey, D. and Aguas, T. (2025). Continued increases in TANF benefit levels are critical to helping families. Technical report, Center on Budget and Policy Priorities.
- Bartik, A. W., Rhodes, E., Broockman, D. E., Krause, P. K., Miller, S., and Vivalt, E. (2024). The impact of unconditional cash transfers on consumption and household balance sheets: Experimental evidence from two us states. Technical report, National Bureau of Economic Research.
- Benabou, R. and Ok, E. A. (2001). Social mobility and the demand for redistribution: the poum hypothesis. *The Quarterly journal of economics*, 116(2):447–487.
- Berliant, M. and Gouveia, M. (2025). On the political economy of nonlinear income taxation. *Working Paper*.
- Best, M. and Kleven, H. (2012). Optimal income taxation with career effects of work effort. *Available at SSRN 2051214*.
- Bierbrauer, F. J. and Boyer, P. C. (2016). Efficiency, welfare, and political competition. *The Quarterly Journal of Economics*, 131(1):461–518.
- Bodoh-Creed, A. L., Hickman, B. R., List, J. A., Muir, I., and Sun, G. K. (2026). Salience and (non-) buyer’s remorse: Optimal nonlinear pricing with cognitively constrained consumers. Technical report, National Bureau of Economic Research.

- Brooks, T., Touschner, J., Artiga, S., Zuvekas, J., and Rosenbaum, S. (2015). Medicaid and CHIP eligibility, enrollment, renewal, and cost-sharing policies as of january 2015. Technical report, Kaiser Family Foundation.
- Buchanan, S., Pai, D., Wang, P., and Ma, Y. (2025). *Learning Deep Representations of Data Distributions*. Online. <https://ma-lab-berkeley.github.io/deep-representation-learning-book/>.
- Capozza, F. and Srinivasan, K. (2024). Who should get money? estimating welfare weights in the us. *URPP Equality of Opportunity Discussion Paper Series*.
- Cappelen, A. W., Haaland, I. K., and Tungodden, B. (2018). Beliefs about behavioral responses to taxation. *Unpublished manuscript. Norwegian School of Economics, Bergen*.
- Cappelen, A. W., Hole, A. D., Sørensen, E. Ø., and Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, 97(3):818–827.
- Cappelen, A. W., Konow, J., Sørensen, E. Ø., and Tungodden, B. (2013). Just luck: An experimental study of risk-taking and fairness. *American Economic Review*, 103(4):1398–1413.
- Centers for Medicare and Medicaid Services, Office of the Actuary (2015). 2015 actuarial report on the financial outlook for medicaid. Technical report, U.S. Department of Health and Human Services.
- Cesarini, D., Lindqvist, E., Notowidigdo, M. J., and Östling, R. (2017). The effect of wealth on individual and household labor supply: evidence from swedish lotteries. *American Economic Review*, 107(12):3917–3946.
- Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018.
- Chetty, R. and Saez, E. (2013). Teaching the tax code: Earnings responses to an experiment with eitc recipients. *American Economic Journal: Applied Economics*, 5(1):1–31.
- Chien, N. (2016). Estimates of child care eligibility and receipt for fiscal year 2012. Technical report, Office of Planning, Research and Evaluation, Administration for Children and Families, U.S. Department of Health and Human Services.
- Child Care Aware of America (2015). Parents and the high cost of child care: 2015 report. Technical report.
- 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.
- de Boor, C. (2001). *A Practical Guide to Splines Revised Edition*. New York: Springer-Verlag.
- Diamond, R., McQuade, T., and Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–3394.
- Dreyfuss, B. (2025). Equilibrium neglect and political feasibility. *Working Paper*.
- Eden, M. and Piacquadio, P. G. (2025). *The ethical mirror*. Centre for Economic Policy Research.
- Elmendorf, C. S., Nall, C., and Oklobdzija, S. (2025). The folk economics of housing. *Journal of Economic Perspectives*, 39(3):45–66.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

- 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.
- Fershtman, C., Gneezy, U., and List, J. A. (2012). Equity aversion: Social norms and the desire to be ahead. *American Economic Journal: Microeconomics*, 4(4):131–144.
- Fisman, R., Gladstone, K., Kuziemko, I., and Naidu, S. (2020). Do americans want to tax wealth? evidence from online surveys. *Journal of Public Economics*, 188:104207.
- Fleurbaey, M. and Maniquet, F. (2006). Fair income tax. *The Review of Economic Studies*, 73(1):55–83.
- Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2023). How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *Quarterly Journal of Economics*.
- Golosov, M., Graber, M., Mogstad, M., and Novgorodsky, D. (2024). How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income. *The Quarterly Journal of Economics*, 139(2):1321–1395.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- Gray, K. F. and Cunyningham, K. (2016). Trends in supplemental nutrition assistance program participation rates: Fiscal year 2010 to fiscal year 2014. Technical report, USDA Food and Nutrition Service.
- Grootendorst, M. (2022). Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Guzman, G. and Kollar, M. (2024). Income in the United States: 2023. Current Population Reports P60-282, U.S. Census Bureau.
- Hajargasht, G. and Griffiths, W. E. (2013). Pareto–lognormal distributions: Inequality, poverty, and estimation from grouped income data. *Economic Modelling*, 33:593–604.
- Hampole, M., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2025). Artificial intelligence and the labor market. Technical report, National Bureau of Economic Research.
- Hvidberg, K. B., Kreiner, C. T., and Stantcheva, S. (2023). Social positions and fairness views on inequality. *Review of Economic Studies*, 90(6):3083–3118.
- Imbens, G. W., Rubin, D. B., and Sacerdote, B. I. (2001). Estimating the effect of unearned income on labor earnings, savings, and consumption: Evidence from a survey of lottery players. *American economic review*, 91(4):778–794.
- Internal Revenue Service (2014). EITC participation rate by states. Technical report, U.S. Department of the Treasury.
- Internal Revenue Service (2015a). 2015 instructions for form 1040. Technical report, U.S. Department of the Treasury.
- Internal Revenue Service (2015b). Child tax credit parameters for tax year 2015. Technical report, U.S. Department of the Treasury. Instructions for Form 1040, Schedule 8812.
- Internal Revenue Service (2015c). Earned income tax credit parameters for tax year 2015. Technical report, U.S. Department of the Treasury. Publication 596.
- Jaroszewicz, A., Jachimowicz, J., Hauser, O., and Jamison, J. (2024). How effective is (more) money? randomizing unconditional cash transfer amounts in the us. *Randomizing Unconditional Cash Transfer Amounts in the US (July 20, 2024)*.
- Jiang, C., Lauletta, M., Levy, R., Shapiro, J. S., and Taubinsky, D. (2026). Understanding support for inefficient environmental policy instruments. *Working Paper*.

- Keane, M. P. and Wasi, N. (2016). Labour supply: the roles of human capital and the extensive margin. *The Economic Journal*, 126(592):578–617.
- Kenney, G. M., Copeland, V. L., and Resnick, D. M. (2012). Medicaid/chip participation among children and parents. Technical report, Urban Institute.
- 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.
- Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American economic review*, 90(4):1072–1092.
- Kosar, G. and Moffitt, R. A. (2017). Trends in cumulative marginal tax rates facing low-income families, 1997–2007. *Tax Policy and the Economy*, 31(1):43–70.
- 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.
- List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political economy*, 115(3):482–493.
- Lockwood, B. B. (2020). Optimal income taxation with present bias. *American Economic Journal: Economic Policy*, 12(4):298–327.
- Maag, E. (2018). Who benefits from expanding the child tax credit? Technical report, Urban Institute.
- McInnes, L., Healy, J., and Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Meltzer, A. H. and Richard, S. F. (1981). A rational theory of the size of government. *Journal of political Economy*, 89(5):914–927.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Miller, S., Rhodes, E., Bartik, A. W., Broockman, D. E., Krause, P. K., and Vivalt, E. (2024). Does income affect health? evidence from a randomized controlled trial of a guaranteed income. Technical report, National Bureau of Economic Research.
- Mirrlees, J. A. (1971). An exploration in the theory of optimal taxation. *The Review of Economic Studies*, 38(2):175–208.
- Norton, M. I. and Ariely, D. (2011). Building a better america—one wealth quintile at a time. *Perspectives on psychological science*, 6(1):9–12.
- of Housing, U. D., of Policy Development, U. D. O., and Research (2006). Effects of housing vouchers on welfare families. Technical report, U.S. Department of Housing and Urban Development, Office of Policy Development and Research.
- Patt, A. and Zeckhauser, R. (2000). Action bias and environmental decisions. *Journal of Risk and Uncertainty*, 21(1):45–72.
- Piketty, T. and Saez, E. (2013). Optimal labor income taxation. In *Handbook of public economics*, volume 5, pages 391–474. Elsevier.
- Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 3982–3992.

- Roberts, K. W. (1977). Voting over income tax schedules. *Journal of public Economics*, 8(3):329–340.
- Rodemeier, M. and Sun, G. (2026). Revealed preference for redistribution and the role of government. *Available at SSRN 5687044*.
- Roemer, J. E. (1999). The democratic political economy of progressive income taxation. *Econometrica*, 67(1):1–19.
- 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.
- Samuelson, W. and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1):7–59.
- Shapiro, B. T., Drango, S., and Moshary, S. (2026). Firearm policy buyers would accept. *Working Paper*.
- Snyder, J. M. and Kramer, G. H. (1988). Fairness, self-interest, and the politics of the progressive income tax. *Journal of Public economics*, 36(2):197–230.
- Stantcheva, S. (2021). Understanding tax policy: How do people reason? *The Quarterly Journal of Economics*, 136(4):2309–2369.
- Sun, G. (2026). Nonparametric identification in optimal income taxation models. *Working Paper*.
- U.S. Census Bureau (2020). Annual social and economic supplement (ASEC). Current Population Survey. Accessed: June 9, 2026.
- U.S. Department of Agriculture, Food and Nutrition Service (2015). Supplemental nutrition assistance program: Benefit allotments for fiscal year 2015.
- U.S. Department of Health and Human Services (2015). Annual update of the HHS poverty guidelines. January 22, 2015.
- U.S. Department of Housing and Urban Development (2015). Housing choice vouchers fact sheet. Technical report.
- Vayansky, I. and Kumar, S. A. (2020). A review of topic modeling methods. *Information Systems*, 94:101582.
- Vivalt, E., Rhodes, E., Bartik, A. W., Broockman, D. E., Krause, P., and Miller, S. (2024). The employment effects of a guaranteed income: Experimental evidence from two us states. Technical report, National Bureau of Economic Research.
- Zedlewski, S. and Holland, J. (2003). Work activities of current welfare recipients. Technical report, Urban Institute.

A Text Analysis Details

A.1 Text Analysis Pipeline

In this appendix, I describe the text analysis workflow in more detail. I scrape Reddit using an open source Python package built specifically for Reddit, praw. For a given keyword and subreddit pair, this package allows the user to find the universe of all search results of message boards that show up if an individual were to search for that keyword within that subreddit. As mentioned in the body of the paper, I search three sets of keywords within three subreddits. For each message board found, I save the text of the prompt of the original post, as well as all top-level comments on that message board directly responding to that initial prompt.²²

After scraping the data from Reddit, I then proceed to conduct two preprocessing steps to help standardize the text data. I begin by filtering out comments that are not explicitly about debating/discussing UBI as a policy. I implement this at scale by making use of an LLM, specifically, Claude 3 Haiku. For each piece of text in my data, I prompt the LLM with the following question:

```
The following is text from a question asked on Reddit:
[Message Board Prompt]
The following is the answer to that question:
[Comment Text]
Based on this question/answer pair, tell me if the answer provides
reasons for/against UBI that appeal to questions about the likely
consequences of UBI for economic outcomes
```

I filter to include only those comments where the LLM responded with “yes”. From manual inspection of the first one hundred observations, I was unable to find an example where I disagreed with the classification of the LLM.

As a second preprocessing step, I prompt another LLM to give me a summary of the

²²Reddit allows for a recursive structure for comments, where some comments can be written as replies to other comments. For ease of analysis, and to ensure that the individual data points in my dataset are non-overlapping, I focus just on the top-level comments.

contents of each comment not filtered out from the first step, using a similar prompt, displayed below. This step ensures that the tone and style of each text example was relatively standardized.

```
The following is text from a question asked on Reddit:
[Message Board Prompt]
-----

The following is the answer to that question:
[Comment Text]
-----

Summarize the arguments the answer makes about the desirability of
universal basic income (UBI) in your own words. Describe only the
SUBSTANTIVE argument the author makes about UBI. If the answer
uses appeals to authority (e.g.,
referencing the arguments of some particular other individual,
remove these references from your summary). Similarly,
if the answer uses a historical example, summarize the insights
learned from that historical example, but do NOT reference
the historical example itself.
```

Given this final dataset of summarized answers, I then proceed to implement my clustering algorithm on the resulting text data. As described in the body, I begin by mapping each text summary into an embedding using the all-MiniLM-L6-v2 model from the Sentence-Transformers Python library. I then dimension reduce the resulting embeddings using the UMAP algorithm with the following hyperparameters: `n_neighbors=15`, `n_components=3`, `min_dist=0.1`, `metric=cosine`. This maps each text example into a vector in $\mathbb{R}^3$. Finally, I use a Gaussian Mixture Model with 20 mixture components to cluster the resulting dimension-reduced vectors.

Two topics from my topic analysis (18, 19) are general comments and do not necessarily debate aspects of UBI policy specifically. I thus exclude them when constructing the distribution of topics found in Figure 1.

A.2 Topics

Table 6: Representative Comments by Topic

Topic 1: Implementation and Government Capacity Concerns
The answer argues that universal basic income (UBI) is undesirable because the government does not have the financial resources to provide a \$1,000 check to every adult. The author expresses skepticism that the government can effectively “take care of everybody” through such a policy. The answer argues that while universal basic income (UBI) is a simple way to redistribute wealth and stimulate the economy, it has limitations. The main argument is that UBI can only be additive to existing social programs, and that many adults who require those programs make poor choices. Providing them with additional money, such as \$1,000, may not address their underlying needs, like ensuring they have access to basic necessities like food. The answer suggests that UBI may simply expand the group of people dependent on government assistance, while reducing the pool of people contributing to the operation of this additional bureaucracy. The answer raises concerns about the feasibility of funding a universal basic income (UBI) program, given the challenges in funding existing social security programs. The answer suggests that implementing a UBI could have negative impacts on the overall GDP, implying that the economic consequences of a UBI may outweigh its potential benefits. The answer raises concerns about the potential risks of a universal basic income (UBI) system. It suggests that a UBI could give the government, particularly a leader like Donald Trump, significant control over the main income source for a large number of people. Additionally, the answer expresses concern that a UBI system could become too expensive, and points to the poor outcomes of austerity measures in Europe as an example of the potential challenges of managing the costs of such a system. The answer argues that while universal basic income (UBI) is a nice idea in theory, it would be logistically difficult to implement in practice. The main concerns are that corruption would find a way to make UBI harmful for everyone except the rich, and that the downsides of UBI would outweigh any potential benefits, leaving people only slightly better off at best.
Topic 2: Social Stratification and Work Ethic
The answer suggests that the wealthy and middle class may fear receiving universal basic income because it would undermine the value they place on their jobs and the effort they have put into obtaining them. The implication is that UBI could be seen as devaluing the work they have done to secure their current employment and income level. The answer argues that a universal basic income would lead to a larger population of homeless drug addicts and elderly people who never worked or moved out of their parents’ homes. The answer raises a concern that universal basic income (UBI) could create a permanent underclass of citizens if there are not enough jobs for everyone. The concern is that UBI could lead to a “have and have not” society, where some people rely on the UBI payment while others have jobs and higher incomes. The answer suggests that while universal basic income could be a privilege of citizenship, it may also create a new poor class. The answer argues against the desirability of universal basic income (UBI) in the United States. The key argument is that not all workers are equally productive or willing to work hard, and that providing a basic income would reward those who are unwilling to work diligently. The answer contends that as long as the economy is strong, there will be jobs available for good, hardworking workers, and that those who are unproductive or unwilling to work do not deserve a basic income or unemployment benefits.

Continued on next page

Table 6 (continued)

Topic 3: Automation and Job Displacement

UBI would provide a safety net to bridge the gap for people who are not paid enough through their jobs to live off of, even for necessary jobs that cannot be replaced by robots.

UBI would accelerate automation and self-service, as businesses would seek to reduce labor costs. The overall conclusion is that UBI would result in people being paid less in real terms and still unable to afford basic necessities.

As technology improves, machines can do more work, reducing the need for human labor. UBI could provide a basic standard of living, allowing people to pursue activities they enjoy rather than being forced to work.

It provides a buffer against job losses due to automation, which is making the job market more competitive. The answer acknowledges there are challenges in implementing UBI, such as accounting for regional differences in cost of living and convincing people to support providing “free money.” It also notes that while UBI may be expensive, it could decrease strain on other public systems.

Initiatives like Mike Rowe’s non-profit that trains people for trades work could complement UBI, as certain jobs like plumbing are less likely to be automated in the near future compared to other professions.

Topic 4: Future Necessity Due to Technological Change

The answer suggests that universal basic income (UBI) may become necessary in the future due to the potential challenges posed by automation. The key argument is that if automation is not properly regulated, it could lead to significant societal challenges, and UBI may be a potential solution to assist with these upcoming challenges.

The answer suggests that while the author currently does not support a universal basic income (UBI), they foresee a future need for it due to the increasing rate of automation and AI. The author’s argument is that as automation and AI continue to advance, there may be a growing need for a UBI to address the potential job displacement and economic disruption caused by these technological changes.

The answer suggests that while universal basic income (UBI) likely has significant problems, it may be necessary as a response to increasing automation in society, as the author cannot think of any other way for society to move forward in that context.

The author anticipates that some older people may view UBI as a “grab to the lazy millennial” and argues that this perception is misguided. The author contends that millennials, including the author, are not lazy and work hard (e.g., the author works six days a week). The author suggests that UBI is necessary due to automation replacing many jobs in the future, rather than being a form of charity.

The answer suggests that universal basic income (UBI) would be desirable only if people’s current employment ended due to technological changes, such as the introduction of autonomous vehicles replacing semi-truck drivers. The argument is that UBI could help support those who lose their jobs due to such technological advancements.

Topic 5: Inflation and Price Increases

The answer presents two potential scenarios for implementing universal basic income (UBI). Best case scenario: people continue working to earn more than the UBI amount, leading to labor shortages in minimum wage jobs and increased automation of those positions. Worst case scenario: either most people stop working completely, or those still working demand higher wages due to the low labor supply. This could cause significant inflation and devaluation of the currency. The answer argues that Universal Basic Income (UBI) would lead to increased prices, such as a burger meal at McDonald’s costing \$30, and therefore would not help raise people out of poverty.

Continued on next page

Table 6 (continued)

The answer argues that for a universal basic income (UBI) to be a good idea, it must be accompanied by fixed and controlled rates for basic necessities such as housing, utilities, and food. This is to prevent further inflation that could negate the benefits of the UBI.

The answer argues that universal basic income (UBI) would lead to inflation, which would widen the gap between the normal and extremely rich by dragging down the low and middle class. Instead, the answer suggests that the focus should be on finding those truly in need and motivating those who can work but do not.

The answer argues that advocates of universal basic income (UBI) do not understand the concepts of inflation and supply and demand, which would undermine the desirability of UBI.

Topic 6: Cost and Funding Feasibility

The answer suggests that implementing a universal basic income (UBI) of \$1,000 per month for all adults in the United States would cost twice the total annual tax revenue of the country, requiring a tripling of all federal taxes to fund it. This indicates that the proposed UBI program would be prohibitively expensive and financially infeasible for the U.S. government to implement.

The alternative of funding UBI by taxing all 158 million working Americans would require an additional \$329 per month in taxes per person to cover the \$1,000 monthly UBI payment, which the author argues is not viable. The answer concludes by questioning how UBI could be practically implemented and funded on a large scale, given the financial challenges outlined.

The answer argues that a universal basic income (UBI) of \$1,000 per month in the US would cost over \$4 trillion per year, effectively doubling the size of the federal budget. Given the author's high income, they expect their taxes would increase significantly more than double under a progressive tax system to fund such a UBI program.

The scale of wealth required to provide a meaningful UBI payment to all Americans is extremely large. For example, confiscating all of Jeff Bezos' wealth would only provide a one-time payment of \$600 per person, whereas most UBI proposals aim for a monthly or bimonthly payment. Providing a more substantial UBI payment of \$100 per month would require producing enough money to increase the number of billionaires in the US by 40 per year.

The answer argues that even without tax increases, there would be enough money to pay a \$1,000 per month universal basic income (UBI) to 60% of U.S. households, which would cost around \$925 billion per year. The author points out that the U.S. government has previously spent similar or larger amounts, such as \$693 billion on the military in one year and \$700 billion to bail out Wall Street in 2008. The author suggests that forcing companies to pay their fair share and closing tax/employment loopholes could further increase the available budget for a UBI program.

Topic 7: Social Benefits and Welfare System Reform

The answer argues that a universal basic income would make education more accessible and remove the necessity of work for basic survival, thereby reducing bureaucratic work.

The answer argues that a universal basic income could reduce financial stress, leading to better mental health and reduced healthcare costs.

Conditional benefits, such as requiring training and job interviews, can provide valuable skills and opportunities for the unemployed, which would be lost with a universal income.

The answer argues that a universal basic income would create hundreds of thousands of local jobs. The reasoning is that having more money in people's hands would lead to increased local spending and economic activity, which would in turn create more jobs in local communities.

Continued on next page

Table 6 (continued)

The answer argues that universal basic income could be desirable in the aspect of supporting families, as it would allow a stay-at-home parent to provide better emotional support and upbringing for children. This, in turn, could potentially reduce childcare costs and lower crime rates by empowering parents to raise their children more effectively.

Topic 8: COVID-19 Crisis Response

The answer argues that the coronavirus crisis is unlikely to lead to the implementation of universal basic income (UBI) based on the difficulties and delays in distributing the \$1,200 government payments, which suggests the government would struggle to effectively implement a more comprehensive UBI program.

The answer argues that universal basic income (UBI) may not be a desirable policy in the context of an economic crisis like the coronavirus fallout, because UBI has to be funded through taxes, and if the economy experiences a massive downturn, the tax revenue required to sustain UBI may not be available.

The answer argues that while universal basic income (UBI) could be a good short-term solution to address the economic fallout from the coronavirus pandemic, more permanent structural changes to the economic system are needed. The key point is that UBI is seen as a quick fix, but the underlying issues require more fundamental reforms to the way the system operates.

The answer argues that universal basic income (UBI) would be an effective policy response to the economic fallout from the coronavirus pandemic, in contrast to the “free market pull yourself up by your bootstraps system” favored by conservative Republicans. The author suggests that the current crisis demonstrates the need for social safety nets and that UBI would “work really well” in addressing the economic challenges posed by the pandemic.

The answer argues that universal basic income (UBI) would be a desirable policy response to the economic fallout from the coronavirus pandemic. The key substantive argument is that UBI can provide a temporary injection of money directly to people, rather than to corporations, which the author suggests would be a more effective way to address the economic crisis.

Topic 9: Work Disincentives and Economic Collapse

The answer argues that universal basic income is undesirable because paying people to do nothing will result in a large number of people not working, while taking a significant amount of money from others to fund the UBI program.

The answer argues that universal basic income is not desirable because it is dependent on a reliable source of funding, such as oil revenue. When that funding source dries up, as happened in Venezuela, there would be no money to continue providing the basic income payments, which could lead to negative consequences.

The answer argues that the adoption of a universal basic income system would have a negative impact on the global economy. It suggests that a substantial number of people would stop working or work less, forcing employers to increase wages substantially. This would lead to the collapse of many businesses, reducing the supply of goods and services, including essential items like food. The decreased supply coupled with increased costs would trigger global inflation and a cycle of stagflation, similar to the experiences of some countries in the 1970s and 1980s. The answer suggests that this would result in the wiping out of life savings and potentially lead to a world war.

The answer suggests that the introduction of a universal basic income would lead to the collapse of the economy, and therefore the author would move away before that happens.

The answer argues that universal basic income would be pointless in a capitalist society, as it would likely just lead to increased inflation rather than meaningfully improving people’s economic situations.

Topic 10: Inflation and Purchasing Power

Continued on next page

Table 6 (continued)

In a capitalist economy, prices are set based on what customers are willing to pay. Increasing the amount of money customers have (through UBI) would lead to prices rising, negating the intended benefits of UBI.

UBI is likely to cause inflation, similar to the effects of raising the minimum wage.

This could lead to a cycle where the UBI amount is no longer sufficient to cover the increased cost of living, requiring the UBI to be increased, which would then further drive up inflation.

As UBI is increased over time to keep up with rising prices, this further fuels the inflationary cycle, making the UBI less effective at improving quality of life.

One of the major problems with UBI is not inflation, but rather that the options are either rampant inflation or raising taxes to the point where the average person is paying more than they receive.

Topic 11: Tax Evasion and Wealth Redistribution

The answer argues that the Panama Papers demonstrate that rich people will attempt to hide their money from taxation, rather than accept higher tax rates to fund a universal basic income. The implication is that a universal basic income may not be financially feasible if the wealthy can avoid contributing through tax evasion.

The answer argues that the wealth hidden by the global elite, as revealed by the Panama Papers, could easily fund a universal basic income (UBI) program in the United States. It estimates that a modest UBI of \$2,000 per year would cost around \$563 billion annually, while a more substantial UBI of \$12,000 per year would cost around \$3.6 trillion. The answer contends that these costs are reasonable when considered in the context of the trillions of dollars in untaxed assets held by the global elite, which the author claims have been “stolen” from the public.

Even if the funds were all legitimately earned, the estimated \$80 billion in global tax evasion revealed would only provide a UBI of around \$11 per person, which is unrealistic to support a meaningful UBI program. The answer suggests that university payments may be a better approach than UBI, but does not provide a substantive argument in favor of that alternative.

The answer argues that the claim that the Panama Papers prove we can afford a universal basic income is not substantiated. Specifically, the answer states that the article does not use any numbers from the Panama Papers, but rather relies on a 2012 report that estimates governments miss out on \$280 billion in tax revenue. However, the answer contends that this amount is not sufficient to cover the costs of a universal basic income program.

The answer argues that the Panama Papers scandal reveals that there is significant wealth being hoarded by large corporations and billionaires, who are able to avoid paying their fair share of taxes. The author suggests that if this hoarded wealth were instead taxed and circulated back into the economy, it could provide the resources necessary to fund a universal basic income program, which would greatly stimulate the economy and benefit the majority of the population.

Topic 12: Personal Financial Impact and Adequacy

The answer suggests that a \$13,000 universal basic income would be desirable, as it would provide additional income that could be saved or spent, thereby contributing to the economy. The author indicates that this amount would represent a significant portion of their overall income, implying that the UBI would be a meaningful supplement.

The answer suggests that a Universal Basic Income of \$1,000 per month would be desirable because it would cover essential expenses such as rent, food, and gas. The implication is that a UBI at this level would provide a basic standard of living and help meet fundamental needs.

Continued on next page

Table 6 (continued)

The answer suggests that a \$1,000 per month universal basic income would lead to an increase in monthly spending of approximately \$800. This implies that the author believes a UBI of this magnitude would be desirable, as it would allow for a significant increase in discretionary spending.

The answer suggests that a \$1,000 per month universal basic income would allow the author to quit working and live off the UBI payment and their savings.

The answer suggests that a universal basic income (UBI) of \$1,000 per month would be more desirable than the person's current financial situation, and they would therefore accept it.

Topic 13: Housing Costs and Rent Inflation

The median income house no longer costs the same as a person's annual salary.

The answer argues that universal basic income (UBI) would not be desirable because it would lead to higher housing costs, as landlords and real estate companies would quickly raise rents and home prices to absorb the additional income from UBI. The key points are that housing is essential, leases must be renewed annually, housing prices can adjust quickly to increased disposable income, many housing markets are constrained by local zoning laws, and UBI does not make construction cheaper.

The answer suggests that the desirability of universal basic income (UBI) is limited because landlords would likely increase rents to capture the additional income, requiring people to work long hours to avoid starvation despite the UBI.

The answer suggests that universal basic income (UBI) would not be a good idea without accompanying price caps on living expenses. The argument is that without such price caps, landlords, businesses, and other providers of essential goods and services could simply raise prices to absorb the additional income from UBI, negating its intended benefits.

Concerns about the poor misusing UBI funds could be addressed by converting a portion of UBI payments into universal basic benefits (UBB) such as rent vouchers, though this adds complexity.

Topic 14: Generational Differences in Support

The answer suggests that economic anxiety among younger Americans contributes to their support for universal basic income (UBI). The survey indicates that a majority of young adults are worried about their ability to afford unexpected expenses, which may drive their desire for an expanded social safety net, including UBI.

The answer argues that older voters are less likely to support universal basic income (UBI) because they already receive a form of UBI in the form of social security payments, which they claim to have contributed to throughout their lives. However, the author suggests that the social security "bank account" is not actually empty, but rather contains a large amount of IOU notes totaling \$22 trillion, implying that older voters are already receiving a form of UBI and thus are less inclined to support a new UBI program.

The answer argues that younger voters support universal basic income (UBI) because the current economic system has not been beneficial for them, with highly educated millennials struggling to find secure employment and financial stability. The author, who is 61 years old, supports UBI despite being part of the older generation, recognizing that the policies favored by baby boomers have not served the interests of younger generations. The author acknowledges that baby boomers may be indifferent to issues like climate change since they will not experience the worst effects, and argues that younger generations should have a greater say in shaping policies that will impact their future.

Continued on next page

Table 6 (continued)

The answer argues that while young Americans tend to support universal basic income and public healthcare, they are less likely to vote for such policies as they get older and their economic circumstances change. The author suggests that as young people pay off student loans, establish careers, and acquire assets like homes, they become more inclined to favor the status quo over radical policy changes, even if those changes would benefit them in the present.

The answer argues that younger voters are more in favor of universal basic income (UBI) compared to older voters because younger people are facing more economic challenges, such as underemployment, high debt, stagnant wages, and skyrocketing housing prices. The answer suggests that the older generation has not fully recognized the changed economic conditions that younger people are experiencing, and that younger people do not believe that hard work alone will lead to good pay, which is their lived experience.

Topic 15: Andrew Yang Campaign and Implementation Questions

The answer argues that while the commenter plans to vote for Andrew Yang due to his proposal for universal basic income (UBI), the commenter is concerned that UBI alone may not be effective in addressing the issue of high rents. The commenter suggests that in addition to UBI, there should be policies to limit what landlords can charge for rent, similar to rent control laws in places like Berlin, Germany. The commenter believes that without such rent control measures, landlords may simply raise rents to offset the additional income from UBI, negating its intended benefits.

The answer expresses support for some form of universal basic income (UBI), but questions the rationale of singling out the tech industry to fund it.

The answer expresses a positive opinion on universal basic income (UBI), describing it as “fantastic.” However, the answer also raises the question of how to pay for it, indicating a potential concern about the feasibility or implementation of UBI.

The answer suggests that the desirability of universal basic income (UBI) is questionable, as the author was banned from a subreddit dedicated to UBI for asking how much it would cost and how it would be paid for. This implies that the author believes the feasibility and affordability of UBI are important considerations that are not being adequately addressed.

The answer expresses a mixed view on universal basic income (UBI). It acknowledges the potential benefit of a base-level salary to ensure survival, but also notes that the funding for such a program would need to come from somewhere, implying potential challenges or tradeoffs.

Topic 16: UBI Amount and Work Incentives Design

The answer suggests that under a UBI system, everyone would receive a fixed amount of income, which would only be enough to cover the cost of a bedroom in a shared apartment. The author argues that this level of UBI would not provide for any additional welfare or other benefits, and that individuals would need to work in order to afford a higher standard of living beyond the basic level provided by the UBI.

Implementing a costly UBI system would deplete the wealth that working people have built, undermining the goal of supporting the poor. The answer concludes that the current progressive tax system is a better approach than UBI for supporting everyone while maintaining incentives to work.

Instead of implementing UBI, it would be better to restructure the tax system to make it cheaper for businesses to employ low-income workers, rather than making it more expensive through measures like minimum wage hikes and mandatory benefits.

Continued on next page

Table 6 (continued)

UBI should not disincentivize work. The answer suggests a Negative Income Tax (NIT) system, where the tax rate is a continuous function of income, as a better approach to avoid discontinuities that could create undesirable tax schemes. A potential compromise could be a UBI that supplements a minimum wage, funded by a corporate profit tax.

Topic 17: Alternative Approaches and Complementary Policies

The answer argues that universal basic income (UBI) could be less costly than the current social welfare system, as it would reduce the burden of expensive medical bills for the homeless and uninsured. It also suggests that UBI could increase demand for various goods and services, leading to job growth, and enable more people to afford higher education, generating increased tax revenue over their lifetimes.

The answer argues that instead of providing a universal basic income (UBI), it would be more desirable to provide universal basic utilities, such as a free basic level of electricity and clean water. The key argument is that providing these essential utilities directly would be more beneficial to people's needs than simply giving them cash through a UBI.

The answer suggests that universal basic income (UBI) is a good idea in principle, as it is a progressive policy. However, the author argues that for UBI to be effective, policymakers would need to regularly increase the UBI amount to match the cost of living, and they would need to avoid replacing other welfare programs with UBI, as that could lead to people starving.

The answer argues that universal basic income (UBI) is a good idea, but cautions that it should not be seen as a catch-all solution. The author believes that in addition to UBI, people should have access to healthcare, water, food, and housing if they cannot afford it themselves. UBI should be a way for people to have more freedom over their lives and work, rather than an excuse to ignore those other essential needs.

The answer argues that universal basic income (UBI) is desirable because it would provide financial support for two groups: 1) people who cannot work due to mental or physical issues, and 2) people who work multiple jobs but still cannot afford basic necessities. The answer suggests that if the government can afford payouts to large corporations and politicians, then it should also be able to provide basic financial support to these groups in need.

Topic 18: Freedom and Ideal Society Vision

It would provide people with enough money to meet their basic needs, such as food and shelter, which would allow them to focus on other activities beyond just survival.

A free and efficient labor market would provide a significant boost to the economy and people's lives.

This would improve the lives of those living in poverty and help young people get started.

To achieve an ideal society, there should be a single world government system that assigns the right people to the right jobs based on merit and performance, not popularity. This would remove inefficiencies from competitive ventures and the finance/sales/marketing industries.

This system would create incentives for businesses to pay their employees better or provide for the programs their employees rely on, and would make wealthy individuals more accountable for their actions.

Topic 19: Healthcare Integration and Tax Policy

The answer argues that the feasibility of implementing universal basic income (UBI) and Medicare for All in the United States would depend on whether the current tax system could support the costs of providing these benefits to the over 300 million citizens. It also raises the question of whether undocumented individuals in the U.S. trying to become legal would be eligible for these benefits.

Continued on next page

Table 6 (continued)

The answer argues that implementing a universal basic income (UBI) would require raising taxes to fund it, and that this could potentially increase taxes for lower-income families. The answer suggests that the author is interested in understanding how Andrew Yang, as a proponent of UBI, would address these concerns about increased taxes for lower-income individuals.

The feasibility of funding both UBI and universal healthcare, and what specific tax increases would be required to support these programs.

The author asks how the candidate plans to fund a Medicare for All proposal in addition to UBI.

The answer expresses support for Andrew Yang’s proposal of a universal basic income (UBI) plan, and inquires about how it would interact with the existing sales tax in Illinois. The key substantive argument made is that the author is a “huge fan” of Yang’s UBI plan and wants to see him on the debate stage, suggesting the author finds the UBI proposal desirable.

Topic 20: Government Control and Alternative Solutions

A UBI could increase government control and dependence, as the government would be providing the basic income and could potentially use that as leverage.

With a UBI, the government is essentially giving money to people, when it might be easier for the government to just not take the money in the first place.

The U.S. has already been paying for a form of UBI through corporate welfare, and the public would be better served by a direct UBI rather than these other expenditures.

While public infrastructure employment programs can complement UBI, a guaranteed jobs program alone is not a complete solution.

Nothing is free, and a UBI would require high taxes on everyone to fund it, without a clear benefit to those paying the taxes.


Notes: This table contains the various topics identified by the topic analysis, and provides the text summarizing the 5 Reddit comments most similar to the respective topic centroid.

A.3 Robustness

I show here that the main results in Figure 1 are not an artifact of my focus on text data from Reddit.com. In particular, I replicate my analysis pipeline on a sample scraped from transcripts of YouTube videos discussing UBI. One modification of my basic pipeline is that a typical video is longer—and thus contains more content—than a typical Reddit comment. As a result, I prompt the LLM to list all distinct points made in the YouTube video, and to summarize each point separately. The unit of analysis for the topic analysis is thus a sub-point of a video. My YouTube sample consists of 897 sub-points derived from 481 videos. In Figure 10, I show the distribution of topics, computed on the YouTube sample. While the

YouTube sample touches on similar issues as the Reddit sample, the ordering of priorities is different. YouTube transcripts are more likely to emphasize topics traditionally implicated in the Mirrlees model, namely redistributive motivations for tax/transfer reform and the effects of UBI on work. They are also more likely to discuss automation/AI concerns specifically, while being less likely to cite concerns about inflation or prices. In both populations, implementation details about cash versus in-kind transfers are moderately important.

Figure 10: Topic Distribution: YouTube Sample



Notes: This figure plots the proportion of sub-points corresponding to each high-level category of topic, using a sample of YouTube video transcripts.

B Welfare Weight Survey Experiment Details

In the survey experiment described in Section 2.2, respondents chose between 10 hypothetical pairs of transfers. In this appendix, I describe the details of how the characteristics of these transfers were randomized. Transfers varied based on the amount the recipient received, whether it was cash or in-kind, and the cost to the government to make the transfer. Recipients varied based on their post-tax income level, their labor-supply reduction if they received the transfer, and their susceptibility to AI automation risk.

All characteristics were randomized independently and identically across the 20 transfers described in the experiment. AI automation risk and the cash vs in-kind distinction, being binary variables, were randomized to take on each of their two values with equal probability. Transfer amounts, T , were randomized to take on values uniformly in the set $\{1000, 2000, \dots, 20000\}$. The cost to make the transfer C was then given by the formula $C = (1 + M) \times T$, where the multiplier M took on values uniformly in $\{0.05, 0.1, \dots, 1\}$.

C Welfare Weight Survey Experiment in Europe

In this appendix, I describe the additional survey experiment that I ran with a sample of 709 European citizens, whose results are reported in Section 2.2. This experiment replicates the design of the survey described in Appendix B with only minor adaptations to fit the European context (for example, displaying incomes in euros rather than in USD, asking about *country* rather than *state* of residence, etc.). I again recruited the sample using Prolific, filtering on subjects who stated that their country of residence was a European country.

D Additional Information About Counterfactuals

D.1 Extrapolating Utility Beyond the Estimation Domain

The B-spline approximations estimated in Section 3 are defined over finite domains, $[0, \bar{c}]$ for the consumption utility function u and $[0, \bar{\ell}]$ for the labor disutility function v , where $\bar{c}$ and $\bar{\ell}$ correspond to consumption and labor-supply levels at approximately the 99th percentile of the income distribution. Because behavior in the upper tail of the income distribution matters for counterfactuals, I extrapolate both utility functions beyond these domains using a constant relative risk aversion (CRRA) functional form. Specifically, at the boundary $\bar{c}$, I compute the implied coefficient of relative risk aversion $\sigma_u \equiv -\bar{c} \cdot u''(\bar{c})/u'(\bar{c})$ from the fitted B-spline, and for $c > \bar{c}$, I extend u according to $u(c) = u(\bar{c}) + \int_{\bar{c}}^c u'(\bar{c})(x/\bar{c})^{-\sigma_u} dx$. I extend

the labor disutility function v analogously. This approach keeps the utility functions well behaved in the tails while preserving the flexibility of the B-spline approximation for the majority of the income distribution.

D.2 Parameterization of the Tax Function

I parameterize the tax function T as a fourth-degree B-spline with 6 knots uniformly spanning the interval $[0, \bar{z}]$, where $\bar{z} = 300$ (in thousands of dollars) corresponds approximately to the 99th percentile of the US income distribution. For incomes above $\bar{z}$, the tax function is extended linearly according to $T(z) = T(\bar{z}) + T'(\bar{z})(z - \bar{z})$, so that the marginal tax rate is constant in the upper tail of the income distribution.

D.3 Definition of Status Quo Tax-and-Transfer Policy

The status quo tax-and-transfer policy T_0 combines federal income taxation with major means-tested transfer programs. I construct $T_0(z) = \tau(z) - B(z)$, where $\tau(z)$ represents federal income tax liability and $B(z)$ represents the total value of transfer benefits.

Federal income taxes are given by a piecewise linear function with marginal tax rates of 10%, 15%, 25%, 28%, 33%, 35%, and 39.6% applying to taxable income brackets with kinks at \$9,225, \$37,450, \$90,750, \$189,300, \$411,500, and \$413,200. This matches the 2015 federal income tax schedule for a single filer (Internal Revenue Service, 2015a).

The transfer system $B(z)$ aggregates seven major programs, calibrated to an average household consisting of a single parent with two children in 2015. The first program is the Earned Income Tax Credit (EITC), which phases in at a 40% rate between \$0 and \$13,870 of earned income, provides a maximum credit of \$5,548 between \$13,870 and \$18,110, and then phases out at a 21.06% rate until benefits reach zero at \$44,454 (Internal Revenue Service, 2015c). The second program is the Supplemental Nutrition Assistance Program (SNAP) for a three-person household, which provides a maximum annual benefit of \$6,132 (U.S. Department of Agriculture, Food and Nutrition Service, 2015). SNAP eligibility requires

income below 130% of the Federal Poverty Level (FPL), which is \$20,090 for a three-person family in 2015 (U.S. Department of Health and Human Services, 2015). The benefit equals the maximum minus 30% of net income, where net income is computed by subtracting 20% of earnings and a standard deduction of \$1,860 from gross income.

The Child Tax Credit (CTC) provides \$1,000 per child, for a maximum of \$2,000 for two children. The credit is partially refundable at a 15% rate on earnings above \$3,000, and phases out at 5% of income above \$75,000 (Internal Revenue Service, 2015b). Housing assistance takes the form of a Section 8 voucher, which I model assuming annual rent of \$12,000. The household pays 30% of its income toward rent, with the government covering the remainder, consistent with Housing Choice Voucher program rules (U.S. Department of Housing and Urban Development, 2015). Medicaid and CHIP provide health insurance coverage valued at \$4,986 annually for the parent if income is below 138% of FPL, and \$3,389 annually for each child if income is below 250% of FPL. These coverage values are based on average per-enrollee Medicaid expenditures reported by the Centers for Medicare and Medicaid Services (Centers for Medicare and Medicaid Services, Office of the Actuary, 2015). The 138% FPL threshold for parents reflects the Affordable Care Act Medicaid expansion, while the 250% threshold for children represents typical CHIP eligibility limits (Brooks et al., 2015). Temporary Assistance for Needy Families (TANF) provides a maximum cash benefit of \$6,624 at zero income, which phases out linearly to zero at \$18,000 of income. This represents an approximation of typical state TANF benefit schedules, which vary considerably across states (Azevedo-McCaffrey and Aguas, 2025). Finally, childcare subsidies cover a portion of childcare costs, which I assume total \$18,000 annually for two children, consistent with estimates of average childcare costs from Child Care Aware of America (Child Care Aware of America, 2015). The subsidy phases out linearly as income rises from \$0 to 200% of FPL, reflecting typical Child Care and Development Fund (CCDF) eligibility rules.

For each program, I apply empirically estimated take-up rates to account for incomplete program participation. The take-up rate for EITC is 78%, based on IRS estimates (Internal

Revenue Service, 2014). For SNAP, I use a take-up rate of 83%, based on USDA reports (Gray and Cunyningham, 2016). The Medicaid take-up rate of 92.8% comes from estimates by Kenney et al. (2012). The TANF take-up rate of 34% reflects the low participation rates documented in administrative data (Zedlewski and Holland, 2003). For housing vouchers, I use a take-up rate of 65%, consistent with estimates from HUD administrative data accounting for waitlist dynamics (of Housing et al., 2006). The take-up rates for CTC (85%) and childcare subsidies (15%) are approximations based on program participation data and eligibility estimates (Maag, 2018; Chien, 2016). The total expected transfer function $B(z)$ is the sum of expected benefits across all seven programs, where expected benefits equal statutory benefits multiplied by the relevant take-up rate.

Both $\tau(z)$ and $B(z)$ are approximated using fourth-degree B-spline functions. For the tax function, I fit a B-spline with two interior knots to minimize squared deviations from the true piecewise linear tax schedule, subject to the constraint that marginal tax rates are approximately constant in the top 20% of the income distribution. For the transfer function, I fit a B-spline with one interior knot, imposing that the marginal phase-out rate at zero income does not exceed 60% in absolute value and that the transfer function exhibits convexity. These spline approximations ensure that T_0 and its derivatives are smooth functions that can be readily incorporated into the computational optimization routines used in the counterfactual analysis. In Appendix D.5, I report results showing that my results do not seem to depend too strongly on exactly how the pre-existing tax-and-transfer function is defined, by constructing T_0 through an entirely different method.

D.4 AI Automation Risk

The specification of preferences over redistributive policy in the body of the paper required a specification of how exposed individuals at different income levels are to automation risk. This appendix describes the specification of this automation risk function. I use the results from Hampole et al. (2025), who estimate equilibrium employment effects from automation

technology, accounting for task- and job-level substitution patterns among workers. I use in particular their Figure 7, which estimates net effects for different income percentiles.


Given their specification, they are able to make inferences about *relative* effects of automation on employment changes, but their methodology is uninformative about the average *level* of employment effects. I therefore assume that the highest earners neither gain nor lose employment from automation technology and compute AI effects from that baseline. In their Figure 7, they report employment gains of +3% among highest earners and employment losses of -2% among lowest earners, in relative terms. I thus assume that employment risk among the lowest earners is 5% and 0% for high earners, and interpolate linearly by income quantile in the middle. Thus, for instance, the median earner faces a 2.5% employment risk. Due to the small magnitude of these effects, my end results are highly insensitive to the precise treatment of automation risk.

D.5 Robustness Checks

Here, I show that the main takeaways from my counterfactuals are not very sensitive to assumptions about behavioral responses to tax reform. I separately test both robustness to assumptions about elasticities of taxable income and assumptions about the pattern of income effects.

A first robustness check concerns a potential criticism of many estimates of elasticity of taxable income that they are too “short-run”. For instance, in models with human capital accumulation, it is often the case that a permanent increase in marginal tax rates (which are the relevant quantity for long-run tax reform) should induce a larger effect on labor supply than a temporary, one-year increase (Keane and Wasi, 2016; Best and Kleven, 2012). I test for robustness to this concern by re-running my entire estimation and counterfactual procedure, but using as inputs estimates of ETI that are three times larger than the ones used in the body of the paper. With such high elasticities, I find that the social planner’s policy is considerably more sensitive to the choice of penalty multiplier λ on labor-supply

Figure 11: Desirability and Popularity of Modified Social Planner’s Policies Under Alternative ETI Specification




Notes: This figure examines the welfare and popular support implications of adding a labor-supply penalty to the social planner’s optimization problem (horizontal axis), under an alternative ETI specification. Panel A shows how the social planner’s objective changes as the labor penalty increases. Panel B shows how popular support changes as a function of the labor penalty.

effects. However, the overall takeaway when considering the whole family of policies as a function of λ remains the same: the social planner’s preferred policy (at $\lambda = 0$) is unpopular, while popular policies involve mitigating the labor-supply impacts by delaying the claw-back of income support to higher income levels. Figures 11 and 12 are analogous to Figures 6 and 7, and look qualitatively very similar.

A second potential concern is that all but one of my MPE estimates use evidence from lottery players, who might be an especially strange subset of the US population. Unfortunately, I only have one additional estimate of MPE, which is the UBI field experiment of Vivalt et al. (2024). In my robustness check, I assume that the MPE estimates from this one data point are representative for the entire population instead, and assume that MPEs are constant, at the level found in this experiment. Under this assumption, I once again re-run my entire estimation and counterfactual procedure, with analogous counterfactual Figures 13 and 14. Once again, results look very similar to what I find in the body of the paper.


In the body of the paper, my main random coefficient utility model did not make any

Figure 12: Modified Social Planner's Policies Under Alternative ETI Specification



Notes: This figure plots various versions of the social planner's policy compared to the status quo tax-and-transfer system, under an alternative specification for the ETI. The dotted line represents the unconstrained social planner's policy. The dashed line represents the policy required to obtain majority support, while the solid blue line represents the policy required to obtain super-majority (more than two-thirds) support. Panel A plots post-tax income as a function of pre-tax income while Panel B plots implied marginal tax rates.


Figure 13: Desirability and Popularity of Modified Social Planner's Policies Under Alternative MPE Specification



Notes: This figure examines the welfare and popular support implications of adding a labor-supply penalty to the social planner's optimization problem (horizontal axis), under an alternative MPE specification. Panel A shows how the social planner's objective changes as the labor penalty increases. Panel B shows how popular support changes as a function of the labor penalty.

use of respondent-level observables, despite my finding that liberal and conservative survey respondents answer differently from one another. In this appendix, I thus consider an

Figure 14: Modified Social Planner's Policies Under Alternative MPE Specification




Notes: This figure plots various versions of the social planner's policy compared to the status quo tax-and-transfer system, under an alternative specification for the MPE. The dotted line represents the unconstrained social planner's policy. The dashed line represents the policy required to obtain majority support, while the solid blue line represents the policy required to obtain super-majority (more than two-thirds) support. Panel A plots post-tax income as a function of pre-tax income while Panel B plots implied marginal tax rates.

alternative specification where I estimate separate random coefficient distributions for the two subsamples of self-identified liberals and self-identified conservatives. I then re-compute popular support with this new random coefficient distribution. Figure 15 plots the analogue of Panel A of Figure 6, using this alternative random coefficient specification. I find that the popularity profile of the various policies is almost identical to the main specification in the paper. I do not reproduce Figure 7 here. Because in this robustness check, I only vary the model of policy preferences among the public, the underlying counterfactual policies are identical to the ones in the paper.

Because the existing system of government transfer programs is quite complex, I made a number of judgment calls when trying to construct what the status quo policy looks like. I probe for robustness to my modeling choices here by considering an alternative parameterization of the transfer function that targets marginal tax rates from Kosar and Moffitt (2017) rather than directly approximating the program-by-program transfer schedule described in Section D.3. Specifically, I target the mean effective marginal tax rates for

Figure 15: Popularity Under Alternative GSMWW Model



Notes: This figure shows how popular support (vertical axis) changes as the penalty on labor in the social planner’s optimization problem (horizontal axis) increases, under an alternative specification of policy preferences among the general public.

households receiving Medicaid, SNAP, TANF, and housing assistance, as reported in Table A8 of Kosar and Moffitt (2017) for married childless adults in 2007. Since these MTRs reflect the combined effect of transfer phase-outs and federal income taxation, I subtract the approximated federal tax schedule’s marginal rate from each target to isolate the transfer-induced MTR component. Let $\tau'(z)$ denote the marginal tax rate from the federal income tax approximation described in Section D.3, and let $\bar{m}_j^{\text{data}}$ denote the mean effective MTR from Table A8 for income band $j \in \{[0, 50\%], [50\%, 100\%], [100\%, 150\%], [150\%, 200\%]\}$ of the FPL. The target transfer MTR for band j is then $\bar{m}_j^{\text{target}} = \bar{m}_j^{\text{data}} - \bar{\tau}'_j$, where $\bar{\tau}'_j$ denotes the average federal marginal tax rate in band j .

I parameterize the alternative transfer function as a fourth-degree B-spline, $\hat{T}(z; \phi_T) \equiv \sum_{k=1}^{K_T+4} \omega_{Tk} \mathcal{B}_{Tk}(z)$, where the basis functions $\mathcal{B}_T = \{\mathcal{B}_{T1}, \dots, \mathcal{B}_{T, K_T+4}\}$ are determined by a knot vector κ_T of 13 knots that partitions the support into $K_T = 12$ sub-intervals.²³ The knot placement is non-uniform, with greater density in the region where the MTR peaks (around 50–100% FPL) and in the phase-out region (200–300% FPL) to ensure adequate flexibility where the function exhibits the most curvature. Using a fourth-degree spline guarantees that

²³The exact choice of knots is $\kappa_T = \{0\%, 25\%, 50\%, 75\%, 100\%, 125\%, 150\%, 200\%, 250\%, 255\%, 270\%, 285\%, 300\%\} \times \text{FPL}$.

the transfer function is thrice continuously differentiable, which ensures that the marginal tax rate $-\hat{T}'(z)$ and its derivative are both continuously differentiable—a desirable property for the numerical stability of the welfare calculations. Throughout, I set the FPL to \$26,650, the 2025 guideline for a family of three. Although Table A8 of Kosar and Moffitt (2017) reports rates for married childless adults, I interpret it as informative about the shape of effective marginal tax rates across the income distribution, and I express income bands relative to the poverty line for a family of three, the household type underlying the main specification, so that the two status quo specifications describe comparable households.

The B-spline coefficients are estimated by solving the following constrained optimization problem:

$$\begin{aligned}
\hat{\phi}_T = \underset{\phi_T}{\operatorname{argmin}} \quad & \sum_j \left(\bar{m}_j^{\text{model}}(\phi_T) - \bar{m}_j^{\text{target}} \right)^2 + \lambda_1 \int_0^{\bar{z}} \left[\hat{T}''(z; \phi_T) \right]^2 dz + \lambda_2 \int_0^{\bar{z}} \left[\hat{T}'''(z; \phi_T) \right]^2 dz \\
& + \lambda_3 \left[\hat{T}(\bar{z}; \phi_T) \right]^2 + \lambda_4 \left[\hat{T}'(\bar{z}; \phi_T) \right]^2 + \lambda_5 \left[\hat{T}''(\bar{z}; \phi_T) \right]^2 \\
\text{s.t.} \quad & -\hat{T}'(z; \phi_T) \geq 0 \quad \forall z \in [0, \bar{z}] \\
& -\hat{T}'(z; \phi_T) \leq 0.80 \quad \forall z \in [0, \bar{z}] \\
& -\hat{T}'(z; \phi_T) \geq 0.40 \quad \forall z \in [0, 0.5 \times \text{FPL}] \\
& -\hat{T}'(z; \phi_T) \text{ is single-peaked,}
\end{aligned} \tag{16}$$

where $\bar{m}_j^{\text{model}}(\phi_T) \equiv \frac{1}{z_j^{\text{max}} - z_j^{\text{min}}} \int_{z_j^{\text{min}}}^{z_j^{\text{max}}} \left[-\hat{T}'(z; \phi_T) \right] dz$ is the model-implied average MTR in band j , and $\bar{z} = 3 \times \text{FPL} = \$79,950$ is the income level at which transfers phase out completely. The smoothness penalties (weighted by λ_1 and λ_2) discourage excessive curvature in the MTR schedule, while the endpoint penalties (weighted by λ_3 , λ_4 , and λ_5) ensure that the transfer function, its first derivative, and its second derivative all converge smoothly to zero at the phase-out threshold.

For computational tractability, I impose three constraints on the shape of the transfer schedule. First, the upper bound of 80% prevents implausibly high effective marginal tax rates that would imply near-confiscatory phase-out rates, as this makes counterfactual ex-


trapolations non-smooth and therefore difficult to optimize. Second, the lower bound of 40% in the lowest income band ensures that transfer phase-outs begin immediately as income rises from zero, consistent with the empirical evidence that low-income households face substantial work disincentives even at very low earnings levels. Finally, the single-peakedness constraint—implemented via monotonicity penalties on the MTR before and after its maximum—ensures that the effective marginal tax rate rises smoothly to a peak (occurring around 75–100% FPL) before declining monotonically thereafter. This shape is consistent with the empirical pattern documented in Kosar and Moffitt (2017), while imposing additional smoothness constraints to aid in numerical tractability.

The resulting marginal tax rate schedule from this procedure looks similar to the hand-constructed one I use in the main analyses of the paper, with similar levels of generosity for the poorest households.²⁴ The overall schedule of marginal tax rates differs slightly. This alternative transfer specification implies slightly smaller marginal tax rates for the very poorest Americans, compared to the main specification, while implying slightly larger marginal tax rates for Americans earning roughly \$10,000–\$25,000/year.

A consequence of these differences in marginal tax rates is that the status quo tax policy looks worse under this alternative specification compared to the main specification. As a result, while the social planner’s proposed reform still is unpopular, it is less unpopular than in the main specification. About 40% of citizens support it without modifications. Under this new specification, it also turns out that simply replacing in-kind transfers with equivalent cash transfers may already be appealing to a majority of the general public. Nonetheless, similar broad lessons continue to apply: to attract super-majority support, a transfer policy that phases out more slowly than the social planner would like is effective at gaining supporters while not sacrificing much on the planner’s objective. These results are summarized in Figures 16 and 17, which, once again, respectively replicate Figures 7 and 6 in the main analysis.


²⁴Notably, I construct the above transfer function using only implied marginal tax rates, yet continue to find similar transfer *levels*.

Figure 16: Modified Social Planner's Policies Under Alternative Status Quo Transfer Specification



Notes: This figure plots various versions of the social planner's policy compared to the status quo tax-and-transfer system, under an alternative specification for the status quo tax-and-transfer system. The dotted line represents the unconstrained social planner's policy. The dashed line represents the policy required to obtain majority support, while the solid blue line represents the policy required to obtain super-majority (more than two-thirds) support. Panel A plots post-tax income as a function of pre-tax income while Panel B plots implied marginal tax rates.

Figure 17: Popularity Under Alternative Status Quo Transfer Specification



Notes: This figure shows how popular support (vertical axis) changes as the penalty on labor in the social planner's optimization problem (horizontal axis) increases, under an alternative specification of the status quo tax-and-transfer system.

E Maximizing The Public's Social Welfare Function

Throughout the body of the paper, my counterfactuals took the perspective of the social planner and asked how she would set policy in the face of a popular support constraint. The social planner's objective function is constructed by taking a weighted average of household utilities, where each household's utility is empirically derived from revealed preference analysis of their labor/leisure tradeoffs. A potential drawback of this view is that it only allows for narrow self-interest to enter into the social planner's calculus. This rules out the possibility that the ethical views of citizens enter into the social planner's utility. As argued by Eden and Piacquadio (2025), if these ethical views are taken into consideration, under reasonable calibrations, they will likely dominate the social welfare calculation. Under this view, the normatively relevant social welfare function is arguably not the social planner's, but rather, the preferences of the average citizen.


While not the main focus of this paper, I show in this appendix that the empirical estimates of this paper facilitate an analysis under this alternative notion of social welfare. Specifically, for each citizen i in society, let $\nu_i(T)$ be that citizen's utility in evaluating policy T , as defined in Equations (13) and (14). I normalize preferences such that all citizens value giving a dollar to a random member of society equally. Specifically, let $\bar{\omega}_i \equiv \int_{\theta} \omega_i(\theta) dF_{\theta}(\theta)$. Define the public's social welfare function over policies as the average utility from that policy, i.e., $\text{SWF}_{\text{pub}}(T) \equiv \mathbb{E}[\nu_i(T)/\bar{\omega}_i]$, where the expectation is taken over the random coefficients drawn from my survey estimates. Then the optimal policy with respect to this welfare function is given by

$$\begin{aligned} \max_T \text{SWF}_{\text{pub}}(T) \quad \text{s.t.} \quad & \int_{\theta} T(z(T, \theta)) dF_{\theta}(\theta) = \int_{\theta} T_0(z(T_0, \theta)) dF_{\theta}(\theta), \\ & \frac{\int_{\theta} (T'(z(T_0, \theta)) - T'_0(z(T_0, \theta)))^2 z(T_0, \theta)^2 dF_{\theta}(\theta)}{\int_{\theta} z(T_0, \theta)^2 dF_{\theta}(\theta)} \leq 0.01. \end{aligned} \tag{17}$$

The results from solving this alternative optimal policy problem are plotted in Figure 18. While the body of the paper shows that the social planner largely prioritizes helping the

poorest members of society, the public finds reducing work disincentives to be a higher-order concern. As a result, the public's preferred policy involves, if anything, slightly reducing generosity to the poorest members of society as well as a slower phase-out of benefits in order to substantially decrease implicit marginal tax rates at the bottom of the income distribution.

Figure 18: The Public's Optimal Policy vs The Mirrleesian Optimum




Notes: This figure plots the Mirrleesian optimal policy and the policy that maximizes the public's preferences, compared to the status quo tax-and-transfer system. The dotted line represents the Mirrleesian optimum, while the solid line represents the public's preferences. Panel A plots post-tax income as a function of pre-tax income while Panel B plots implied marginal tax rates.

F Policy Preference Survey Experiment Details

This appendix provides additional details about the survey experiment described in Section 4.1. The sample was recruited from the same population as the surveys designed to elicit welfare weights. After standard demographic questions, the main body of the survey consists of asking respondents three binary yes/no questions about whether, given a description of a hypothetical tax/transfer reform, they would support that policy. As described in more detail in Section 4.1, the three policies described are, respectively: (A) the optimal policy from the point of view of the Mirrleesian social planner, (B) the modified policy that achieves

majority support, and (*C*) the modified policy that achieves majority support and is framed to address concerns about AI. All respondents are asked about all three of the policies, but to mitigate concerns about order effects, the order in which the three policies are presented is randomized. The three orderings I present are $A \rightarrow B \rightarrow C$, $B \rightarrow C \rightarrow A$, and $C \rightarrow A \rightarrow B$, consistent with a cyclic Latin square design.

Figure 19: Support for Various Policies, By Order



Notes: This figure plots the proportion of respondents who choose to support each of the three policies, separated by the order in which that policy was presented. Error bars represent 95% confidence intervals.

Figure 19 shows how the proportion of people supporting each policy varies depending on whether the policy was shown first, second, or third. It is clear from this figure that order effects exist. In particular, the first policy shown to each respondent is more likely to be rated favorably. Nonetheless, it is clear that holding order of presentation fixed, policies *B* and *C* are more popular than policy *A*, in line with the predictions implied by the estimated welfare weights as well as the results presented in the body of the paper. In particular, even focusing on choices made by respondents over the first policy they encounter (this mimics the result from a between-subject design where each subject is only shown a single one of the three policies at random), policy *A* is statistically significantly less popular than policy *C*.