1 Introduction

A standard approach to optimal income taxation theory going back to Mirrlees (1971) posits a utilitarian social planner who is able to set tax policy at will in order to maximize a weighted average of utility within the population. This framework has been highly productive in generating insights about the tradeoffs involved when the government wishes to raise revenue and redistribute by taxing income. In the real world, however, especially in democratic countries, governments may not be free to set whatever tax policy it wants, and may be constrained by popular opinion. The goal of of this project is to empirically estimate the popular opinion constraint the government must contend with when setting its tax-and-transfer policy and to use these empirical estimates to calculate how public policy is distorted away from the standard Mirleesian second best optimum due to these constraints. Previously, I estimated popular opinions about income distribution will be measured via a survey experiment among members of the US voting public. This PAP describes the same design, but using a sample of participants sampled across Europe instead.

2 Survey Design

The survey consists of three parts. The first part asks about basic demographic information (gender, age, marital status, dependents, ethnicity/race, citizenship status, state of residence, education, employment status, income, political alignment). The second part, which is the meat of the survey, contains ten hypothetical questions. Each question will present the survey respondent with a hypothetical choice over two transfers. These transfers differ both in terms of the size of the transfer, the cost to the government to implement the transfer, and whether the transfer is cash or in-kind. Transfers will also differ in terms of the recipient's income, labor supply response, and whether or not their job is at risk of being automated due to AI. These different characteristics of the transfer will all be randomized from question to question, and how survey-takers respond to this variation will help to identify preferences over redistribution.

In the final part of the survey, I ask respondents about whether or not they supported Andrew Yang's universal basic income proposal. I then ask to what extent redistribution concerns, automation, cost, inflation, disincentives for work, and crowd-out of pre-existing programs contributed to their answer.

The data from this survey is of the form $\{\boldsymbol{Z}_i, \{Y_{iq}, \boldsymbol{X}_{iq1}, \boldsymbol{X}_{iq2}\}_{q=1}^{1}0\}_{i=1}^{N}$, where i indexes individuals, and q indexes questions, $\boldsymbol{Z}_i$ represents individual-level covariates, Y_{iq} is an indicator for whether or not the survey-taker preferred the first transfer in question q, and

$\boldsymbol{X}_{iqt}$ represents the characteristics of transfer $t \in \{1, 2\}$ in question q of survey-taker i. Let $\Delta\boldsymbol{X}_{iq} \equiv \boldsymbol{X}_{iq1} - \boldsymbol{X}_{iq2}$ represent the difference between the transfer characteristic in the first compared to the second transfer for question q of respondent i.

3 Reduced Form Analysis

Before estimating the structural model that I will use for the policy counterfactual stage of the project, I will summarize my results using a number of reduced form regressions to get a sense for the variation in the data. These reduced-form regressions will be aimed at documenting answers to the following question: 1) what aspects of redistribution do people most care about, and 2) are there systematic patterns in how these relative preferences differ by observables?

To answer the first question, I will estimate a simple linear probability model of the form:

$$Y_{iq} = \alpha + \boldsymbol{\beta}'\Delta\mathbf{X}_{iq} + \boldsymbol{\gamma}'\mathbf{Z}_i + \varepsilon_{iq}$$

where i indexes individual, q indexes question, $\mathbf{X}_{iq}$ denotes the characteristics of the redistribution situation described in question q to household i, and $\mathbf{Z}_i$ indexes household characteristics. Standard errors will be clustered at the household level.

To answer the second question, I will further discretize $\mathbf{Z}_i$ such that for continuous variables, I split $\mathbf{Z}_i$ by above/below median and then regress

$$Y_{iq} = \alpha + \boldsymbol{\beta}'\Delta\mathbf{X}_{iq} + \boldsymbol{\gamma}'\mathbf{Z}_i + \boldsymbol{\delta}'(\Delta\mathbf{X}_{iq} \otimes \mathbf{Z}_i) + \varepsilon_{iq}$$

I will report the coefficient $\boldsymbol{\delta}$ as a matrix, with one dimension indexing household characteristic and the second dimension indexing question characteristic.

4 Structural Analysis

To motivate the structural analysis of the survey, I will first describe a theoretical framework for modeling how households think about redistribution. I assume that preferences for redistribution can be described via a generalized marginal social welfare weight framework.

In the framework, *voters*, indexed by i form preferences over public policies, which affect a distribution of *citizens* in society more broadly. To describe the heterogeneity amongst the *citizens*, let $\mathbf{W}$ be a vector of household characteristics which varies across citizens. A tax and transfer reform can be thought of as a pair of functions, $(\Delta T(\mathbf{W}), \Delta K(\mathbf{W}))$, which describes how much the (net) tax burden and in-kind transfer, given to a household with characteristic $\mathbf{W}$ changes.

I assume that the re-distributive preferences of survey respondents are separable across households. I also assume that households may dislike increasing government spending, holding all else fixed. Formally, the preferences of a given household i can be represented by a welfare function of the form

$$\mathcal{S}_i(\Delta T, \Delta K, \Delta G) = \int_{\boldsymbol{W}} V_i(\Delta T(\mathbf{W}), \Delta K(\mathbf{W}), \mathbf{W})\, \mathrm{d}F(\mathbf{W}) - \alpha_i \Delta G,$$

where ΔG is the change in the government's budget constraint. Separability is convenient here, because it implies that to learn about preferences for re-distributive policies in general, it suffices to study preferences over transfers to individual households in isolation.

To first order approximation, the social welfare function can thus be written as:

$$\mathcal{S}_i(\Delta T, \Delta K, \Delta G) \approx \int_{\boldsymbol{W}} \left[V_i(0,0,\mathbf{W}) + \frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta T} \Delta T(\mathbf{W}) + \frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta K} \Delta K(\mathbf{W}) \right] \mathrm{d}F(\mathbf{W}) - \alpha_i \Delta G.$$

Here, we can think of $\frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta T}$ as the generalized marginal social welfare weight (Saez and Stantcheva, 2016) on a household with characteristic $\mathbf{W}$ while $\frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta K} \big/ \frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta T}$ represents the relative efficiency of in-kind transfers compared to cash transfers. Denote $\frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta T}$ and $\frac{\partial V_i(0,0,\mathbf{W})}{\partial \Delta K}$ respectively as $U_{i0}(\mathbf{W})$ and $U_{i1}(\mathbf{W})$. I assume that for the range of policies I will be evaluating, V_i is approximately linear in its first two arguments, so to evaluate the share of people who prefer a given policy to the status quo, it suffices to estimate the functions $U_{i0}(\mathbf{W})$ and $U_{i1}(\mathbf{W})$.

Auxiliary evidence from a text analysis I conducted suggests that the three most salient features entering $\mathbf{W}$ are (Y, δ, A) where

1. Y is the status-quo household income
2. δ is the change in labor supply in response to the transfer.
3. A is an indicator for whether or not the household's current occupation choice is at risk of automation.

The survey experiment varies these three salient features of the household as well as the the amount transferred to the household (denoted τ), the in-kind status of the transfer (K), and the cost to make the transfer (M). Thus, $\mathbf{X}_{iq} = (Y_{iq}, \delta_{iq}, A_{iq}, \tau_{iq}, A_{iq}, M_{iq})$.

I parameterize the generalized marginal social welfare weights as

$$U_{i0}(Y,\delta,A) = \beta_{is}[\gamma_i^{-1}\exp(-\gamma_i Y) + \beta_{id}\delta + \beta_{ia}A],$$
$$U_{i1}(Y,\delta,A) = U_{i0}(Y,\delta,A) + \beta_{is}\beta_{ik}K,$$

and denote $U_i(Y,\delta,A,K) \equiv (1-K)U_{i0}(Y,\delta,A) + KU_{i1}(Y,\delta,A)$.

When responding to a survey, I assume that survey respondents make choices according to their preferences over redistribution, subject to an i.i.d. standard logistic error, that is, survey-taker i will choose to transfer to household 1 in the survey if and only if

$$F(\mathbf{X}_{iq1},\mathbf{X}_{iq2}) \equiv [U_i(Y_{iq1},\delta_{iq1},A_{iq1})\tau_{iq1} - \alpha_i M_{iq1}] - [U_i(Y_{iq2},\delta_{iq2},A_{iq2})\tau_{iq2} - \alpha_i M_{iq2}] \geq \varepsilon_{iq},$$

where $\varepsilon_{iq} \overset{i.i.d.}{\sim} \mathit{Logistic}(1)$. In other words, I assume that survey responses are described by a mixed-logit model, where the random coefficients component of the model corresponds to the true redistributive preferences of households in the sample. To aid tractability, I further assume that the distribution of random coefficients takes the form

$$\boldsymbol{\theta}_i \equiv \begin{pmatrix}\beta_{is}\\ \gamma_i\\ \beta_{id}\\ \beta_{ik}\\ \beta_{ia}\\ \alpha_i\end{pmatrix} = H(\boldsymbol{G}_i) \equiv \begin{pmatrix}\log(1+\exp(G_{i1}))\\ \log(1+\exp(G_{i2}))\\ G_{i3}\\ G_{i4}\\ G_{i5}\\ \log(1+\exp(G_{i6})),\end{pmatrix} \quad \text{where} \quad \boldsymbol{G}_i \sim \mathcal{N}(\boldsymbol{\mu},\boldsymbol{\Sigma}).$$

The transformation $\log(1+\exp)$ applied to β_{iy}, γ_i, and α_i ensures that these coefficients are positive, but I will be agnostic about the sign of the other coefficients.[1] I represent the variance matrix $\boldsymbol{\Sigma}$ via its lower Cholesky factor $\boldsymbol{L}$ such that $\boldsymbol{L}\boldsymbol{L}' = \boldsymbol{\Sigma}$. The structural primitives to be estimated are therefore $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

To derive an estimator for these primitives, I now derive the likelihood of the model. I begin by describing the likelihood the choices of a single individual, conditional on their survey questions, and conditional on knowing the draw of $\boldsymbol{\theta}_i$. This likelihood function is given by

$$\mathcal{L}(\boldsymbol{\theta}_i;\{Y_{iq},\boldsymbol{X}_{iq1},\boldsymbol{X}_{iq2}\}_{q=1}^{10}) \equiv \prod_{q=1}^{Q} Y_{iq}^{p_{iq}(\boldsymbol{\theta}_i;\{Y_{iq},\boldsymbol{X}_{iq1},\boldsymbol{X}_{iq2}\}_{q=1}^{10})}(1-Y_{iq})^{1-p_{iq}(\boldsymbol{\theta}_i;\{Y_{iq},\boldsymbol{X}_{iq1},\boldsymbol{X}_{iq2}\}_{q=1}^{10})}$$

[1] A somewhat more standard specification is to simply take any component that is positive to be log-normal. In analyzing a pilot version of the experiment, I found that the log-normal specification was numerically unstable, which motives my use of the $\log(1+\exp)$ transformation.

where the probability p_{iq} is given by the logit rule

$$p_{iq}(\boldsymbol{\theta}_i; \{Y_{iq}, \boldsymbol{X}_{iq1}, \boldsymbol{X}_{iq2}\}_{q=1}^{10}) \equiv \frac{1}{1 + \exp(-F(\boldsymbol{X}_{iq1}, \boldsymbol{X}_{iq2}))}.$$

The overall likelihood function integrates over the distribution of $\boldsymbol{\theta}_i$ and aggregates over individuals in the survey and is given by

$$\boldsymbol{L}(\boldsymbol{\mu}, \boldsymbol{L}) = \prod_{i=1}^{N} \mathbb{E}_{\boldsymbol{G}\sim\mathcal{N}(\boldsymbol{\mu},\boldsymbol{L}\boldsymbol{L}')}[\boldsymbol{L}(H(\boldsymbol{G}_i); \{Y_{iq}, \boldsymbol{X}_{iq1}, \boldsymbol{X}_{iq2}\}_{q=1}^{10})].$$

The expectation in the definition of the likelihood is intractable, and hence, I approximate it by simulation. Specifically, for $s = 1, \ldots, S$, let $V_s \sim \mathcal{N}(0, \mathbf{I}_6)$ be a draw from a standard 6-variate normal. Define $\theta_s(\boldsymbol{\mu}, \mathbf{L}) \equiv H(\boldsymbol{\mu} + \mathbf{L}\mathbf{V}s)$. Then $\theta_s(\boldsymbol{\mu}, \mathbf{L}) \sim H(\mathcal{N}(\boldsymbol{\mu}, \mathbf{L}\mathbf{L}'))$. The simulated likelihood can then be written

$$\tilde{\mathcal{L}}(\boldsymbol{\mu}, \mathbf{L}) \equiv \prod_{i=1}^{N} \frac{1}{S} \sum_{s=1}^{S} \boldsymbol{L}(\theta_s(\boldsymbol{\mu}, \boldsymbol{L}); \{Y_{iq}, \boldsymbol{X}_{iq1}, \boldsymbol{X}_{iq2}\}_{q=1}^{10}),$$

and the estimates are formed by finding $\boldsymbol{\mu}$ and $\boldsymbol{L}$ values to maximize $\tilde{\mathcal{L}}$. Standard errors are computed using the Huber sandwich formula.

References

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

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