Pre-Analysis Plan for "Taxing Income: Scientific Uncertainty versus Normative Disagreement"

Matthias Rodemeier and Gregory Sun

March 26, 2025

1 Introduction

Since the pioneering work of Mirrlees (1971), optimal income taxation has become a well-developed field of study within economics. Over the years, a large number of theoretical elaborations of the basic framework have contributed to a deep understanding of the salient trade-offs faced by a government in designing their income tax system while empirical studies have provided a wealth of evidence on the likely causal impacts of tax policy reform. In light of the scientific maturity of the field, it is desirable to develop methodology for synthesizing the accumulated knowledge from decades of research on optimal income taxation and to take stock of what this accumulated knowledge can teach us about policy-making. This paper represents a first step towards this end. We focus on the literature which estimates elasticities of taxable income by looking at bunching around "notches" and "kinks" in many real-world tax schedules, which represents arguably the cleanest causal evidence of elasticities of taxable income. We synthesize information in this literature by first developing a *structural*, *Bayesian* meta-analysis framework and using this framework along with a hand-crafted dataset of existing studies using tax kinks/notches to identify structural primitives of an optimal income taxation model in the spirit of Mirrlees (1971).

Our meta-analysis produces a mean estimate of the elasticity of taxable income implied by the literature, and arguably more importantly, also produces uncertainty estimates surrounding this mean estimate. However, to translate this *positive* information into a set of policy implications, these elasticities must be integrated with information about *normative* judgements about which objectives are desirable. We thus next survey the general US public to elicit *tastes* for redistribution. Comparing the variation in redistributive tastes of with the uncertainty estimates implied by our meta-analysis, we ask a central but unexplored question: *To what extent does policy disagreements stem from scientific uncertainty as opposed to from normative disagreement?* We argue that the answer to this question may have important implications for how economists communicate their findings with policymakers and the general public.

1.1 Meta Analysis Details

We begin our study by generalizing the basic kink/notch framework pioneered by Saez (2010) to allow for non-constant elasticities of taxable income (ETI). Within this generalized framework, we formalize what we view as a "folk wisdom" in the literature that even if elasticities are not literally constant across the income distribution, the parametric specifications adopted in the literature approximately estimate the local ETI around the kink/notch under study.

Next, given a collection of pointwise elasticity estimates at various incomes, we provide a framework that integrates information from these various estimates into a single, coherent, structural model of income and how it reacts to tax incentives. We conceptualize the main structural parameter of interest as being a (continuous) *function* mapping location within the income distribution to ETIs around that income level. We can then think about uncertainty over this structural primitive in terms of a *probability distribution over continuous functions* representing the ETI function. Given a "prior" over the ETI function, we can then think of the scientific process as forming "posterior" updates based on the observations of pointwise ETI estimates found in the literature. Within machine-learning, *Gaussian Processes* (GPs), as implemented for instance by Matthews et al. (2017), are a flexible methodology for parameterizing probability distributions over the infinite dimensional space of continuous functions and thus provides precisely the language to formalize this idea. We thus theoretically adapt the GP framework for our purposes before turning to the data and using it to draw from a posterior distribution over the ETI function given our meta-analytical data.

1.2 Survey Experiment

We combine our structural Bayesian meta-analysis framework with survey data to generate novel insights about the scientific and normative drivers of disagreement in what optimal tax policy should look like. In particular, we survey the general public and elicit *i*) preferences over redistribution, *ii*) beliefs about elasticities of taxable income, *iii*) uncertainty about those beliefs about elasticities of taxable income.

After exploring normative disagreements and the data on elasticities of taxable income, we next explore the policy implications of these disagreements. Specifically, for different values of redistributive tastes elicited in the population, and for different draws of elasticity of taxable income from the posterior distribution of our meta-analysis, we compute an optimal income tax schedule. We then ask whether, at various income levels, whether the optimal marginal tax rate varies more fixing redistributive tastes, but sampling from the posterior distribution of elasticites or varying redistributive tastes but fixing a single elasticity estimate.

The rest of this Pre-Analysis Plan elaborates on our empirical strategy and the structural estimation.

2 Descriptive Analysis

The data from our experiment will contain columns depicting:

1. Respondent identifier
2. Respondent demographics
3. Income of recipients 1 and 2
4. Cost of transfers 1 and 2

Our descriptive analysis will seek to answer the following questions: *i*) by how much on average do people trade off transfers to different recipients of different income groups, *ii*) how much heterogeneity in general is there in redistributive preferences, and *iii*) how does this tradeoff vary by observable demographic information.

We can answer question 1 using a simple linear regression model. Specifically, our transfers will vary in terms of the income of recipients as well as in the transfer amount to each of the two hypothetical recipients, $j \in \{1, 2\}$. Denote by I_{iqj} the income of the j^{th} recipient in question q faced by survey-taker i. Similarly, let T_{iqj} denote the transfer amount. Finally, let Y_{iq} be an indicator such that $Y_{iq} = 1$ if survey-taker i chooses to give the transfer to household 1 on question q. Then we will estimate a linear regression model of the form

$$Y_{iq} = \alpha + \beta[\log(T_{iq1}/T_{iq2}) - \gamma[I_{iq1} - I_{iq2}]] + \varepsilon_{iq}.$$

The above model will be estimated by an OLS regression of Y_{iq} on $\log(I_{iq1}/I_{iq2})$ and $[T_{iq1} - T_{iq2}]$ and transforming the resulting coefficients accordingly.

The interpretation of the above coefficients is as follows: α and β are nuisance parameters. The intercept α accounts for any bias towards always picking the first recipient (note that by our randomization, the actual characteristics of households 1 and 2 are drawn from the same distribution, so $\alpha = 0.5$ indicates no bias). The slope β measures how sensitive choices are to recipient characteristics. Finally, γ is the parameter of interest and characterizes the tradeoff between giving a *larger* transfer and giving a transfer to a *preferred recipient*. For example, if the transfer given to household 1 increases by 1% ($\log(T_{iq1}/T_{iq2})$ increases by 0.01, then the income difference between household 1 and 2 must also increase by $0.01/\gamma$ in order to leave survey respondent behavior unchanged. Higher values of γ thus correspond to survey takers who are more responsive to transferring money towards the poor. As will be seen in the next section, it will in particular relate to the CARA parameter describing redistributive preferences.

To answer questions *ii*) and *iii*), we estimate a random coefficients model of the form

$$Y_{iq} = \alpha + \beta[\log(T_{iq1}/T_{iq2}) - \gamma_i \log(I_{iq1}/I_{iq2})] + \varepsilon_{iq}, \quad \gamma_i = \boldsymbol{\theta}'\mathbf{X}_i + \delta_i.$$

Here, $\mathbf{X}_i$ is a vector of demographic characteristics asked about in the survey. To answer question *ii*), we report the overall $\text{Var}(\gamma_i)$ implied by model estimates, as well as the implied "R^2" of the regression of γ_i on $\mathbf{X}_i$. We will also report the $\boldsymbol{\theta}$ coefficients, which tells the slope of the relationship between redistribution preferences and $\mathbf{X}_i$.

In addition to asking about redistribution preferences. We also ask individuals to try to predict the behavior response to changes in a 50% decrease in tax burden. We will report an OLS regression of these predictions on the same demographic observables $\mathbf{X}_i$ and also report mean values of this quantity.

3 Structural Analysis

We model survey-takers as making choices about their preferred redistribution according to marginal social welfare weights parameterized by a CARA function in baseline consumption. More specifically, let i index an individual survey taker. The marginal social welfare weights of individual i are assumed to be of the form

$$g_i(C) = \exp(-\gamma_i C)$$

where C is baseline consumption and γ_i is the absolute risk aversion parameter of individual i. A hypothetical question in our survey asks the following question: consider two households who respectively have an income after taxes and transfers from the government of C_1 and C_2. Would you rather transfer T_1 to household 1 or T_2 to household 2? If the survey taker answers this question according to her re-distributive preferences, she would transfer to household 1 over household 2 if and only if

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

We allow for the possibility that there is some cognitive noise in survey responses so in reality, the survey taker chooses redistribution to household 1 if and only if

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

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

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

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

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

Before describing estimation, we discuss the sampling distribution of the data conditional on individual level parameters γ_i and σ_i. As a reminder, our data comes from a survey experiment where we ask respondents 10 questions where we randomly sample values of (C_1, C_2, T_1, T_2). Let the draws of question parameters be given by

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

The outcomes of the survey can be represented by a vector of dummy variables

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

where Y_{iq} represents whether or not survey respondent i chose to distribute to household 1 on question q. The likelihood function conditional on $\mathbf{X}_i$ and the individual-level parameter values is given by

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

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

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

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

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

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

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

3.1 Choice of Grids

To apply FKRB-style estimators, we must choose the grid of points which (γ_i,σ_i) is supported on. There is little formal guidance in the literature on how to do this, so we pick a datadependent procedure based on what seemed reasonable in a pilot version of our study, which we describe in this subsection.

In the first step, we fit a standard logistic regression without an intercept. The outcome variable is an indicator for whether or not redistribution to family 1 was chosen, and the regressors are $(C_1 - C_2)$ and $[\log(T_1) - \log T_2]$.

Let $\bar{\gamma}$ be the ratio of the coefficients from this regression, and let $\bar{\sigma}$ be the reciprocal of the coefficient in front of $[\log(T_1) - \log T_2]$. Let $\tilde{\boldsymbol{\gamma}}$ be a grid of 40 points that is evenly spaced in log units between $\bar{\gamma}/10$ and $40\bar{\gamma}$. Let $\boldsymbol{\gamma} = (0, \tilde{\boldsymbol{\gamma}}, 100)$.[1] Similarly, let $\boldsymbol{\sigma} = (\bar{\sigma}/3, (\bar{\sigma}/3 + 2\bar{\sigma})/2, 2\bar{\sigma})$ be a grid of error variances. Then we assume that the joint distribution of (γ_i, σ_i) is supported on $\boldsymbol{\gamma} \times \boldsymbol{\sigma}$. This discrete distribution can be represented by a matrix $\boldsymbol{\theta}$ where $\theta_{ab} \equiv \Pr[\gamma_i = a, \sigma_i = b]$, where $\theta_{ab} \geq 0$ for all a, b and $\sum_{a,b} \theta_{ab} = 1$.

4 Sample Size

We initially plan on rolling out the experiment to 2,500 survey participants. This should be sufficient to estimate heterogeneity in the population. In our pilot experiment, we estimate that the 10^{th} percentile of $\log \gamma$ is -10, while the 90^{th} percentile is -8. We aim to collect a large enough sample size to detect this difference with high probability. With our initial sample, we will therefore calculate the standard error on the estimated 90-10 percentile difference of $\log \gamma$, and ensure that this standard error is smaller than $0.5 = (10 - 8)/4$. If the standard error turns out to be above 0.5, we will collect another sample of 2,500 subjects to increase precision.[2]

References

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

[1]The case of $\gamma = 0$ corresponds to "efficiency maximizing" preferences, and someone with $\gamma = 0$ will always choose the largest transfer. The case of $\gamma = 100$ corresponds to Rawlsian redistributive preferences in that the respondent will always transfer to the poorer household. In a pilot experiment, we find that a sizable fraction of the population behaves according to one of the above two "extreme" preferences. An issue with interpreting these as true preferences is that these two extreme cases are also observationally equivalent to survey responses corresponding not to true preferences, but instead, corresponding to simple "rules of thumb". The $\gamma = 0$ case corresponds to only looking at transfer amount, while the $\gamma = 100$ case corresponds to only looking at incomes. Our downstream analysis will take the distributional of γ_i conditional on $\gamma_i \notin \{0, 100\}$ as our measure of true preferences. This will likely understate the true dispersion in redistributive preferences.

[2]One practical reason why we would collect another 2,500 subjects is that this is the maximal sample size with a *representative* sample that our survey platform Prolific offers per trial.

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

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

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

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