Designing, not Checking, for Policy Robustness: An Example with Optimal Taxation

Benjamin B. Lockwood, Afras Sial, and Matthew Weinzierl*

July 13, 2020

Abstract

Economists typically check the robustness of their results by comparing them across plausible ranges of parameter values and model structures. A preferable approach to robustness—for the purposes of policymaking and evaluation—is to design policy that takes these ranges into account. We modify the standard optimal income tax model to include uncertainty over parameter values and characterize robust optimal policy as that which maximizes expected social welfare. After calibrating uncertainty over the elasticity of taxable income from past empirical work and novel survey data on economists' beliefs, we compare the implied robust optimal marginal tax rates to two benchmarks: a probability-weighted average of optimal policies computed for each possible set of parameter values, and the optimal policy based on the best estimates of them. Preliminary results indicate that robust optimal marginal tax rates are smoother and more progressive than in benchmark analyses.

Policymaking must proceed in the face of widespread uncertainty.[1] Despite our best efforts, economists cannot provide policymakers with definitive estimates of most, if any, of the inputs to optimal policy models (much less guarantee that we are using the right models). Therefore,

*University of Pennsylvania Wharton School and NBER; University of Pennsylvania Wharton School; Harvard Business School and NBER.

[1]We will use the term uncertainty as a shorthand for risk, uncertainty, and ignorance, three classes of limits to understanding (see Zeckhauser 2014). Our analysis has the policy designer face risk (known probabilities of known states) and uncertainty (unknown probabilities of known states), and we use the latter to encompass both. We would be eager to incorporate ignorance (unknown probabilities of unknown states), but doing so would take us too far from standard policy evaluation methods. See Weinzierl (2017) for a possible approach to managing ignorance in policy design.

economists studying optimal policy design have a responsibility to ensure that their results are robust to plausible variation in the model inputs. In this paper, we argue that the current approach to robustness used by most economists is flawed as a guide to policymaking and evaluation, however natural it may be for academic research, and we propose an alternative.

A typical robustness analysis is designed for the use of journal editors and referees, not policymakers. In particular, the results obtained in a baseline case are typically labeled "robust" when they are close—in some unspecified but generally understood sense—to the results obtained at several points within a plausible space of parameter values and range of model specifications. This approach provides reassurance that the results are not special or fragile; that is, that they do not rely on a particular calibration or modeling choice. But this approach is at odds with how economists themselves advise policymakers to confront limits to their understanding, and it is therefore a disservice to policymakers (and their economic advisors). Instead, most economists believe rational policymaking in the face of uncertainty is nothing other than expected social welfare maximation,[2] in which policymakers ought to take into consideration the probabilities and implications of the full range of plausible parameter values when choosing policy. In other words, policies ought to be robust to parameter and model uncertainty *by design*.

The expected social welfare maximization approach has at least two advantages as a guide to robust policy. First, it makes it more difficult to ignore nonlinearities in "outlier" results from far-off parts of the parameter space. Under the current approach, if some combination of parameter values is unlikely but not impossible, and yet has substantial implications for the results, it may easily escape the analyst's attention and encourage a false sense of security. This risk is especially great in complex situations with a number of uncertain parameters, when relying on the researcher's best judgment may be insufficient. Expected welfare maximization avoids this problem, as it is designed to be sensitive to any extreme welfare implications. Second, expected welfare maximation delivers a rigorously determined optimal policy compromise in the face of uncertainty, whereas the current method provides no guidance on how the results under baseline and "outlier" cases are to be combined. It is striking, upon reflection, how seldom the results of robustness checks under the current approach lead to recommendations for *any* adjustment to the baseline optimal policy.

[2]We address primarily economists working in the welfarist tradition widely assumed in modern economic analysis. A concern for robustness may lie behind support for nonwelfarist principles of policy design, though we do not explore that possibility here. See Weinzierl (2018, 2019) for an elaboration of the case for nonwelfarist principles.

To demonstrate how using expected social welfare maximation modifies standard analytical and quantitative optimal policy results, we focus on the much-studied optimal income tax problem. Working within the standard Mirrlees (1971) model, we assume that its key parameters may be state-dependent and that the tax designer, who maximizes expected social welfare, is uncertain over the true state of the world. The resulting necessary conditions for a robust optimal tax policy, and the policy implied by them under a specific calibration, can be compared to two benchmarks: an "average optimal policy" which equals the probability-weighted average of policies computed for each possible value of these parameters, and the "best estimate optimal policy" based on the expected values of these parameters.

Our new analytical results clarify how the robust optimal tax policy more effectively responds to relationships among uncertain parameters than do the two benchmark policies. For example, in one scenario considered below, we find that the necessary conditions for robust optimal marginal income taxes depend on the inverse of the expected value (across states) of the product of the elasticity of taxable income and the density of the income distribution $\left(\text{using the notation below, } \frac{1}{E_s\left[\zeta_s h_{s,T}(y)\right]}\right)$, not the expected value of the inverse of the product of those parameters ($E_s\left[\frac{1}{\zeta_s h_{s,T}(y)}\right]$, upon which depends the average optimal policy) or the inverse of the product of their expected values ($\frac{1}{E_s[\zeta_s]E_s\left[h_{s,T}(y)\right]}$, upon which depends the best estimate optimal policy).

Our quantitative results indicate that, for plausible levels of uncertainty, the robust optimal tax policy remains broadly similar to the policy chosen under certainty (equivalently, to the best estimate optimal policy), but instructive differences exist. To calibrate the model, we must construct probability distributions over key parameter values. We start with the long-term elasticity of taxable income, for which we draw on two data sources: the vast existing empirical literature estimating it, and a novel survey of academic economists.[3] Then, for each value of that elasticity, we infer an income-earning ability distribution from current data on income and the tax system. We are thus able to construct state-dependent joint probability distributions for two of the dimensions of uncertainty facing tax policy designers and use them to compare the robust optimal policy to the two benchmark policies. We find that the robust optimal tax schedule retains the now-standard U-shape discussed in Diamond (1998) and Saez (2001) and sets rates at roughly

[3]We therefore combine many individuals' beliefs, having the policy designer treat uncertainty in the same formal manner as risk. If some sources of beliefs are more credible than others, their influence would be greater on the policy designer's probability distribution over states.

similar levels with or without uncertainty. At the same, time, robust optimal rates are smoother, more progressive, and welfare-improving. Our calibration of uncertainty, based on results from our pilot survey of public finance economists, indicate that the robust optimal policy raises marginal tax rates across the upper middle class by about 5 percentage points relative to the best estimate policy, and generates annual welfare gains of $114.16 billion in the United States.

Finally, we note that our alternative approach in no way undermines, and may improve, the most important benefit of conventional robustness checks. Both scholars and makers of policy want to know whether, in what ways, and by how much a policy based on the best of their knowledge will be substantially in error if they turn out to have been wrong. This desire makes it natural to compare (in a variety of ways) two policies: the optimal policy given baseline parameter values (what we call the best estimate optimal policy), and the optimal policy given substantially different values. Under our approach, the first of these two policies is replaced by the optimal *robust* policy, which maximizes expected social welfare over the possible range of parameter values. Therefore, the results of a traditional robustness analysis can still teach us how costly are deviations of the actual parameter values from our beliefs, but we now measure these costs relative to a rationally constructed robust policy.

1 Analytical results without and with uncertainty

To make our approach tangible and demonstrate its value, we work within the currently dominant optimal tax framework based on Mirrlees (1971) and developed by Diamond (1998) and Saez (2001), among many others. The Mirrleesian literature has expanded along a number of dimensions, with each expansion introducing greater sophistication and complexity to the (already demanding) model. But the basic model's core result on marginal tax rates (that is, distortions to individuals' consumption-leisure tradeoffs) relies on a remarkably short list of parameters: the long-term[4] elasticity of taxable (labor) income, the shape of the income (or income-earning ability) distribution, and marginal welfare weights along the income (or ability) distribution.[5]

[4]As Saez, Slemrod, and Giertz (2012) write in their literature review on this elasticity, "The long-term response is of most interest for policy making."

[5]Note that we designate the elasticity of taxable income as containing within it all responses to tax policy that affect income, including (for example) human capital accumulation decisions. Changes to the income distribution not captured by this elasticity may be due to non-tax policy changes (such as education) and changes to the productive environment (such as technological). As there is no other income in the model, and the model is static (it can be

Uncertainty is rampant, however, over even the short list of parameters on which Mirrleesian optimal tax results hinge. Taxable income elasticities have been exhaustively studied for decades, but their long-term values in response to substantial tax changes remain elusive given the limits of econometrics in a complex and dynamic world.[6] The income distribution is relatively well understood at a point in time, but its long-term evolution and (non-tax-policy) determinants remain only vaguely understood.[7] Even marginal social welfare weights, which for decades were viewed as arising in a straightforward way from a utilitarian social welfare function, have recently come up for debate, with some authors (including this paper's) finding support for quite different values.[8]

Our model policymaker takes this uncertainty into account by imagining each possible collection of parameter values as a state of the world, specifying a probability distribution over states of the world, and choosing tax policy to maximize expected social welfare.

1.1 Linear taxation

To build intuition, we begin with an optimal linear tax model. The tax system is defined by a lump-sum grant b and constant marginal rate t.

1.1.1 Certainty benchmark

Individuals differ in their unobservable income-earning ability type, indexed by $i = \{1, 2, ..., I\}$, and choose labor effort to maximize utility $U\left(c^i, y^i\right)$,where c^i is individual i's after-tax income (i.e., consumption) and y^i is pre-tax income (i.e., the product of ability and effort). We assume utility is quasilinear (that is, linear in c but not in y), so there are no income effects of taxation

thought of as a lifetime model), there is no income shifting in response to a tax change.

[6] As Saez, Slemrod, and Giertz (2012) point out, "the long-term response is more difficult to identify empirically. The empirical literature has primarily focused on short-term (one year) and medium-term (up to five year) responses, and is not able to convincingly identify very long-term responses." In part, they argue, this difficulty is because the elasticity of taxable income is not a "structural parameter" with a value independent of context. Instead, it depends on the existing tax and broader economic system, such that "an elasticity estimated around the current tax system may not apply to a hypothetical large tax change." Factors of potential importance on which empirical evidence is quite limited, and which thereby complicate the econometrics, include the long-term responsiveness of investment in human capital and interdependent social norms.

[7] This uncertainty is closely related to the assumed unobservability of income-earning ability at the heart of the Mirrlees model. Both the income and ability distributions will depend, in the long term, on the evolution of a range of unforeseeable factors such as production technology, the educational system, fertility choices, and various non-tax policies.

[8] The inverse-optimum literature (see Bourguignon and Spadaro 2012, Bargain et al. 2014, and Lockwood and Weinzierl 2016) seeks to infer welfare weights from existing policy choices, but its results are subject to substantial uncertainty. See also Weinzierl (2014) and references therein.

on an individual's optimization. We use σ^i to denote the (compensated and uncompensated, as income effects are absent) elasticity of the individual's taxable income with respect to $1-t$. An individual earning income y^i pays tax $T\left(y^i\right) = -b + ty^i$, so the individual's budget constraint is $c^i = (1-t)\, y^i + b$, and individual i's indirect utility can be denoted $V^i\left(1-t, b\right)$. We denote average income with $\bar{y} = \sum_i p^i y^i$, where p^i is the fixed population proportion of type i in the economy.

The tax authority's objective is to maximize a simple-sum utilitarian measure of social welfare: $\sum_i p^i V^i\left(1-t, b\right)$. This form for the objective is standard—but not uncontroversial—in the literature. We will use the term $g^i = \frac{\partial V^i}{\partial b}$ to denote the marginal social welfare of income for individual i, and we normalize such that $\sum_i p^i g^i = 1$ at the optimum policy. The tax authority's budget constraint is $t\sum_i p^i y^i = b + R$, where R is exogenous required public goods spending.

We adopt a perturbation method to derive the necessary conditions for optimal policy made familiar (in the nonlinear tax context) by Saez (2001). The idea is to consider a small increase dt in the tax rate t and add up its effects. Only if these effects sum to zero can the tax policy (off of which the perturbation was considered) be optimal.[9]

In our setting, the effects of the perturbation dt can be collected into two classes: mechanical and behavioral. The mechanical effect includes a revenue increase of $y^i dt$ from each individual i and a consequent welfare decrease of $g^i y^i dt$ for each individual i. The total mechanical effect, dM, is therefore

$$dM = \sum_i p^i (1 - g^i) y^i dt. \tag{1}$$

The behavioral effect arises from each individual i adjusting earnings by $-\frac{dy^i}{d(1-t)}dt$, which thereby changes tax revenue. The total behavioral effect, dB, is

$$dB = -t \sum_i p^i \frac{dy^i}{d\left(1-t\right)} dt. \tag{2}$$

If the initial policy is optimal, then $dM + dB = 0$. With some simplification (and recalling that $\frac{dy^i}{d(1-t)}\frac{1-t}{y^i} = \sigma^i$), the optimal linear tax rate t satisfies

$$\frac{t}{1-t} = \frac{-Cov[g^i, y^i]}{E\left[\sigma^i y^i\right]}. \tag{3}$$

[9] In Appendix A, we show that this perturbation method yields the same results for optimal linear taxation as does the direct optimization of the tax authority's objective.

We additionally assume $\sigma^i \equiv \sigma$ is constant across individuals, and we retain this assumption throughout the rest of the paper. Then, result (3) can be expressed more simply as:

$$\frac{t}{1-t} = \frac{-Cov\left[g^i, \frac{y^i}{\bar{y}}\right]}{\sigma}. \tag{4}$$

In words, expression (4) says that the optimal linear tax rate t is positive if the value society places on an additional unit of income is lower for higher-income individuals ($Cov\left[g^i, \frac{y^i}{\bar{y}}\right] < 0$), and that tax rate is larger the more that value declines with income. On the other hand, the optimal linear tax rate is smaller the greater the elasticity (and thus efficiency costs) of taxation (i.e., the larger is σ).

This paper is motivated by the observation that result (4) is obtained under the assumption that its terms are known with certainty by the tax authority. In reality, all aspects of the right-hand side of this expression are estimable only with great uncertainty, if at all. We turn now to understanding how robust optimal linear taxation manages that uncertainty.

1.1.2 Setting with uncertainty

Now, we consider a setting in which multiple possible states of the world, indexed by $s = \{1, 2, ..., S\}$, arise with known probabilities π_s, where $\Sigma_s \pi_s = 1$.[10] Each state corresponds to different realized values of the components of expression (4) over which the tax authority is uncertain when designing policy. In particular, we now assume state-dependence of social welfare weights g_s^i, relative and average incomes $\frac{y_s^i}{\bar{y}_s}$, and the (uniform) elasticity of taxable income σ_s.

The tax authority must specify one tax system $\{b^*, t^*\}$ that will apply to whichever state is eventually realized (the asterisks are used to highlight the robust optimal tax policy). Conceptually, we have in mind choosing a tax policy that is expected to apply, unchanged, over a long time horizon. The tax authority's objective is to maximize expected social welfare as before, but now the expectation is also taken over states: $\sum_s \pi_s \sum_i p_s^i W\left(V_s^i\left(a, b^*\right)\right)$.[11] We constrain the tax authority to a pre-specified, state-invariant policy; that is, to a fixed linear rate t and lump-sum grant $b(t) = t\sum_s \pi_s \sum_i p_s^i y_s^i - R$, so that the budget constraint is met in expectation. Note that this

[10] Uncertainty over π is, of course, another potential dimension the tax authority may need to manage. For simplicity, we abstract from it.

[11] We assume the tax authority is risk neutral in social welfare levels.

setting allows, in effect, the tax authority to transfer resources across states, in that its budget constraint may not hold in a given state. But, it does not allow the tax authority to choose state-specific lump-sum grants optimally, for example by increasing the grant in states where labor disutility is high on average.[12]

Before characterizing robust optimal policy, we note that our approach to robust policy design contrasts with the prominent approach of Hansen and Sargent (2001) in macroeconomics, who introduce techniques of robust optimal control in which the primary concern is to avoid worst-case scenarios (see Williams 2008). Formally, Hansen and Sargent's policymakers proceed as if they are playing a game against a malevolent Nature that chooses the worst realization of parameter or model uncertainty given the chosen policy. As a result, the robust policy literature within macroeconomics typically emphasizes maximin policy objectives. Our expected social welfare approach reacts less strongly to worst-case scenarios. For example, while policy under our approach is chosen with the risks of worst-case scenarios included in the calculations, it also takes into account the likelihood of best-case scenarios (and all scenarios in between).[13]

Applying the same perturbation method as above, we obtain parallel *expected* mechanical and behavioral effects, which we denote dM^e and dB^e. The expected mechanical effect, analogous to equation (1) above, is

$$dM^e = \sum_s \pi_s \left[\sum_i p_s^i (1 - g_s^i) y_s^i dt \right]. \tag{5}$$

The expected behavioral effect, analogous to equation (2) above, is

$$dB^e = -t \sum_s \pi_s \left[\sum_i p_s^i \frac{dy_s^i}{d(1-t)} dt \right]. \tag{6}$$

[12]One technical implication of this assumption is that the normalization of the welfare weights g_s^i sets their expectation equal to unity across states, so that the expectation of these weights within any given state may differ from unity. Two additional technical subtleties raised by the presence of multiple states—namely the choice of whether to tie welfare weights to utility levels or types across states and the choice of how to normalize cardinal utility levels across states—will be addressed in the quantitative section below, where their resolution has implications for the results.

[13]The contrast between our approach and that of Hansen and Sargent is analogous to the contrast, pointed out by Arrow (1973), between economists' conventional utilitarianism and the influential Difference Principle of Rawls (1971). Rawls reasoned that individuals designing society's institutions from behind a veil of ignorance, and thereby ignorant of their own positions within that society, would maximize the well-being of the worst-off individual in that society. Arrow argued against this inference as reflecting too extreme a level of risk aversion, in particular relative to Harsanyi's (1953; 1955) inference (using a similar thought experiment designed to emphasize impartiality) that social welfare ought to be the probability-weighted sum of individual utilities within society. We take much the same view of how policymaking should respond to parameter uncertainty.

In this scenario, with uncertainty, a necessary condition for optimality is that the tax authority does not view such a perturbation as yielding a net gain *in expectation*. Setting $dM^e + dB^e = 0$, the robust optimal linear tax rate t^* satisfies:

$$\frac{t^*}{1-t^*} = \frac{1}{\sum_s \pi_s \left[\sigma_s \bar{y}_s\right]} \left(\sum_s \pi_s \bar{y}_s - \sum_s \pi_s \left[\sum_i p_s^i g_s^i y_s^i \right] \right). \tag{7}$$

Again, we can express this result more simply:

$$\frac{t^*}{1-t^*} = \frac{E_s \left[-Cov \left[g_s^i, \frac{y_s^i}{\bar{y}_s} \right] \bar{y}_s \right]}{E_s \left[\sigma_s \bar{y}_s\right]} \tag{8}$$

To understand the adjustments the tax authority makes to achieve robustness in the face of uncertainty, we compare result (8) to two alternatives. The tax authority might solve for the optimal linear tax rate given each possible state of the world (using a state-specific version of 4) and then take the expectation of those values across states. We denote the resulting policy with $\left\{\bar{b}, \bar{t}\right\}$ and refer to it as the "average optimal linear tax rate," which satisfies:

$$\frac{\bar{t}}{1-\bar{t}} = E_s \left[\frac{-Cov \left[g_s^i \frac{y_s^i}{\bar{y}_s} \right] E_s \left[\bar{y}_s\right]}{\sigma_s E_s \left[\bar{y}_s\right]} \right]. \tag{9}$$

Alternatively, the tax authority might set policy based on the best estimates of the uncertain parameters, which would yield the same result as the benchmark analysis with certainty from above, result (8). This "best estimate optimal linear tax rate" satisfies:

$$\frac{t}{1-t} = \frac{-Cov \left[E_s \left[g_s^i \right], E_s \left[\frac{y_s^i}{\bar{y}_s} \right] \right] E_s \left[\bar{y}_s\right]}{E_s \left[\sigma_s\right] E_s \left[\bar{y}_s\right]}. \tag{10}$$

Comparing results (8), (9), and (10) reveals the main difference between robust optimal linear taxes and the two alternatives.

Mathematically, understanding this difference starts with the observation that expressions (8), (9), and (10) are of nearly the same structure. That is, all three rates (t^*, $\bar{t}$, and t) depend in the same directions on the same multiplicative terms: the product of the covariance (between welfare weights and relative income) and mean income in the numerator and the product of the elasticity

of taxable income and mean income in the denominator. How they differ is in the level at which the expectation over states is taken: at the level of these multiplicative terms for the robust optimal rate t^*, at the level of the ratio of these terms for the average optimal rate $\bar{t}$, and at the level of each factor within these terms for the best estimate optimal rate t. As a result, neither the average optimal linear tax rate in (9) nor the best estimate linear tax rate in (10) utilizes information most effectively to maximize expected welfare.

Economically, the intuition behind the difference is that the robust tax policy is sensitive—to the optimal degree—to complementarities in the effects of uncertainty's resolution across states. For example, suppose there are three states $s = \{1, 2, 3\}$, and their elasticities are symmetrically distributed, such as $\sigma_1 = 0.25$, $\sigma_2 = 0.50$, and $\sigma_3 = 0.75$. Suppose further that the high-elasticity state also has a substantially greater average income than in the other states, such that $E_s\left[\sigma_s \bar{y}_s\right] > E\left[\sigma_s\right] E\left[\bar{y}_s\right]$. That is, the robust optimal rate reflects that the distortionary costs of taxation are greater in expectation than what the best estimate optimal rate would suggest because the state in which the elasticity of taxable income is high is also that in which incomes are greater.[14] The conceptual contrast between the robust optimal rate and the average optimal rate is more subtle: while the former takes these complementarities into account separately in the numerator and denominator, the latter takes them into account through (i.e., by taking the expectation over states of) their ratio. If the values of the multiplicative terms in the numerator and denominator are correlated across states, either positively or negatively, the average optimal linear tax rate will respond not to the ratio of their expected values but to the expected value of their ratio, missing opportunities for more sophisticated robustness to uncertainty.

More colloquially, the robust optimal tax policy tries to protect against worst-case scenarios, take advantage of best-case scenarios, and manage the cases in between, rather than just do best in the expected case or choose a simple average of its best state-specific values.

Alternative assumption on the tax authority's budget constraint We also consider a scenario in which the tax authority must satisfy its budget constraint state-by-state, such that the

[14] If uncertainty is asymmetric in other ways—for example, if the distribution of elasticities has a mean greater than its median—this difference between the robust optimal policy and the best estimate policy is potentially even greater. And a similar logic applies to the terms in the numerator: to the extent that variation across states in parameters generates asymmetric uncertainty in the welfare effects of taxation, the robust optimal policy will differ from the best estimate policy.

expression for the state-specific lump-sum grant is $b_s(t) = t\sum_i p_s^i y_s^i - R$. This setting constrains the tax authority more severely than our previous setting, as it disallows transfers of resources across states. Which setting is preferable depends on, among other factors, how great a share of the tax authority's overall fiscal constraint we believe this model captures. If it captures a small share, the first setting may be preferable, as the tax authority may then choose a tax policy that is both self-funding in expectation and not self-funding in the true state but can be offset by opposing imbalances elsewhere. If it captures a large share, the second setting may be preferable because it more accurately reflects the necessary adjustment to the lump-sum transfer that the tax authority will have to make once the true state obtains. Note that both settings reflect our assumption that the tax authority cannot learn about the true state, for example by monitoring the extent to which the policy is self-funding, and thus reduce its uncertainty.

In this setting, the tax authority must take into account that the chosen tax rate will have implications for the lump-sum grant that vary—perhaps dramatically—by state. The same process as above yields the following result for robust optimal policy:

$$\frac{t^*}{1-t^*} = \frac{E_s\left[-Cov\left[g_s^i, \frac{y_s^i}{\bar{y}_s}\right]\bar{y}_s\right] + Cov[\bar{g}_s, \bar{y}_s]}{E_s\left[\sigma_s \bar{y}_s\right] + Cov[\bar{g}_s, \sigma_s \bar{y}_s]} \tag{11}$$

Compared to result (8), result (11) incorporates an additional covariance term into both the numerator and the denominator. These additions capture two forces: on the one hand, if increasing the tax rate generates extra revenue in those states where marginal utilities are high on average (large $\bar{g}_s$) then the tax should be higher; on the other hand, if increasing the tax rate depresses income in those states especially strongly, it should be lower. (Future versions of this paper may explore this case in greater depth.)

1.2 Nonlinear optimal taxation

We now remove the linearity restriction on the tax system to characterize robust optimal nonlinear income taxes, thus connecting more closely to the main Mirrleesian literature.

1.2.1 Certainty benchmark

Individuals are indexed by their (possibly multi-dimensional) type $\theta \in \Theta$, which is distributed according to $F(\theta)$ with associated density $f(\theta)$. Individuals choose, taking the tax system as given, a level of labor effort to maximize individual utility. The product of labor effort and type is income y, which is distributed (conditional on the tax function T) according to $H_T(y)$ with associated density $h_T(y)$. Let $\zeta(y)$ denote the compensated labor supply elasticity and $\eta(y)$ the income effect. The nonlinear income tax $T(y)$ is chosen by the tax authority to maximize social welfare subject to the feasibility constraint $\int T(y)h_T(y)dy \geq R$, where R is exogenous public goods expenditure.

As with the linear tax case, we use the perturbation method to derive our results. Consider a small increase in the marginal tax rate $d\tau$ in an interval of size ε around some earnings level y^*. The mechanical effect of this increase has three components. First, it raises revenues equal to $\varepsilon d\tau \int_{y^*}^{\infty} h_T(z)dz$ from those earning y^* and more. Second, it directly reduces after-tax income among these individuals, with an effect on social welfare of $\varepsilon d\tau \int_{y^*}^{\infty} (-g(z))\, h_T(z)dz$, where $g(y)$ is (as before) the marginal social welfare weight on y-earners in terms of public funds. Third, this after-tax income reduction for individuals with earnings above y^* may cause them to raise their earnings through an income effect, generating an increase in revenue for the tax authority. Denote this fiscal externality as $\varepsilon d\tau \int_{y^*}^{\infty} \frac{-T'(z)}{1-T'(z)}\eta(z)h_T(z)dz$. For notational simplicity, let $\hat{g}(y) = g(y) - \frac{T'(y)}{1-T'(y)}\eta(y)$ denote the combination of the latter two effects.[15] We can write the overall mechanical effect, then, as:

$$dM = \varepsilon d\tau \int_{y^*}^{\infty} (1 - \hat{g}(z))\, h_T(z)dz. \tag{12}$$

This tax increase also generates a behavioral effect, as y^*-earners reduce their effort and, thereby, generate a negative fiscal externality equal to

$$dB = \varepsilon d\tau \frac{-T'}{1-T'} \zeta(y^*)y^* h_T(y^*). \tag{13}$$

At the optimum, the tax system is optimal only if the sum of these effects is zero. That is, if $T(y)$

[15] In the case of quasilinear utility, which we assumed in the linear tax case above but relax here, income effects are absent and $\hat{g}(y)$ reduces to $g(y)$.

is optimal, then the following condition on the marginal tax rate $T'(y)$ holds:

$$\frac{T'(y)}{1-T'(y)} = \frac{1}{y\zeta(y)h_T(y)} \int_y^\infty (1-\hat{g}(z))\, h_T(z)dz. \tag{14}$$

1.2.2 Setting with uncertainty

As in the linear setting, we now allow for uncertainty over the parameters that matter for optimal taxation. In particular, we imagine that there are multiple possible states of the world indexed by s, each arising with probability π_s such that $\sum_s \pi_s = 1$. We denote uncertainty through state-dependence of the individual type distributions $F_s(\theta)$ and densities $f_s(\theta)$, endogenous income distributions $H_{s,T}(y)$ and densities $h_{s,T}(y)$, compensated elasticities of taxable income and income effects $\zeta_s(y)$ and $\eta_s(y)$, and marginal welfare weights $g_s(y)$. We assume that the tax authority must pre-specify the entire tax function—that is, both the nonlinear income tax schedule $T(y)$ and the lump-sum grant $T(0)$—to apply in all states so that the tax authority's budget constraint is satisfied in expectation: $\sum_s \pi_s \left[\int_y T(z)h_{s,T}(z)dz\right] \geq R$.

The perturbation method proceeds as before. Consider a small increase in the marginal tax rate $d\tau$ in an interval of size ε around some earnings level y^*. The *expected* mechanical effect is expression (12) modified for multiple states:

$$dM^e = \varepsilon d\tau \sum_s \pi_s \left[\int_{y^*}^\infty (1-\hat{g}_s(z))\, h_{s,T}(z)dz\right]. \tag{15}$$

The expected behavioral effect is a similar modification of expression (13):

$$dB^e = \varepsilon d\tau \sum_s \pi_s \left[\frac{-T'(y^*)}{1-T'(y^*)}\zeta_s(y^*)y^*h_{s,T}(y^*)\right] \tag{16}$$

The sum of these expected effects must be zero under the robust optimal tax, so:

$$\frac{T^{*\prime}(y)}{1-T^{*\prime}(y)} = \frac{E_s\left[\int_y^\infty (1-\hat{g}_s(z))\, h_{s,T}(z)dz\right]}{yE_s\left[\zeta_s(y)h_{s,T}(y)\right]}. \tag{17}$$

Again, it is useful to compare this result to the average optimal tax policy

$$\frac{\bar{T}'(y)}{1-\bar{T}'(y)} = E_s\left[\frac{\int_y^\infty (1-\hat{g}_s(z))\, h_{s,T}(z)dz}{y\zeta_s(y)h_{s,T}(y)}\right], \tag{18}$$

and the best estimate optimal tax policy:

$$\frac{T'(y)}{1-T'(y)} = \frac{\int_y^\infty E_s\left[(1-\hat{g}_s(z))\right] E_s\left[h_{s,T}(z)\right] dz}{yE_s\left[\zeta_s(y)\right] E_s\left[h_{s,T}(y)\right]}. \tag{19}$$

As with the results in the linear case, the most noticeable difference among these results is the level over which the expectation across states is taken. The robust optimal policy at income y takes into account the multiplicative interactions within states between the compensated elasticity of taxable income and the density of the income distribution at y as well as those between the marginal welfare weight at y and that density. In contrast, the average optimal nonlinear tax policy turns on the ratio of those multiplicative interactions within states, while the best estimate nonlinear income tax takes expectations of the factors within those multiplicative terms. In words, the robust optimal policy appropriately reacts to how inputs to the optimal policy function may build upon or offset each other in the various states of the world; that is, how uncertainty across inputs to the model may exacerbate or mitigate the optimal policy design problem.

Alternative assumption on the tax authority's budget constraint Again, we also consider a scenario in which the tax authority must satisfy its budget constraint state-by-state. Specifically, we require the tax authority to pre-specify the nonlinear income tax schedule $T(y)$ which will apply across states, but we adjust the state-specific lump-sum grant to ensure state-by-state budget balance. Formally, the expression for the state-specific lump-sum grant is $b_s = \int_{y>0} T(z)h_{s,T}(z)dz - R$. This setting constrains the tax authority more severely than our previous setting, as it disallows transfers of resources across states.

Using same perturbation approach, we obtain expressions for the mechanical effect:

$$dM^e = \varepsilon d\tau \sum_s \pi_s \lambda_s \left[\int_{y^*}^\infty (1-\hat{g}_s(z))\, h_{s,T}(z)dz\right], \tag{20}$$

and the behavioral effect

$$dB^e = \varepsilon d\tau \sum_s \pi_s \lambda_s \left[\frac{-T'(y^*)}{1-T'(y^*)} \zeta_s(y^*) y^* h_{s,T}(y^*) \right]. \tag{21}$$

In both expressions, the term λ_s denotes the state-specific marginal values of public funds, which measures the welfare impact of a small increase in the state-specific lump-sum grant.

At the optimum, the following condition must hold: :

$$\frac{T^{*\prime}(y)}{1-T^{*\prime}(y)} = \frac{E_s\left[\lambda_s \int_y^\infty (1-\hat{g}_s(z))\, h_{s,T}(z) dz\right]}{y E_s\left[\lambda_s \zeta_s(y) h_{s,T}(y)\right]}. \tag{22}$$

Expression (22) captures factors quite similar to those in expression (11): tax rates should be higher if they generate more revenue when the state-specific marginal value of public funds λ_s is high (explaining its presence in the numerator); but tax rates should be lower if that marginal value of public funds is high precisely when the elasticity of taxable is high (explaining the presence of λ_s in the denominator).

2 Evidence on the extent of parameter uncertainty

In this section, we attempt to quantify some of the uncertainty to which the robust optimal tax conditions of the previous section respond. We restrict our attention to the dimensions of uncertainty over which we have relatively good information: that is, over the parameters for which uncertainty is not as great. In particular, we start by focusing on what economists know—and don't know—about the policy-relevant value(s) of the elasticity of taxable income. Then, we use recent data on the tax system and income distribution to infer, for each possible value of that elasticity, an underlying income-earning ability distribution. These exercises produce a joint probability distribution for the elasticity of taxable income and the income distribution, allowing us to simulate and compare robust and benchmark optimal policies in the next section.

We note, however, that we are including only a small portion of the risk, uncertainty, and ignorance[16] that faces a tax designer using the optimal policy model above. In particular, we abstract from the many non-tax factors that will affect the evolution of the income distribution, and we

[16] See the first footnote for a discussion of these three terms.

assume certainty in normative judgments of optimal policy (i.e., in the pattern of welfare weights). In reality, these other dimensions of limits to our understanding—and how they interact with each other and the dimensions we do include—are likely to have additional substantial quantitative implications for robust policy.

2.1 The elasticity of taxable income

To characterize beliefs on the long-term elasticity of taxable income, we turn to two sources of evidence: the existing empirical literature, and a new survey of academic economists.

2.1.1 Existing literature on the elasticity of taxable income

A leading survey of the literature on the elasticity of taxable income—Saez, Slemrod, and Giertz (2012)—summarizes the state of knowledge about this parameter's value over the horizon of relevance to optimal tax theory: "Estimates of the elasticity of taxable income in the long run (i.e., exceeding a few years) are plagued by extremely difficult issues of identification, so difficult that we believe that there are no convincing estimates of the long-run elasticity of reported taxable income to changes in the marginal tax rate."[17] Giertz (2010) reaches a similarly discouraging conclusion: "In most ETI studies, it is more accurate to refer to estimated elasticities as either 'short term' or 'longer term' (as opposed to 'long term'). Long-term responses may be the most important but could take many years before responses are fully observed. These types of changes (which may include some human capital and occupation decisions as well as more traditional investment) are currently beyond the scope of the ETI literature. For that reason, this article eschews the label 'long term' in favor of 'longer term' in order to distinguish these estimates from short-term elasticities while also recognizing that they are not truly long term."

Lacking definitive evidence, what do economists believe is a reasonable value for this elasticity? Three examples will serve to illustrate the range of beliefs. Diamond and Saez (2011) are skeptical of *any* substantial real responses to marginal income tax rates,[18] concluding that "the elasticity $e = 0.57$ is a conservative upper bound estimate" and that a value of 0.9 would be "extremely

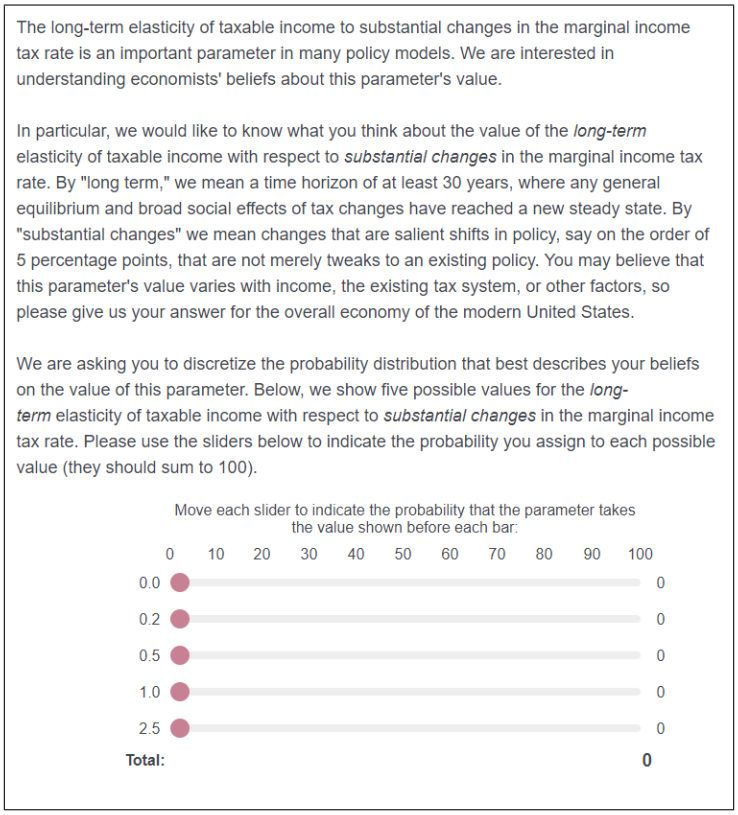
[17] As do we, these authors note that the uncertainty surrounding the parameters on which their work focuses extends to many areas of economics, writings: "Many of these problems are not unique to identifying the long-run ETI, but apply to the estimation of all behavioral responses."

[18] Real responses are those that change income earned, not reported, and are the relevant ones for optimal tax policy (because enforcement is a concern outside the theory).

high." Giertz (2010) estimates values "from 0.78 to 1.46 over the longer term" of three to six years, and thus Giertz (2009) considers values ranging from 0.2 to 1.0 for his calculations of the efficiency consequences of tax reforms. Jones (2019) incorporates idea generation into the model and—though he acknowledges substantial uncertainty—prefers parameterizations that imply a responsiveness of aggregate income corresponding to elasticities as high as 2.36.[19] More generally, the difficulty—even impossibility—of taking into account truly long-term factors such as changes to social norms around work in response to taxation—strongly suggests to us that economists have very limited knowledge of the policy-relevant value for this key parameter.

2.1.2 Novel survey evidence on the elasticity of taxable income

We also take a direct route to gauging economists' beliefs about the value for σ: we ask them. We design a novel survey to elicit beliefs about this parameter, in which we ask following question:



We note two features of this question. First, we describe in some detail the parameter about

[19] These values are imputed by this paper's authors, using Jones's optimal top marginal tax rates from his Table 1 and the standard (see Diamond and Saez 2011) formula for the top marginal tax rate with a Pareto parameter $\alpha = 1.5$.

which we are asking. Though perhaps tedious to the respondent, this detail is essential for eliciting the right beliefs for our purposes. Second, we centered the discretized distribution on 0.5 to align with what we believe most economists are likely to have in mind as the "baseline" value from recent prominent surveys (including Diamond and Saez 2011, Chetty et al. 2011, Chetty 2012, and others). We include 0.0 and 2.5 as "extreme" values in an attempt to span the full range of beliefs, though of course negative and even larger positive values for this parameter are possible. An alternative setup would have asked respondents to specify their own discrete points as well as probabilities; we worried that doing so would tempt respondents to give over-confident answers that require less effort and thought.

This paper reports results based on a small pilot survey with responses from 9 leading public finance economists. The results of this pilot survey are shown in Figure 3, where we plot the mean answer for each point in the discretized distribution.

2.1.3 Combining evidence on the elasticity of taxable income

The evidence on σ shows that robust optimal tax policy design confronts both risk and uncertainty. In particular, a robust tax policy designer faces not only the *risk* posed by the distribution over the possible values for σ that any individual researcher will suggest, but also the *uncertainty* posed by disagreement over that distribution across researchers. In principle, robust optimal policy might respond differently to these two types of limits to understanding.

We will simplify matters and treat risk and uncertainty in the same formal manner: that is, both are managed through expected social welfare maximization. Different possible probability distributions are simply weighted (perhaps uniformly, though differential credence could generate a motive for variation) and summed, yielding a single probability distribution used as the basis for robust policy.

2.2 The shape of the income distribution

A second source of uncertainty in the optimal tax model is the long-term shape of the income distribution and how it varies with the elasticity of taxable income. Forecasting the evolution of aggregate income, much less its distribution, is notoriously difficult, as it requires anticipating major forces—such as technological and cultural changes—whose underlying drivers are poorly

understood.

We will focus on a narrow slice of the broad uncertainty over the income distribution's shape: the range of possible *current* underlying income-earning ability distributions that is implied by different plausible values for the elasticity of taxable income. Given an observed income distribution, an existing tax system, and a model of the individual labor supply decision, we can infer a current underlying distribution of income-earning abilities for each value of the elasticity of taxable income.[20]

These current underlying ability distributions are all consistent with one existing income distribution, given the existing tax system, but if the tax system is changed they will translate into different income distributions. That is, when we reverse the inference process that yielded these ability distributions, instead plugging them into the model of individual labor supply with a new tax system, we will obtain a new income distribution for each value of the elasticity of taxable income. We can then feed the joint probability distribution of these pairs of elasticities and income distributions into the robust optimal tax model from the previous section.

Formally, consider the individual utility maximization problem: $\max_{y^i} U\left(c^i, y^i\right)$, where $c^i = y^i - T\left(y^i\right)$, using the same notation as in the nonlinear case above. To relate our results to the large existing literature on optimal taxation without income effects, we assume that the individual utility function takes the Type 1 form from Saez (2001):

$$U\left(c^i, y^i\right) = G\left(c^i - \frac{1}{1+1/\sigma}\left(\frac{y^i}{w^i}\right)^{1+1/\sigma}\right), \tag{23}$$

where G is a concave transformation. The individual's first-order condition is therefore

$$w^i = \left(y^i\left(1 - T'\left(y^i\right)\right)^{-\sigma}\right)^{\frac{1}{1+\sigma}}. \tag{24}$$

To obtain $T(y)$, we calibrate the current U.S. tax system using the empirical mapping between a broad measure of market income and consumption, as reported in Piketty, Saez, and Zucman (2018). This measure thus incorporates broad features of the tax and transfer including not only income taxes at all levels, but also other tax credits and phaseouts, in-kind transfers and other

[20] We are assuming that income is the product of these abilities and labor supply as determined by the individual's optimization.

social support programs, and an accounting of distributed public goods.

We again use the income distribution from Piketty, Saez, and Zucman (2018) to calibrate the status quo income distribution. Using the calibrated status quo tax function T_{US} and the status quo income distribution $H_{T_{US}}(y)$, we can compute an underlying implied ability distribution using the first-order condition in (24) for any assumed value of the elasticity of taxable income σ. More generally, if the policymaker faces an uncertain *distribution* of possible elasticity values, this procedure permits us to compute a separate underlying ability distribution in each case, which can then be used to compute an optimal expected tax function in the face of such uncertainty. We next compute the results of such an exercise, for a range of assumptions about the elasticity of income with respect to tax rates.

3 Quantitative results

In this section, we use the optimality conditions and calibration procedures described above in order to demonstrate the effect of parameter uncertainty on optimal tax results. In other words, for a given assumption about the *average* value of a model parameter—in this case, the elasticity of taxable income—what is the effect of uncertainty on the optimal tax schedule?

Here we continue to employ the individual utility function in equation (23). To account for the agent's diminishing marginal utility of consumption (or, equivalently in this model, the policymaker's degree of inequality aversion) an assumption must be made about the shape of the concave transformation G. As Saez (2001) notes, this concavity can be described in a transparent way using "social marginal welfare weights" g^i, proportional to $G'\left(c^i - \frac{1}{1+1/\sigma}\left(\frac{y^i}{w^i}\right)^{1+1/\sigma}\right)$ at the optimum. In our baseline simulations, we use a logarithmic functional form $G(x) = \ln(x)$.

We first consider two illustrative examples in which the optimal tax schedule is computed for different calibrations with common average values of the elasticity of taxable income but very different levels of uncertainty. In the first example, we compare the case in which $\sigma = 0.4$ with certainty, relative to the case in which the planner believes that the elasticity of taxable income is likely to be low—i.e., $\sigma = 0.1$ with probability 0.7—but there is a smaller chance that it is much higher—i.e., $\sigma = 1.1$ with probability 0.3. These values result in the same average value of $\sigma = 0.4$. The optimal marginal tax rate schedules for these specifications are plotted in Figure 1. In Figure

2, we present results for even more skewness, in which σ takes the values of 0.1, 0.2, and 2.6 with probabilities 0.4, 0.5, and 0.1, respectively.

These results immediately highlight a few implications of parameter uncertainty.

First, the shape of the optimal tax schedule remains broadly similar with and without uncertainty, and retains the trademark U-shape discussed in Diamond (1998) and Saez (2001), and the overall magnitude of the marginal tax rates is only modestly altered.

Second, the shape of the tax schedule is generally "smoother" under uncertainty. This has an intuitive interpretation. If there is a concentration of the population density at a specific point in the income distribution, then that implies a high density point in the ability distribution, the location of which can be predicted with certainty in the face of a hypothetical tax reform. As a result, the optimal tax schedule will feature a dip right at the point of that higher density, in order to reduce distortionary marginal tax rates for that high density of workers. On the other hand, if the elasticity is uncertain, the location of that point of population density under a reformed tax schedule is less predictable, and so the depressing effects on optimal tax rates are more diffuse.

A third feature of uncertainty is that the tax schedule generally becomes more progressive. This is particularly apparent from the higher skew distribution of uncertainty in Figure 2, which shows that uncertainty tends to substantially reduce marginal tax rates at low incomes, while raising them at higher incomes. The rationale for this "progressivity effect" has a simple intuition. Under both specifications (with and without uncertainty) optimal tax rates are generally higher than in the U.S. status quo. That means that the overall income distribution will be shifted downward, in the optimum, relative to the status quo. The *amount* by which incomes shift down, though, is related to the elasticity. If σ is high, then incomes will shift downward more substantially—and, in that event, taxes are particularly distortionary, so the optimal tax rates should be low. On the other hand, when the income elasticity is low, incomes will not shift down by much, and at the same time, the low elasticity implies that income taxes are not very distortionary. As a result, optimal marginal tax rates are higher. Combining these forces results in an optimal tax rate schedule that is more progressive than in the baseline case with certainty.

While the examples in Figures 1 and 2 are illustrative, in Figure 4 we present results calibrated using the data from our survey of expert economists. Specifically, we assume the elasticity of taxable income may take each of six different values corresponding to intervals we asked about in

the survey (0.10, 0.20, 0.40, 0.75, 1.50, and 2.50), with probabilities corresponding to the average across economists elicited in our survey (normalized to ensure a sum of 100%). The results of this calibration suggest that the robustly optimal income tax is modestly more progressive than the best estimate optimum under the degree of uncertainty prevailing in this group of surveyed economists.

4 Conclusion

In this paper, we propose a method by which to design economic policy analyses for robustness to uncertainty over the models and parameters at their core. We apply this method to the topic of optimal income taxation in the Mirrleesian tradition and use both existing and novel empirical evidence on the range of professional economists' beliefs over the value of the elasticity of taxable income to introduce uncertainty into the analysis. We characterize—both analytically and quantitatively—robust optimal income taxes and compare them to benchmark results, suggesting that this method might be fruitfully applied to a wide range of economic policies.

References

Arrow, Kenneth J. 1973. "Rawls's Principle of Just Saving." *The Swedish Journal of Economics* 75 (4):323–335.

Bargain, Olivier, Mathias Dolls, Dirk Neumann, Andreas Peichl, and Sebastian Siegloch. 2014. "Tax-Benefit Revealed Social Preferences in Europe and the US." *Annals of Economics and Statistics* (113/114):257–289.

Bourguignon, François and Amedeo Spadaro. 2012. "Tax–Benefit Revealed Social Preferences." *The Journal of Economic Inequality* 10 (1):75–108.

Chetty, Raj. 2012. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." *Econometrica* 80 (3):969–1018.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *American Economic Review* 101 (3):471–75.

Diamond, Peter and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives* 25 (4):165–90.

Diamond, Peter A. 1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review* 88 (1):83–95.

Giertz, Seth. 2009. "The Elasticity of Taxable Income: Influences on Economic Efficiency and Tax Revenues, and Implications for Tax Policy." In *Tax Policy Lessons from the 2000s*, edited by Alan Viard. Washington, D.C.: AEI Press, 101–136.

———. 2010. "The Elasticity of Taxable Income during the 1990s: New Estimates and Sensitivity Analyses." *Southern Economic Journal* 77 (2):406–433.

Hansen, Lars Peter and Thomas J Sargent. 2001. "Acknowledging Misspecification in Macroeconomic Theory." *Review of Economic Dynamics* 4 (3):519–535.

Harsanyi, John C. 1953. "Cardinal Utility in Welfare Economics and in the Theory of Risk-Taking." *Journal of Political Economy* 61 (5):434–435.

———. 1955. "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *Journal of Political Economy* 63 (4):309–321.

Jones, Charles I. 2019. "Taxing Top Incomes in a World of Ideas." NBER Working Paper No. 25725.

Lockwood, Benjamin B. and Matthew Weinzierl. 2016. "Positive and Normative Judgments Implicit in US Tax Policy, and the Costs of Unequal Growth and Recessions." *Journal of Monetary Economics* 77:30–47.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *The Review of Economic Studies* 38 (2):175–208.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *Quarterly Journal of Economics* 133 (2):553–609.

Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Belknap Press of Harvard University Press.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *The Review of Economic Studies* 68 (1):205–229.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2012. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." *Journal of Economic Literature* 50 (1):3–50.

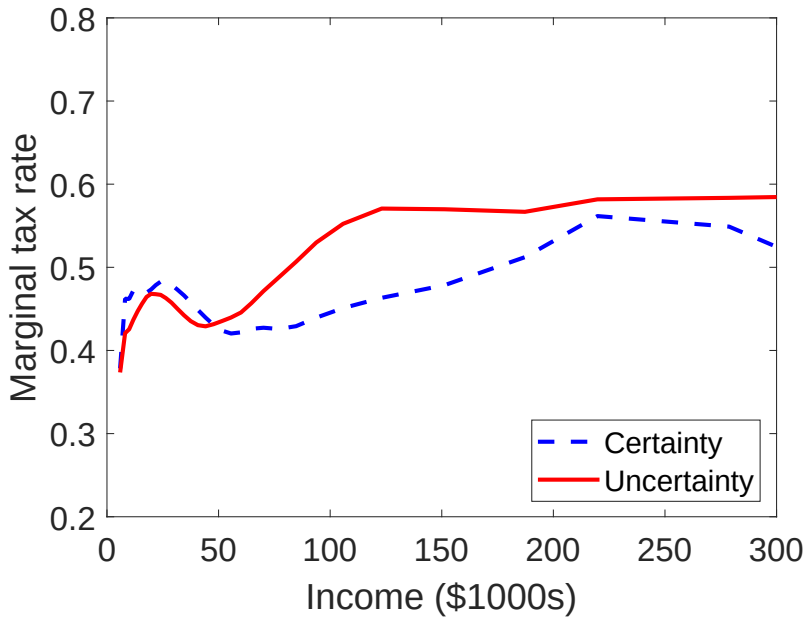
Weinzierl, Matthew. 2014. "The Promise of Positive Optimal Taxation: Normative Diversity and a Role for Equal Sacrifice." *Journal of Public Economics* 118:128–142.

Williams, Noah. 2008. "Robust Control." *New Palgrave Dictionary of Economics* 2 (6).

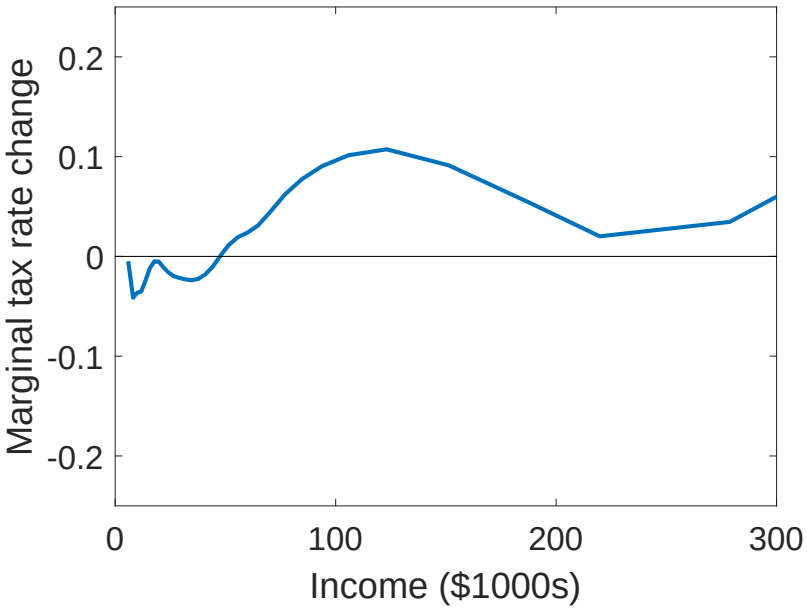
Zeckhauser, Richard. 2014. "New Frontiers Beyond Risk and Uncertainty: Ignorance, Group Decision, and Unanticipated Themes." In *Handbook of the Economics of Risk and Uncertainty*. Oxford, UK: Elsevier.

Figure 1: Optimal Tax Schedules Under Parameter Uncertainty

(a) **Comparison of Optimal Tax Rate Schedules**



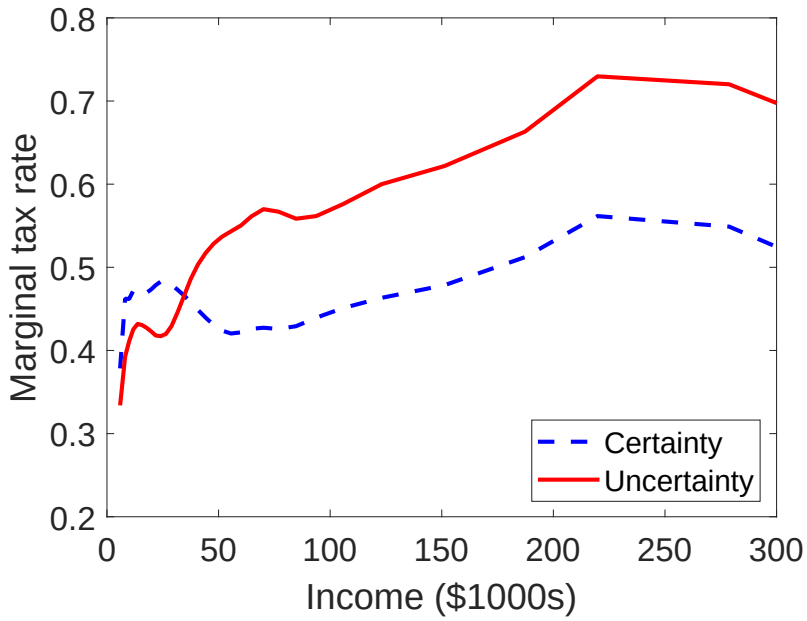
(b) **Change in Optimal Marginal Tax Rates due to Uncertainty**



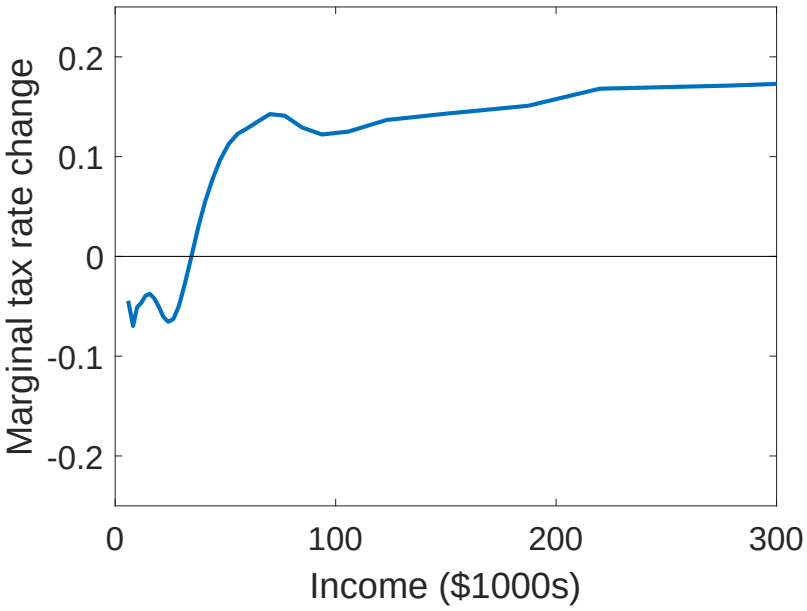
Notes: This figure plots optimal marginal tax rates under a given average value of the elasticity of taxable income, but for different assumptions about uncertainty. The line labeled "Certainty" corresponds to a specification in which the elasticity is 0.4. The lump-sum grant received under this tax schedule is $\$26,160$. The line labeled "Uncertainty" plots the optimal schedule of tax rates when the policymaker is uncertain about the elasticity of taxable income, believing there is a 70% chance that it is 0.1, and a 30% chance that it is 1.1. The lump-sum grant received under this tax schedule is $\$27,666$.

Figure 2: Optimal Tax Schedules Under Parameter Uncertainty (Skewed)

(a) **Comparison of Optimal Tax Rate Schedules**

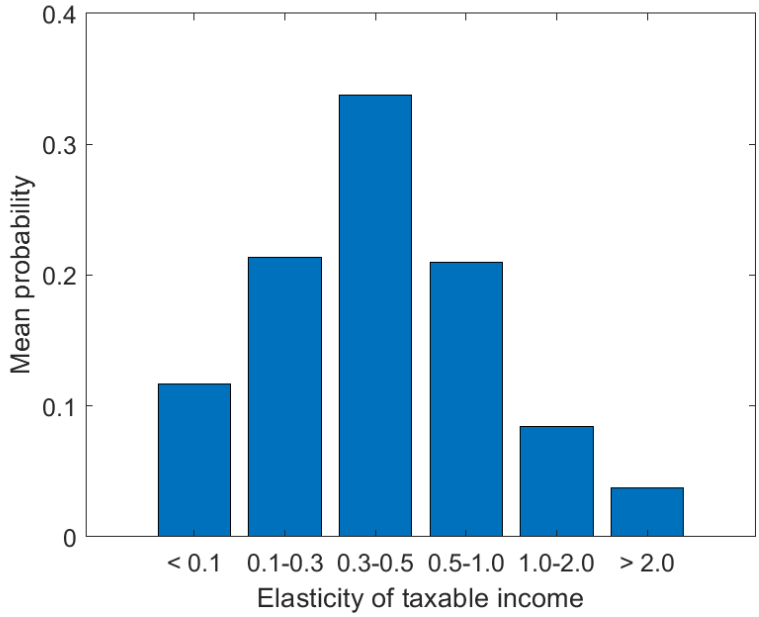


(b) **Change in Optimal Marginal Tax Rates due to Uncertainty**



Notes: This figure plots optimal marginal tax rates under a given average value of the elasticity of taxable income, but for different assumptions about uncertainty. The line labeled "Certainty" corresponds to a specification in which the elasticity is 0.4. The lump-sum grant received under this tax schedule is $\$26,160$. The line labeled "Uncertainty" plots the optimal schedule of tax rates when the policymaker is uncertain about the elasticity of taxable income, believing there is a 40% chance that it is 0.1, a 50% chance that it is 0.2, and a 10% chance that it is 2.6. The lump-sum grant received under this tax schedule is $\$29,453$.

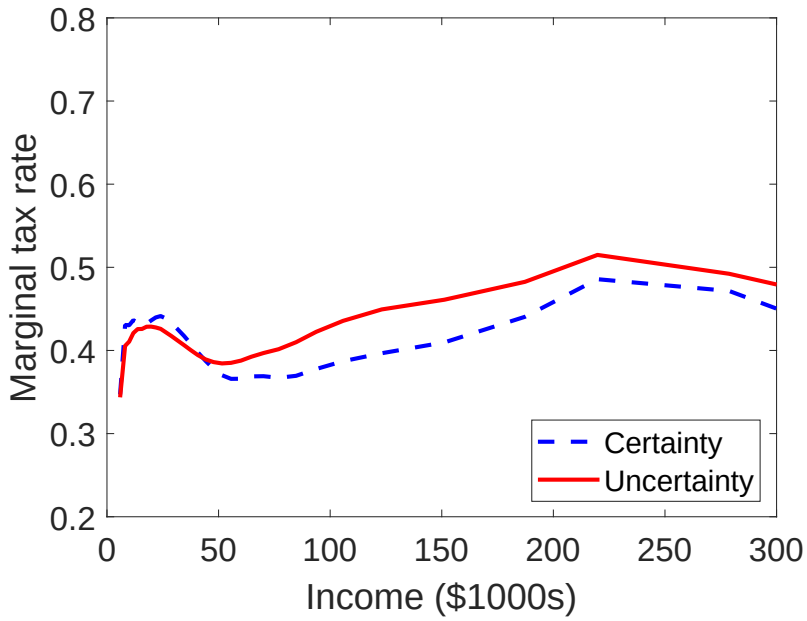
Figure 3: Probability Distribution of the Elasticity of Taxable Income (Pilot Survey Data)



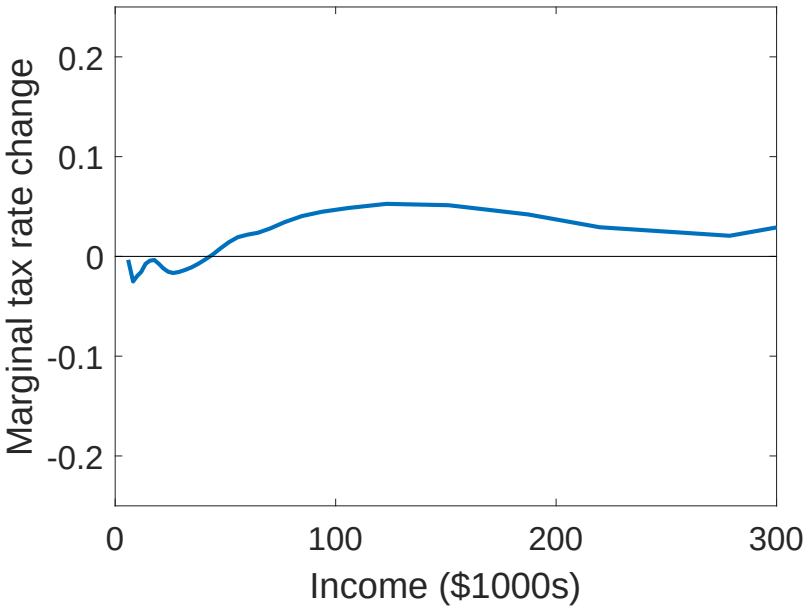
Notes: This figure plots the mean probability assigned to each range of values of the long-term uncompensated elasticity of taxable income for which we elicited beliefs in our survey of academic economists. Data are based on a pilot survey of 9 public finance economists.

Figure 4: Optimal Tax Schedules Under Parameter Uncertainty (Pilot Survey Data)

(a) **Comparison of Optimal Tax Rate Schedules**



(b) **Change in Optimal Marginal Tax Rates due to Uncertainty**



Notes: This figure plots optimal marginal tax rates under a given average value of the elasticity of taxable income, but for different assumptions about uncertainty. The line labeled "Certainty" corresponds to a specification in which the elasticity is 0.57, the expected value of the elasticity from our survey of academic economists. The lump-sum grant received under this tax schedule is $\$22,893$. The line labeled "Uncertainty" plots the optimal schedule of tax rates when the policymaker is uncertain about the elasticity of taxable income, with the same beliefs as the average academic economist from our survey. The policymaker believes there is a 11.7% chance that the elasticity is 0.10, a 21.3% chance that it is 0.20, a 33.8% chance that it is 0.40, a 21.0% chance that it is 0.75, a 8.4% chance that it is 1.50, and a 3.8% chance that it is 2.50. The lump-sum grant received under this tax schedule is $\$23,611$.

Appendix

A Derivations with Direct Maximization

In this appendix, we show how the results derived for the optimal linear tax through the perturbation method can also be obtained through a direct maximization of the tax authority's objective.

A.1 Certainty benchmark

Individual utility is $U^i(c,y) = c - v^i(y/w^i)$, with v^i increasing and convex. Types are indexed by i with population fractions p^i. The social welfare function is $\sum_i p^i \Phi\left(U^i(c,y)\right)$, with Φ increasing and concave. Tax policy is defined by a linear rate t, which determines a lump-sum grant $b(t) = t\sum_i p^i y^i$. Individuals choose $y^i(t) = \arg\max_y \left\{y(1-t) + b - v^i(y/w^i)\right\}$, and we denote $\sigma^i = \frac{dy^i(t)}{d(1-t)} \cdot \frac{1-t}{y^i}$.

The tax authority's policy problem is $\max_t \left\{\sum_i p^i \Phi\left(y^i(t)(1-t) + b(t) - v^i(y^i(t)/w^i)\right)\right\}$. The first-order condition for the optimum is:

$$\sum_i p^i \Phi'(U^i) \left(-y^i(t) + \sum_j p^j y^j(t) - t \sum_j p^j \frac{dy^j(t)}{d(1-t)} \right) = 0, \tag{25}$$

or, rearranging,

$$\frac{t}{1-t} \sum_i p^i \Phi'(U^i) \cdot \sum_j p^j \sigma^j y^j = \sum_i p^i \Phi'(U^i) \cdot \sum_j p^j y^j - \sum_i p^i \Phi'(U^i) y^i. \tag{26}$$

Now, we add notation for simplicity: $\lambda = \sum_i p^i \Phi'(U^i)$ is the social value of marginally augmenting the lump-sum grant (also known as the MVPF); $\bar{x} = \sum_i p^i(x^i)$ is the expected value of an individual-specific variable x^i; and $g^i = \frac{\Phi'(U^i)}{\lambda}$ is the social marginal welfare weight on type i. Note that $\bar{g} = \sum_i p^i \frac{\Phi'(U^i)}{\lambda} = \frac{\sum_i p^i \Phi'(U^i)}{\sum_j p^j \Phi'(U^j)} = 1$. With these terms, we can rewrite the first-order condition as

$$\frac{t}{1-t} = \frac{E\left[y^i\right] - E\left[g^i y^i\right]}{E\left[\sigma^i y^i\right]} \tag{27}$$

$$= \frac{-Cov[g^i, y^i]}{E\left[\sigma^i y^i\right]}. \tag{28}$$

If we additionally assume $\sigma^i \equiv \bar{\sigma}$ is constant, then this FOC simplifies further to

$$\frac{t}{1-t} = \frac{-Cov\left[g^i, \frac{y^i}{\bar{y}}\right]}{\bar{\sigma}}. \tag{29}$$

A.2 Uncertainty, state-invariant lump-sum grant

In this scenario, states are indexed by s, with associated probabilities π_s. Individual utility is $U_s^i(c,y) = c - v_s^i(y/w^i)$; population fractions are p_s^i. The social welfare function is $\sum_s \pi_s \sum_i p_s^i \Phi\left(U_s^i(c,y)\right)$, with Φ increasing and concave. Tax policy is made up of a state-invariant linear rate t, which determines the state-invariant lump-sum grant: $b(t) = t\sum_s \pi_s \sum_i p_s^i y_s^i$. Note that while this setup implies resources can be transferred across states, it is *not* the same as imposing the optimal state-specific lump-sum grants, as that would generally entail making the grant higher in states where labor disutility $v_s^i(y/w^i)$ is high on average. In other words, this is the optimal tax and transfer when both must be pre-specified.

The individual optimization problem is $y_s^i(t) = \arg\max_y \left\{y(1-t) + b_s - v_s^i(y/w_s^i)\right\}$, where we denote $\sigma_s^i = \frac{dy_s^i(t)}{d(1-t)} \cdot \frac{1-t}{y_s^i}$. The tax authority's policy problem is

$$\max_t \left\{ \sum_s \pi_s \sum_i p_s^i \Phi\left(y_s^i(t)(1-t) + b(t) - v_s^i(y_s^i(t)/w_s^i)\right) \right\}. \tag{30}$$

The optimum is characterized by the following first order condition:

$$\sum_s \pi_s \sum_i p_s^i \Phi'(U_s^i) \left(-y_s^i(t) + \sum_{s'} \pi_{s'} \sum_j p_{s'}^j y_{s'}^j - t \sum_{s'} \pi_{s'} \sum_j p_{s'}^j \frac{dy_{s'}^j(t)}{d(1-t)} \right) = 0. \tag{31}$$

Again, we introduce notation, now adapted for uncertainty: $\lambda = \sum_s \pi_s \sum_i p_s^i \Phi'(U_s^i)$ is the MVPF; $\bar{x} = \sum_s \pi_s \sum_i p_s^i x_s^i$ is the expectation over types and states; $\bar{x}_s = \sum_i p_s^i x_s^i$ is the average across people within a state s; and $g_s^i = \frac{\Phi'(U_s^i)}{\lambda}$, still implying $\bar{g} = 1$. With these in hand, we rewrite the optimality condition as:

$$\frac{t}{1-t} = \frac{E[y_s^i] - E[g_s^i y_s^i]}{E[\sigma_s^i y_s^i]} \tag{32}$$

$$= \frac{-Cov[g_s^i, y_s^i]}{E[\sigma_s^i y_s^i]}. \tag{33}$$

And if we further assume that $\sigma_s^i \equiv \bar{\sigma}_s$, i.e., there is uncertainty about the (homogeneous) value of σ, then the FOC simplifies to:

$$\frac{t}{1-t} = \frac{-Cov[g_s^i, y_s^i]}{\bar{\sigma}\bar{y} + Cov[\bar{\sigma}_s, \bar{y}_s]} \tag{34}$$

$$= \frac{-Cov\left[g_s^i, \frac{y_s^i}{\bar{y}}\right]}{\bar{\sigma}\left(1 + Cov\left[\frac{\bar{\sigma}_s}{\bar{\sigma}}, \frac{y_s^i}{\bar{y}}\right]\right)}. \tag{35}$$

A.3 Uncertainty, budget constraint binds in each state

Finally, we assume that the budget constraint binds within each state, i.e., $b_s(t) = t\sum_i p_s^i y_s^i$. The policy problem is $\max\left\{\sum_s \pi_s \sum_i p_s^i \Phi\left(y_s^i(t)(1-t) + b_s(t) - v_s^i(y_s^i(t)/w_s^i)\right)\right\}$. The optimality condition is:

$$\sum_s \pi_s \sum_i p_s^i \Phi'(U_s^i)\left(-y_s^i(t) + \sum_j p_s^j y_s^j - t\sum_j p_s^j \frac{dy_s^j(t)}{d(1-t)}\right) = 0. \tag{36}$$

We can rewrite this result as

$$\frac{t}{1-t}\sum_s \pi_s \sum_i p_s^i \bar{g}_s \sigma_s^i y_s^i = \sum_s \pi_s \sum_i p_s^i \bar{g}_s y_s^i - \sum_s \pi_s \sum_i p_s^i g_s^i y_s^i, \tag{37}$$

implying

$$\frac{t}{1-t} = \frac{E[\bar{g}_s y_s^i] - E[g_s^i y_s^i]}{E[\bar{g}_s \sigma_s^i y_s^i]} \tag{38}$$

$$= \frac{-Cov[g_s^i, y_s^i] + Cov[\bar{g}_s, \bar{y}_s]}{E[\sigma_s^i y_s^i] + Cov[\bar{g}_s, \sigma_s^i y_s^i]} \tag{39}$$

And if we further assume that $\sigma_s^i \equiv \bar{\sigma}_s$, i.e., there is uncertainty about the (homogeneous) value of σ, then the FOC simplifies to

$$\frac{t}{1-t} = \frac{-Cov[g_s^i, y_s^i] + Cov[\bar{g}_s, \bar{y}_s]}{E[\bar{g}_s \bar{\sigma}_s \bar{y}_s]} \tag{40}$$

$$= \frac{-Cov\left[g_s^i, \frac{y_s^i}{\bar{y}}\right] + Cov\left[\bar{g}_s, \frac{\bar{y}_s}{\bar{y}}\right]}{\bar{\sigma}\left(1 + Cov\left[\frac{\bar{\sigma}_s}{\bar{\sigma}}, \frac{\bar{y}_s}{\bar{y}}\right] + Cov\left[\bar{g}_s, \frac{\bar{\sigma}_s}{\bar{\sigma}}\frac{\bar{y}_s}{\bar{y}}\right]\right)}. \tag{41}$$