# Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions

Benjamin B. Lockwood and Matthew Weinzierl*

**Abstract**

We use official data and standard optimal tax conditions to infer the positive and normative judgments implicit in U.S. tax policy since 1979. We find that explanations within this framework for the time path of U.S. policy require central parameters of the model, namely the elasticity of taxable income or the marginal social welfare weights on top earners, to take unconventional values. We use inferred social preferences to provide novel estimates of the welfare costs of unequal growth and recessions and find that they are sensitive to the assumed distortionary costs of taxation and the year from which preferences are derived. We explore several potential explanations for our results with available data. Among other possibilities, those data are consistent with the hypothesis that unconventionally high marginal social welfare weights on top earners prevail upon policy. To examine whether those inferred weights reflect true public opinion or a biased political outcome, we present new survey evidence on the public's desire for increasing redistribution above current levels. We find that the results of such a survey are highly sensitive to what information is provided to survey respondents.

## Introduction

In this paper, we use official data on U.S. income distributions and standard theoretical conditions on optimal marginal tax rates to infer the combinations of positive and normative judgments implicit in U.S. tax policy over the last several decades.

Our main finding is that explanations within this standard framework for the time path of U.S. policy from 1979 to 2010 require central parameters of the model to take unconventional values. In particular, we identify three possible explanations: first, the perceived costs of redistribution may have risen well above standard estimates; second, the perceived benefits of redistribution from high earners may have fallen below levels widely assumed to hold in modern analyses of optimal taxation, perhaps due to the political influence of the rich; or third, those perceived benefits may have been persistently well below commonly assumed levels–again perhaps due to the political influence of the rich–while the perceived costs increased toward the upper end of the conventional range.

We attempt to use the results of this inference exercise to provide novel quantitative estimates of the welfare costs of two prominent empirical phenomena: rising inequality and recessions. We find, however, that these estimates are highly sensitive to the explanation one adopts for the evolution of U.S. policy (e.g., to the assumed distortionary costs of taxation) and the year from which one infers social preferences.

*Thanks to Nathan Hendren, Louis Kaplow, Greg Mankiw, Andreas Peichl, and Glen Weyl for insightful discussions. Lockwood: Harvard University, lockwood@fas.harvard.edu; Weinzierl: Harvard University and NBER, mweinzierl@hbs.edu.

We then present empirical evidence pertaining to the three possible explanations for our findings, though the limitations of that evidence are admittedly substantial. With regard to the first explanation, we find no evidence for a large increase in the perceived distortionary costs of taxation among either the public or experts, including in official government policy analyses. We show that academic estimates of the overall elasticity of taxable income and labor supply are consistently small, with perhaps a slight increase over this period. We also find that concerns about taxation having negative trickle-down effects do not appear to have increased over this period. With regard to the second explanation, we show evidence that the level of support for government actions to reduce inequality has been stable throughout this period. Though the proper interpretation of that stability is unclear, it may be consistent with a decrease in the perceived benefits of redistribution. Beliefs about how well the political process represents average citizen's interests are stable over this period, providing some very limited evidence that such an inferred fall in the perceived benefits of redistribution is not due to increased divergence between policy outcomes and voter preferences. With regard to the third explanation, we describe existing survey evidence that shows strong popular support for increasing the extent of redistribution above prevailing levels, suggesting a large but stable divergence between policy outcomes and voter preferences. We also provide new survey evidence that shows the sensitivity of this finding to the information provided to survey respondents. In particular, we find that support for increasing redistribution above prevailing levels falls substantially when respondents are presented with information on the distribution of pre-tax income, the distribution of total federal tax payments, or both distributions. The evidence on this third explanation thus remains especially unclear.

Methodologically, our approach is as follows. For each year from 1979 through 2010 we use market (pre-tax-and-transfer) income quantiles reported by the Congressional Budget Office (CBO) to calibrate a Pareto-lognormal income distribution and the NBER's TAXSIM tool to construct the prevailing marginal tax rate schedule. We then invert the Diamond (1998) and Saez (2001) optimal marginal tax rate formula to compute implicit marginal social welfare weights (MSWWs) for a range of elasticities of taxable income in each year.[1] To calculate the welfare costs of rising inequality and recessions, we apply these weights to changes over time in the levels of disposable income experienced at different points in the income distribution.

Our analysis is part of a recent explosion of similar research. A number of papers have assumed an elasticity of taxable income, inverted optimal tax formulas, and estimated the implied MSWWs at points in time for a range of countries. Francois Bourguinon and Amedeo Spadaro (2012) estimate the MSWWs for France in 1995 and find that they are negative on high earners unless the labor supply elasticity is small, similar to our findings in U.S. for the early 1980s. Amedeo Spadaro, Lucia Mangiavacchi, and Luca Piccoli (2012) estimate and compare average MSWWs on five income-earning groups for 26 European countries, finding substantial variation across them. Olivier Bargain, Mathias Dolls, Dirk Neumann, Andreas Peichl, and Sebastian Siegloch (2011; 2013) include the United States (as well as European countries) in their analysis. In Bargain et al. (2011) they examine 2005 U.S. policy and data, estimate the relevant labor supply elasticities, and calculate implicit MSWWs that are quite flat across the income distribution relative to many European countries, consistent with our findings for the same time period if elasticities of taxable income are perceived to be small. In Bargain et al. (2013), they study how tax policy has affected inequality in the United States from 1979-2007 and find partisan effects on policy consistent with the trends we show below. Floris Zoutman, Bas Jacobs, and Egbert L.W. Jongen (2013*a*; 2013*b*) analyze in detail the Dutch tax system and proposals for it by Dutch political parties to infer prevailing and preferred MSWWs in the Netherlands. In Zoutman et al. (2013*a*) they find that top MSWWs are negative unless the elasticity of

[1]The work of Saez and Stantcheva (2013) develops and applies MSWWs as a key tool for policy evaluation.

labor supply is small, consistent with the Bourguinon and Spadaro results for France and ours for the United States. In Zoutman et al. (2013*b*) they show (among other things) that the tax proposals by Dutch political parties have the same feature. Finally, Nathaniel Hendren (2014) formalizes an isomorphic alternative to MSWWs called the inequality deflator that does not rely on the existence of a social welfare function, and he uses it compare alternative economic environments in the spirit of the Kaldor-Hicks compensation principle. He calculates the implicit weights put on marginal income across the U.S. income distribution in 2012, finding a pattern similar to what we infer for 2010. He then estimates the costs of rising income inequality since 1980 at 15–20 percent of total economic growth in his baseline calibration. When we replicate his approach using our data, we find a very similar result (16 percent).

Our paper contributes to this fast-growing literature in three ways. First, our focus on intertemporal analysis allows us to see the puzzling features of positive and normative judgments that are required to explain policy over time, features that can be missed when looking at one point in time. Second, as part of trying to explain these puzzling features, our paper also emphasizes the importance for policy inversion exercises of this sort of the *perceived* distortionary costs of taxation (as well as perceptions of other issues such as trickle-down effects). Finally, as far as we are aware, the application of this approach to the costs of recessions is new to this paper.

Our paper's novel survey evidence also contributes to two active literatures: one in economics on stated preferences toward redistribution (as in Ilyana Kuziemko et al. (2013)) and one in political science on the potential link between economic and political inequality in the modern United States (as in Larry Bartels (2008) and Martin Gilens (2012)). The interpretation of our findings depends in part on how well the evolution of policy over this period represents Americans' true redistributive preferences, and our new survey results demonstrate the sensitivity of evidence on these preferences to the information that is made salient to survey respondents.

The paper is structured as follows. In Section 1 we describe how we calibrate the income distribution and tax schedules, and we discuss in some detail the estimated combinations of MSWWs and ETIs we infer for the United States from 1979 through 2010. We extend the inference of top-income MSWWs, using more limited evidence, back to 1920. We also distill from these results the three possible explanations for policy noted above. Section 2 uses a subset of the possible patterns of MSWWs to compute revealed preference costs of rising inequality and recession. In Section 3 we discuss suggestive evidence pertaining to each of the three explanations. Section 4 concludes.

# 1 Positive and normative judgments inferred from U.S. tax policy

In this section we show how the Diamond (1998) and Saez (2001) formula for optimal marginal tax rates can be used to infer combinations of the perceived costs and benefits of redistributive taxation underlying modern U.S. tax policy. In that famous formula, these perceived costs and benefits depend on specific quantities, namely the elasticity of taxable income (ETI) and the pattern of marginal social welfare weights (MSWWs). We then describe the data we use for the inference exercise. Finally, we present our results on the possible patterns of the ETI and MSWWs that can explain the evolution of U.S. tax policy since 1979.

## 1.1 Inverting the optimal marginal tax rate formula

As described in Bourguignon and Spadaro (2012), the inversion exercise is best understood as the dual of the standard Mirrlees (1971) optimal taxation problem. We follow their example and focus on the special

case of quasilinear individual utility, a concave social welfare function, and a uniform ETI, as considered in Diamond (1998). In this case the first-order condition for optimal marginal tax rates takes a particularly simple and transparent form. We use the version of that expression derived in Saez (2001), written as a function of observable earnings $y$, the earnings distribution $F(y)$, with assumed density $f(y)$, and the elasticity of taxable income $\varepsilon$:

$$\frac{T'(y)}{1-T'(y)} = \frac{1-F(y)}{\varepsilon y f(y)} \int_y^\infty \frac{1-g(z)}{1-F(y)} dF(z). \tag{1}$$

This is analogous to expression (4) in Bourguignon and Spadaro, though by using Saez's expression we avoid the need to back out the underlying skill distribution.[2]

In this expression, $g(y)$ denotes the MSWW of an individual earning $y$, which we will use to characterize the implicit social preference for redistribution, i.e., the social welfare function.[3] In other words, $g(y)$ is the social welfare generated by a marginal increase in consumption for an agent earning $y$, expressed in terms of public funds. By construction, $\int_0^\infty g(z)f(z)dz = 1$ under the optimal tax policy; the planner is indifferent between a marginal dollar of public funds and a dollar equally distributed across the population. Thus a value of $g(y') = 0.5$, for example, indicates that the planner is indifferent between an evenly distributed \$0.50 of income and a rise in consumption of one dollar for an agent earning $y'$.

The standard Mirrlees approach specifies an individual utility function and social welfare function and solves for the tax function that satisfies (1); here we invert that approach, taking the observed tax function as given and solving for the social welfare function that would rationalize it (see Bourguignon and Spadaro for a thorough discussion of the conditions under which this inversion is possible). To implement this inversion, we rearrange (1) to write:

$$\int_y^\infty \frac{g(z)}{1-F(y)} dF(z) = 1 - \frac{T'(y)}{1-T'(y)} \frac{\varepsilon y f(y)}{1-F(y)}. \tag{2}$$

Given the observed tax code, a calibrated income distribution, and an assumed ETI the right side of this equation can be computed as a function of income. Differentiating (2) with respect to $y$ yields

$$g(y) = -\left(\frac{1}{f(y)}\right) \frac{d}{dy}\left[1 - F(y) - \frac{T'(y)}{1-T'(y)}(\varepsilon y f(y))\right], \tag{3}$$

where the expression in brackets can be computed using numerical differentiation.[4] This provides an estimate of the MSWW $g(y)$ as a function of income.

**Limitations** The Diamond and Saez formula captures the costs and benefits of redistributive taxation in quite a general way. Nevertheless, it has a number of potentially important limitations. We briefly discuss a few of them here.

First, the MSWWs implicit in U.S. tax policy do not necessarily correspond to the MSWWs that represent Americans' "true" preferences for the distribution of income: i.e., those that we would infer from policy generated by a political system that perfectly represented Americans' preferences. We directly address

[2] A technical point: as discussed in Saez (2001), $f(y)$ actually denotes the "virtual"earnings density—that which would obtain if the tax code were linearized around $T(y)$. While the virtual density is not identical to the observed income density, it is closely related, and therefore we view the empirical calibration of the observed income distribution, discussed in Section 1.2, as a good approximation of the virtual density. An alternative approach, followed by Bourguignon and Spadaro, is to infer an underlying ability distribution based on observable earnings and an assumed utility function.

[3] Strictly speaking, MSWWs correspond to the derivative of the full social welfare function, but because the intercept of the social welfare function is immaterial, these weights are sufficient statistics for redistributive preferences.

[4] The algebra behind this derivation is shown in the appendix.

this possibility in depth in Section 3, where we use both existing and new evidence to study whether the preferences implied by policy may be due to a political outcome that poorly represents public opinion.

Second, the optimal tax model used in the approach shown above may be misspecified relative to reality, and we may therefore make faulty inferences of the component factors. One possible misspecification is that the model focuses entirely on the intensive labor supply margin (the question of how much to work), while ignoring the extensive margin (the question of *whether* to work). Although extensive margin elasticities have important implications for optimal tax design (Saez, 2002), we view our approach as a useful simplification for two reasons. First, the extensive margin is particularly important for low incomes, while our analysis focuses primarily on MSWWs on high incomes.[5] Second, representing the perceived distortionary costs of taxation through a single parameter simplifies and facilitates the derivation and exposition of our results. Provided that we view the intensive ETI as a proxy for the overall perceived distortionary costs of taxation, we believe our results would be similar in a richer model incorporating other margins of adjustment. Another possible misspecification is that the model we use is static, while the implications of dynamic factors (such as financial and human capital accumulation) for optimal taxation have received much attention in recent years.[6] As with the extensive margin, to the extent that these factors primarily affect the distortionary costs of taxation, our use of a single parameter to capture perceptions of those costs is flexible enough to (at least roughly) accommodate them. In addition, while the theoretical benefits of a tax system that takes into account earnings histories, age, and other complications may be large, actual policy design thus far is arguably better described by the simpler static Mirrleesian model.[7]

Third, actual tax policy may reflect priorities not well-captured in the standard model, for instance a preference for simplicity as in Smith (1776). One way to interpret our finding that actual policy differs substantially from what a conventional optimal tax model would imply is that we are measuring the implicit quantitative importance and distributional implications of such non-standard factors.

## 1.2 Data

As described above, the estimation procedure for MSWWs depends on the market (pre-tax) income distribution $F(y)$ and the schedule of marginal tax rates $T'(y)$. To compute welfare costs of economic changes, these MSWWs are then applied to disposable income (after taxes and transfer). We obtain these data from two sources—the U.S. Congressional Budget Office (CBO) for pre-tax and post-tax income, and the National Bureau of Economic Research's TAXSIM utility for the marginal tax rate schedule.

**Income distribution** Since 1979 the CBO has produced annual data describing the distribution of market and disposable incomes across U.S. households.[8] These data consists of average market and disposable income levels for households in eight quantiles, partitioned by percentiles 20, 40, 60, 80, 90, 95, and 99. Throughout the paper, all real figures are given in 2010 U.S. dollars as computed using the Personal Consumption Expenditures price index from the U.S. Bureau of Economic Analysis.

The CBO data's coarse level of aggregation creates two important limitations for our purposes. First, we are unable to infer variation in MSWWs within quantiles reported by CBO. This is particularly important

[5] Consistent with this, Bourguignon and Spadaro demonstrate that incorporating an extensive margin yields qualitatively similar results, particularly at the top.

[6] The relevant literature is usually said to start with Golosov, Kocherlakota and Tsyvinski (2003). See Stantcheva (2014) for a recent examination of human capital in the model.

[7] See Golosov, Kocherlakota and Tsyvinski (2003), Farhi and Werning (2007), Weinzierl (2011), and Best and Kleven (2014).

[8] We use the term "disposable income" to indicate income after federal taxes and transfers. This implicitly assumes property taxes and other local taxes are a component of consumption.

for the bottom four quantiles, which each represent 20% of the population. Mitigating this concern, much of our analysis will focus on MSWWs of high earners, where the smaller buckets provide a finer picture of the distribution. Second, we are unable to account for variation in marginal tax rates within the quantiles reported by the CBO. Rather, we will feed the average income levels from each quantile into NBER's TAXSIM utility to obtain marginal tax rate estimates, effectively treating each bucket as a representative agent, weighted by its population share. Although we do not believe these shortcomings qualitatively affect the nature of our results, universal data (such as that used by Hendren (2014)) surely provide a more precise measure of the distribution.

Nevertheless, in two respects the CBO data are particularly well-suited for the purposes of this paper. First, to compute the costs of unequal growth and recessions as we do in Section 2, we need to apply MSWWs to changes in disposable (post-tax) income. The CBO data are intended to provide a consistent and carefully-constructed measure of this income, combining internal U.S. Treasury tax return data with data from the Consumer Population Survey (CPS) of the U.S. Census Bureau. As the CBO argues in its documentation, combining these data sources is important because tax records exclude people who do not file federal tax returns as well as information on some government cash transfers and in-kind benefits that are captured by the CPS, while CPS data are sparse at the upper end of the income distribution relative to tax return data.

Second, the CBO data occupy a prominent place in public debates over taxation and the distribution of income in the United States.[9] This paper seeks to infer the combination of positive and normative judgments from existing policy, so using an official data source that is (and has been) salient for policymakers over the last several decades is arguably preferable to using more precise data that was less visible to them.

To compute the MSWWs, we include all market income other than capital gains in our analysis. Thus, we include dividend income and other capital income. In part, this choice is guided by the policy treatment of this income: capital gains are generally taxed separately from other forms of personal income in the United States, so it is natural to exclude them.[10] We also exclude capital gains to increase comparability to the well-known data on the concentration of income produced by Piketty and Saez (2007), as they also compute series excluding them. Table 1 compares the the shares of market income excluding capital gains that each data source assigns to three groups—the 90-95th percent, the 95-99th percent, and the 99-100th percent of the population—at four points over the period.

| Comparison CBO and Piketty & Saez (2007) data | | | | | | |
|---|---|---|---|---|---|---|
| | 99-100 share | | 95-99 share | | 90-95 share | |
| Year | CBO | P&S | CBO | P&S | CBO | P&S |
| 1980 | 8.0 | 8.2 | 12.0 | 13.0 | 10.5 | 11.7 |
| 1990 | 11.7 | 13.0 | 12.5 | 14.1 | 10.7 | 11.8 |
| 2000 | 14.3 | 16.5 | 13.6 | 15.0 | 10.8 | 11.6 |
| 2010 | 14.0 | 17.5 | 14.1 | 16.3 | 11.2 | 12.6 |

**Table 1:** Comparison of income concentration in top quantiles according to CBO data and Piketty and Saez (2007).

[9]For one recent example, the CBO's report on the distributional impact of tax expenditures in May 2013 presented the data in exactly the same form as the data we use in this paper. That CBO report was widely covered in the press and referred to by policymakers, such as the House Ways and Means ranking member Sander Levin, (see http://democrats.waysandmeans.house.gov/press-release/levin-statement-cbo-report-distribution-tax-expenditures)

[10]In any case, our basic results remain the same if capital gains are included in market income.

As the table shows, these series track each other directionally, though the CBO data indicate a lower degree of income concentration at the very top over time. This difference may be due in part to the focus on households in the CBO data rather than tax-paying units, consistent with a calculation in Hendren (2014). To the extent that our use of CBO data thereby underestimates the incomes of higher earners, we will infer lower MSWWs on high earners for any given ETI, especially over time as inequality has risen.

The process for backing out MSWWs described above requires a continuous income distribution. We calibrate a Pareto lognormal distribution to the eight moments reported in the CBO data each year. This parametric form, introduced in Colombi (1990), fits observed income distributions quite well, both in the middle of the distribution and in the top tail (see Reed and Jorgensen (2004)). The distribution has a Pareto upper tail, the importance of which for computing optimal income taxes is highlighted in Saez (2001) and Diamond and Saez (2011). The Pareto-lognormal distribution is characterized by three parameters, one of which is the Pareto parameter for the distribution toward which the Pareto-lognormal converges in the upper tail. We constrain this parameter to equal the Pareto parameter implied by data reported in Piketty and Saez (2007), which is based on tax returns and thus should represent the top tail of the distribution well. We calibrate the remaining two parameters to minimize the squared errors between the observed means within the quantiles reported by CBO and the means within those quantiles predicted by our calibration. Figure 1 displays averages reported by CBO and our best-fit calibration over the range of years for which we have data.

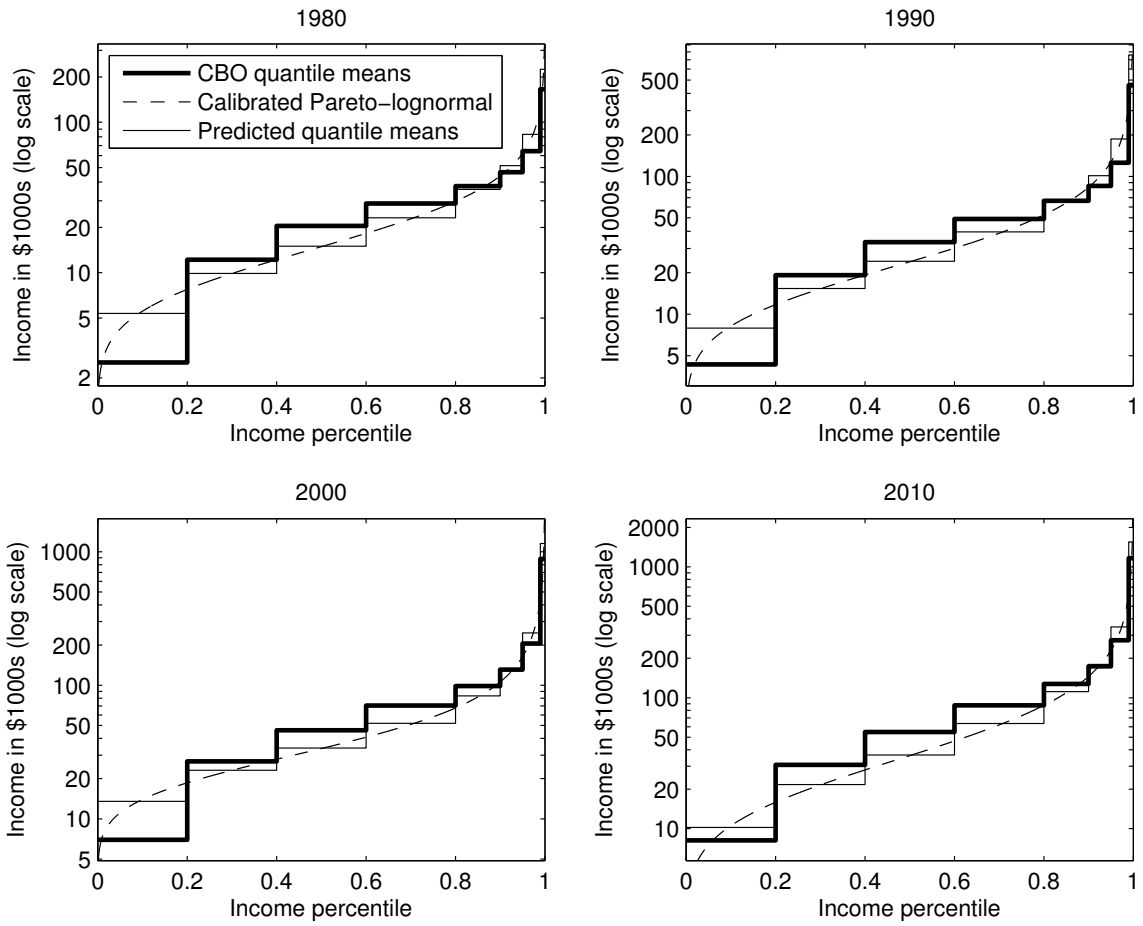



**Figure 1:** The reported and calibrated distribution of market income in the US for years 1980, 1990, 2000, and 2010. The bold lines represent the mean income within the quantiles reported in CBO data. The dashed line is the best-fit Pareto-lognormal distribution, selected by minimizing the sum of squared errors between the reported quantile means and those predicted by the calibrated distribution—denoted here by the thin solid line. (Income is reported in nominal terms.)

The figure indicates that this fit is imperfect—in particular, our calibration underestimates incomes in the center of the distribution, while overestimating incomes at the very bottom. Nevertheless we view this approach as a good approximation given the sparse nature of the CBO data; in particular, since the optimal top tax rate depends only on the ETI, the top limiting MSWW, and the Pareto parameter (see Saez (2001)), our use of the Piketty and Saez Pareto parameter estimates ensures that our computed top welfare weights approach those under the true income distribution.[11] Our calibrations capture the well-known increase in concentration of income over this period, as we show in the Appendix through Lorenz curves for 1980, 1990, 2000, and 2010.

[11] Outside of the top tail, an underestimate of income at a particular percentile is likely to bias upward MSWWs above that income relative to those below. Intuitively, underestimating income at a given percentile is equivalent to overestimating $F(y')$, where $y'$ is the estimated income level. Then raising the marginal tax rate at $y'$ will have smaller redistributive benefits than the calibration suggests, and thus a given tax rate must correspond to a greater commitment to redistribution than the calibration suggests. This effect is complicated by the dependence of the optimal tax rate on the density of earners at $y'$, so a directional claim cannot be made with certainty, but provided the latter effect is small, this reasoning suggests our MSWWs are likely to be biased down at low incomes, biased upward for middle incomes, and approximately correct at high incomes.

**Marginal tax rate schedule** We use the NBER's TAXSIM tool to obtain marginal tax rates at each of the CBO's reported income levels from 1979 through 2010. We obtain the marginal tax rate on taxpayer earnings by assigning all of the household's income to the taxpayer in a joint-filing household with two adults of working age, two dependents under the age of 17, and no state tax liability or capital gains. We construct a piecewise-constant marginal tax schedule from the TAXSIM marginal tax rates reported for a fine grid of incomes between \$1 and \$10 million; Figure 2 displays the marginal tax rate schedule for the years 1980, 1990, 2000, and 2010.

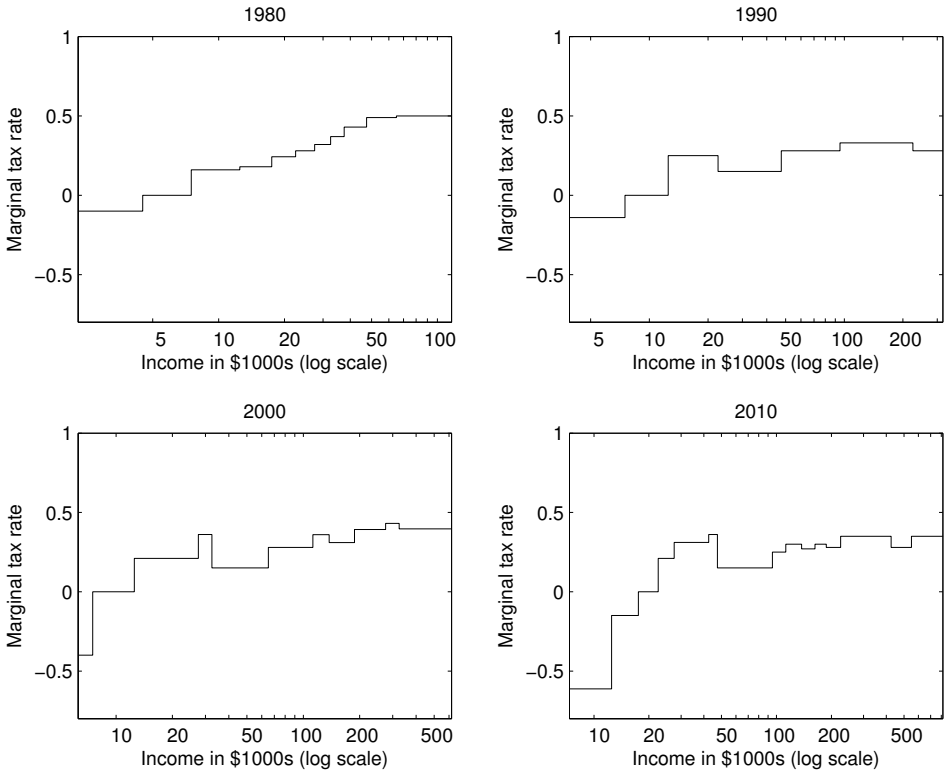



**Figure 2:** Marginal tax rate schedule as reported by NBER's TAXSIM for 1980, 1990, 2000, and 2010. Taxes on capital gains and at the state and local level are not included. Tax rates are computed for a family of two adults, filing jointly, with two dependents under age 17.

The figure shows substantial negative marginal tax rates at the bottom of the income distribution, driven by the Earned Income Tax Credit, and top marginal tax rates between 28% and 50%.

## 1.3 Implicit MSWWs and perceived ETIs: three explanations for U.S. policy

Using the theory and data described in the preceding subsections, we can compute the combinations of MSWWs and perceived ETIs implicit in U.S. tax policy for each year from 1979 through 2010.

We will repeatedly compare the implicit MSWWs to those that are implied by conventional social welfare functions (such as in Diamond (1998) or Saez (2001)). By "conventional" we have in mind two features: MSWWs are everywhere nonnegative, and MSWWs approach zero toward the top of the income distribution. These two features arise from the intuitively plausible and commonly-imposed assumptions that society's preferences are Pareto efficient (i.e., non-negative weight is given to an increase in any person's consumption,

all else the same) and that the marginal social welfare of income is decreasing (which may be due to diminishing marginal utility at the individual level or a social judgment). Similarly, we will compare the implicit ETIs to the "conventional" levels estimated in the empirical labor literature. The Congressional Budget Office (2012) for example, reviews that literature and finds "substitution elasticities for the total population that range from 0.1 to 0.3."

### 1.3.1 General patterns of MSWWs and ETIs

First, we characterize the general features of the inferred MSWWs and ETIs. We begin by assuming a fixed ETI and inferring the MSWW schedules consistent with policy over time. Figure 3 shows the results. To generate this figure, we begin by computing the right side of (2), which represents the average MSWW above a given level of income (called "S weights" in Bourguignon and Spadaro (2012)) and which we denote with the function $S(y) = \int_y^\infty g(z)dF(z)/(1-F(y))$. For the interested reader, we show $S(y)$ for 1980, 1990, 2000, and 2010 in the Appendix. We then differentiate $S(y)$ to compute the MSSW schedule $g(y)$ as given in equation (3). Figure 3 shows the averages of these MSWWs within each of the eight CBO income groups for 1980, 1990, 2000, and 2010.[12]

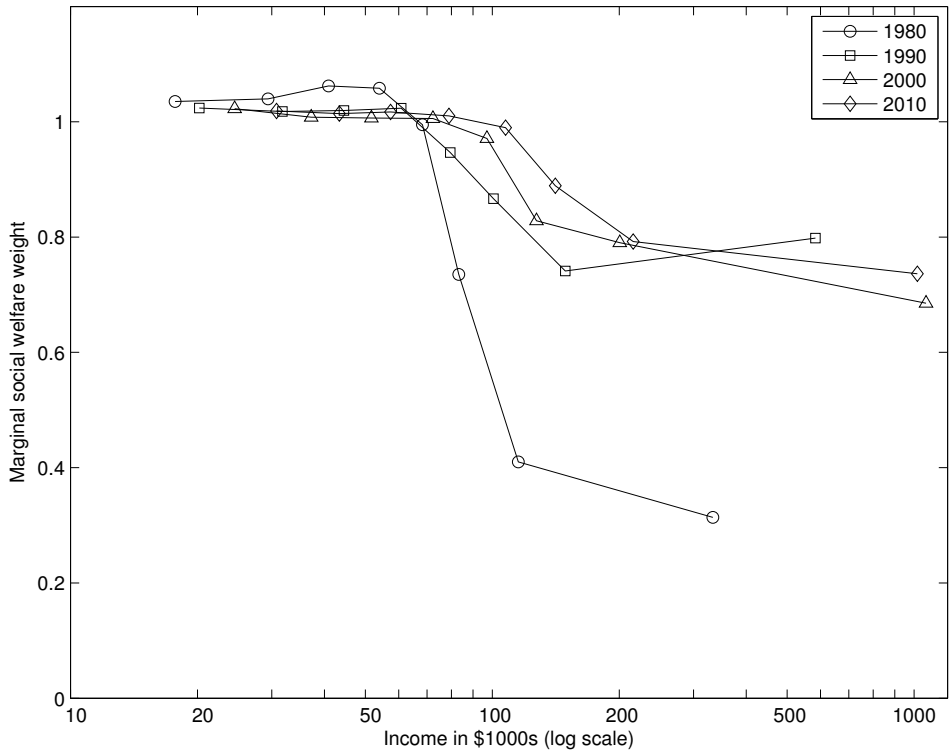



**Figure 3:** Average MSWWs, plotted against disposable income in 1980, 1990, 2000, and 2010. Points represent the average MSWWs within eight partitions of the market income distribution as reported by CBO, bounded by the following percentiles: 20, 40, 60, 80, 90, 95, 99. These averages are plotted against average income after federal taxes and transfers within each partition, in real 2010 dollars. Weights are computed assuming an elasticity of taxable income of 0.3.

[12]Since these weights represent the marginal social value of consumption, they are properly viewed as a function of real, disposable income, and thus are plotted against income after federal taxes and transfers, as reported by the CBO. As with market income, only average disposable income is reported within each quantile, so we plot the average MSWWs within each market income quantile against reported average disposable income within that quantile. For the lowest quintile, we average MSWWs over the values of $y$ such that $0.01 \leq F(y) \leq 0.2$, to avoid numerical issues from MSWWs approaching infinity as $y$ approaches zero.

Figure 3 demonstrates that implicit average MSWWs are generally positive and decreasing with income, as conventionally assumed. It also shows that there is substantial variation over time in the shape of the implicit average MSWW schedules at high incomes, with a marked increase in the average MSWW for high earners between 1980 and 1990 that persists over the next two decades. This change is a result of both rising inequality in pre-tax income and falling top tax rates; the appendix decomposes the influence of these two factors, demonstrating that much of the change can be explained by tax reforms alone.[13]

Figure 4 shows how these results change for different ETI values. Average MSSWs for 1980 and 2010 are displayed for ETIs ranging from 0.1 to 0.6.

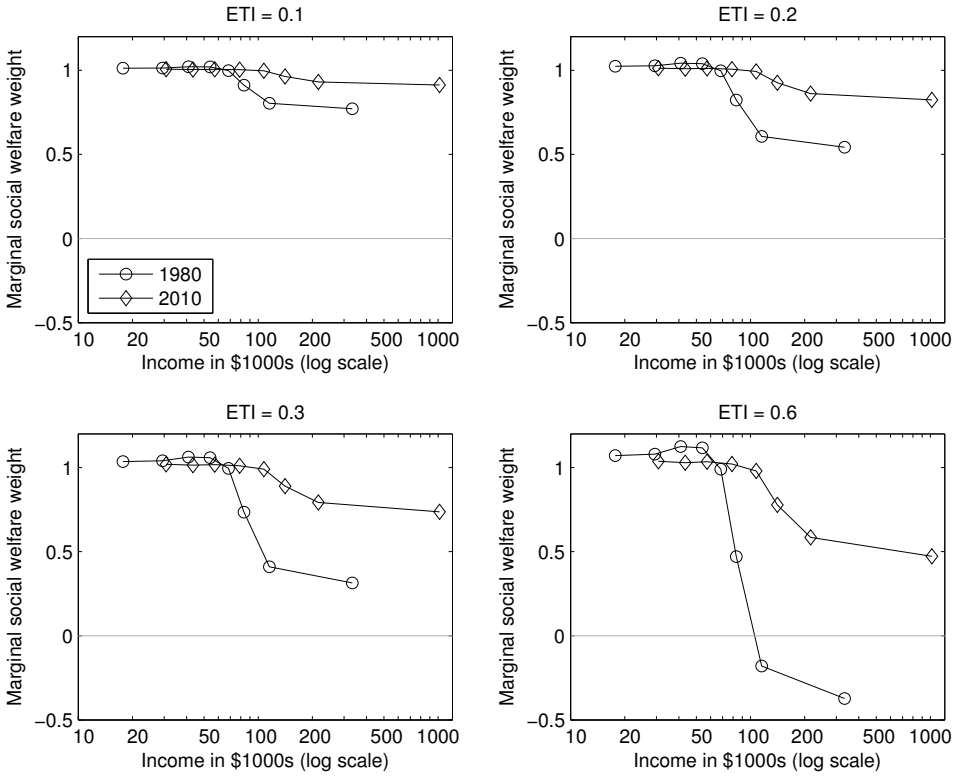



**Figure 4:** Average MSWWs in 1980 and 2000 for a range of elasticity assumptions, plotted against income after federal taxes and transfers within each partition, in real 2010 dollars.

Figure 4 shows that the results are highly sensitive to the assumed ETI, especially at the top of the income distribution. Notably, for an ETI of 0.6, the 1980 tax schedule is not Pareto efficient, placing negative average MSWWs on earners in the top decile. At low ETIs, on the other hand, average MSWWs on high incomes are above 0.5, in contrast with the conventional assumption that marginal social welfare value of consumption approaches zero at high incomes.

### 1.3.2 MSWWs at high incomes

Figures 3 and 4 show that high-income MSWWs provide a simple way to distinguish MSWW schedules. For simplicity, we denote by $g_t^*$ the average MSWW on the top one percent of income-earners in year $t$, corresponding to the rightmost point on the average MSWW schedules in Figures 3 and 4. Formally, let $y_t^*$

[13]Thanks to Diana Moreira for suggesting this decomposition.

be the 99th percentile of market income in year $t$, so that $F_t(y_t^*) = 0.99$, then $g_t^* = \int_{y_t^*}^{\infty} g_t(z)dF_t(z)/0.01$. Figure 5 makes clear the negative relationship between $g^*$ and the perceived ETI required to explain the U.S. data for any one year.

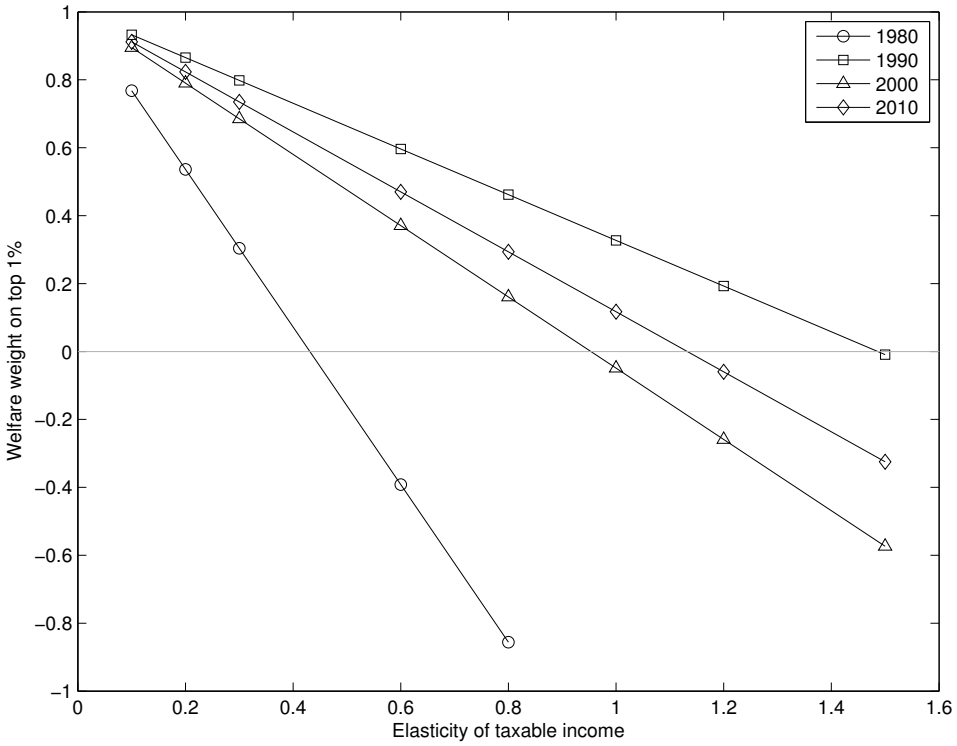



**Figure 5:** The implicit welfare weight attributed to the top 1% of earners under various elasticities of taxable income for years 1980, 1990, 2000, and 2010.

Intuitively, a larger $g^*$ means less willingness to redistribute from high earners and thus exerts downward pressure on high-income marginal tax rates, while a lower perceived ETI means less concern for the distortionary cost of redistribution and thus exerts an offsetting upward pressure on high-income marginal tax rates.

### 1.3.3 Three explanations

We can use Figure 5 to distill the results of our inference exercise into three possible explanations for recent U.S. tax policy, all of which require either MSWWs or ETIs to take unconventional values.

The first possible explanation is that the perceived elasticity of taxable income rose sharply above consensus estimates, while preferences for redistribution were stable and at conventionally-assumed levels over time. This explanation can be seen by looking horizontally across Figure 5 at a given value of $g^*$, and in particular at $g^* = 0.01$. Doing so reveals the variation in the perceived ETI required to explain U.S. policy over this period when high-income MSWWs take their conventional values. Figure 6 illustrates the implications of this possible explanation of our findings in greater detail. It plots the time-varying values of the perceived ETI that would be consistent with each of a number of constant MSWWs on the top 1% of income earners, i.e., $g^* = \{0.01, 0.20, 0.40, 0.60\}$. (Note that the smallest of these values is the "conventional" assumption).




**Figure 6:** Implied elasticities of taxable income which would place a welfare weight on the top 1% of earners, denoted $g_t^*$, of 0.01 (the "conventional" case), 0.2, 0.4, or 0.6.

This figure demonstrates that a substantial increase in the perceived ETI is required to explain the observed evolution of the tax code. In particular, for the the implied $g^*$ to remain at 0.01, the perceived ETI would have to rise from 0.42 to 1.1 between 1979 and 2010. Figure 6 also reveals the impacts of the 1986, 1993, and 2001-2003 tax reforms (which, respectively, lowered, raised, and lowered top marginal tax rates).[14] To the extent that one views these sharp changes as partisan shifts that cross from one side of the true "median" perceived ETI, Figure 6 provides a sense of the corresponding value for the true median.

We can extend the analysis in Figure 6 farther back in time. We use data from Piketty and Saez (2007) to calibrate the Pareto parameters at the top of the income distribution from 1916 to 2012. We then use the U.S. statutory marginal tax rate schedule, an assumed value for the MSWWs on the top 1 percent of income earners, and the simplified formula for the top marginal tax rate from Saez (2001) to back out the implicit perceived ETI over the entire period. Figure 7 shows the results.

[14]Though the CBO data do not extend past 2010, the more recent increases in top marginal tax rates would cause a downward shift in the perceived ETI at each value of $g^*$.

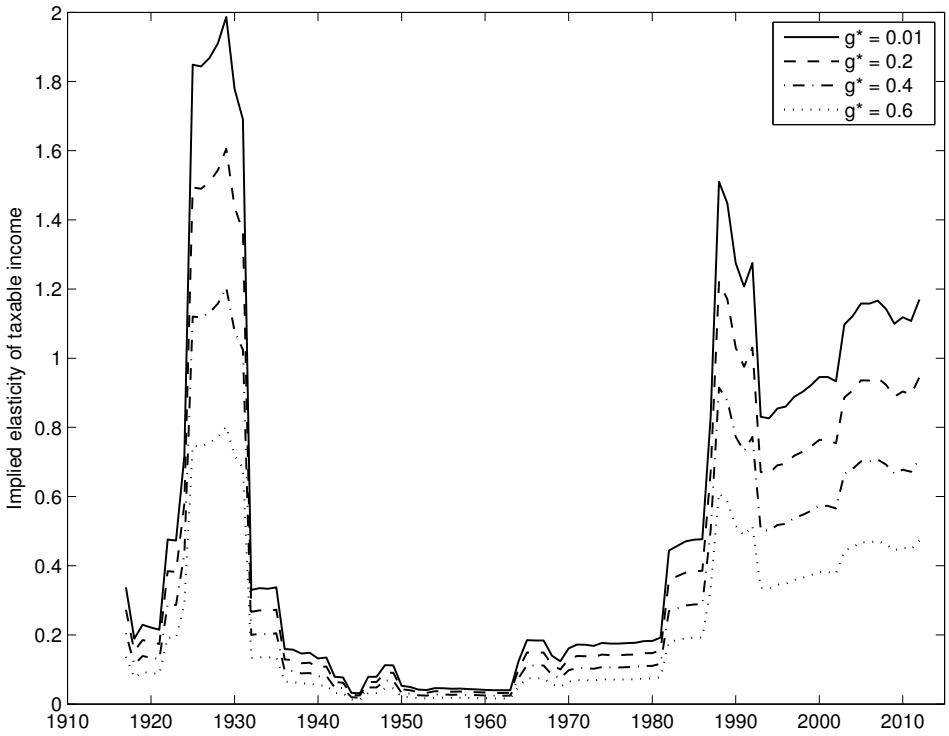



**Figure 7:** Implied elasticities of taxable income which would place a welfare weight on the top 1% of earners of 0.01, 0.2, 0.4, or 0.6, for the history of US income tax. This figure is constructed using the highest marginal tax rate, which at times did not apply to the full top percentile of income earners; in that case this represents the elasticity necessary to rationalize a given top asymptotic MSWW.

As this figure makes clear, the perceived ETI implied by U.S. policy for the mid-20th century was extremely low for a wide range of high-earner MSWWs and varied enormously over the full span of the century unless the value of high-earner MSWWs was very large.

The second possible explanation is that political support for redistribution among Americans may have decreased sharply below conventionally-assumed levels since the 1970s while the perceived distortionary costs of taxation were stable at conventional levels. This explanation can be seen by looking vertically up and down Figure 5 at a given ETI, and in particular at ETIs ranging from 0.1 to 0.3. Doing so reveals the variation in top-income MSWWs that is required to explain policy over this period when ETIs take their conventional values. Figure 8 shows these results in greater detail.

In this figure, the inferred values of $g_t^*$ over the last 25 years substantially exceed those from the early 1980s, for any given ETI. That is, the perceived benefit of redistribution apears to have sharply decreased since the 1980s. For conventional values of the ETI, the implicit political support for redistribution fell well below conventionally assumed values sometime between 1980 and 1990 and has remained there since.




**Figure 8:** Implied social welfare weight on the top 1% for a range of elasticities of taxable income.

The third possible explanation is something of a mixture of the first two. If the perceived distortionary costs of taxation rose a bit over this period, but stayed within conventional ranges, a smaller decrease in the perceived benefits of redistribution is required to explain policy. At the same time, those relatively stable perceived benefits would have been at unconventionally low levels throughout the period. Mathematically, the value of the perceived ETI is directly related to one minus the social welfare weight and inversely related to the marginal tax rate. For a larger weight on high earners, variation in the marginal tax rate translates into smaller (absolute) effects on the perceived ETI, as the levels of the ETI are lower. This explanation can be seen by focusing on the upper left corner of Figure 5, where high values of $g^*$ generate a narrower required range of ETIs to explain the evolution of policy.

We will explore each of these three possible explanations empirically in Section 3.

# 2 Applications: policy-inclusive welfare costs of unequal growth and unequal recessions

One reason for computing MSWWs is that they can be used for making interpersonal comparisons, a crucial component of policy evaluation. Ordinarily such calculations are controversial because they are sensitive to the modeler's assumptions about the social preferences for redistribution (i.e., social welfare function). MSWWs inferred from policy are in principle able to provide a more objective basis for comparison, and in this section we use our results from above to measure of the welfare consequences of changes in the policy-inclusive (i.e., after taxes and transfers) distribution of income due to unequal growth and recessions in the United States since 1979. We show that these calculations are vulnerable to a different form of sensitivity:

they are highly sensitive to which explanation for the policy path is assumed to be correct (e.g., the assumed ETIs) and the year from which MSWWs are derived.

## 2.1 Costs of unequal growth

From 1979 to 2010, US disposable income has grown by an average 1.5% annually. But as is well-known, income growth has been highly concentrated among high earning households. Figure 9 provides a graphical view of this evolution, plotting average income after federal taxes and transfers for the eight quantiles reported by the CBO.

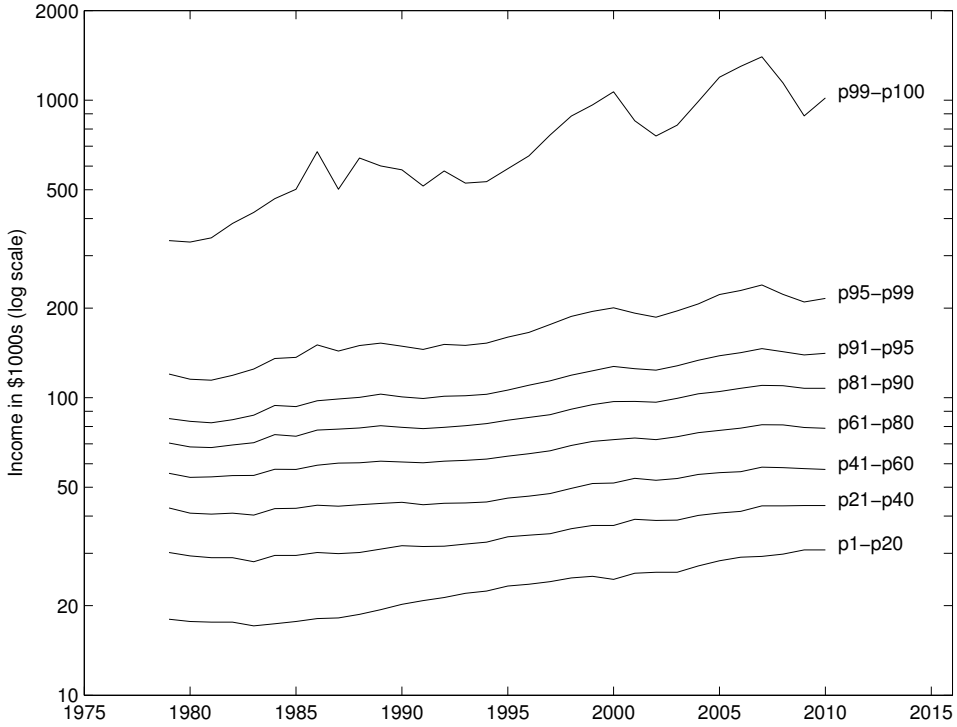



**Figure 9:** Growth in inequality of disposable income, 1979–2010. Each line represents the evolution of average income after federal taxes and transfers within eight quantiles of market income, as reported by CBO. Income is plotted on the vertical axis in log scale. The top percentile exhibits the strongest positive trend, reflecting rising inequality.

In this section, we use revealed social preferences from various years to compute the implied welfare costs of that unequal growth.

We compute the cost of rising inequality by asking a simple question: How much economic growth would be willingly sacrificed in order to prevent rising inequality? Answering this question requires a means of trading off gains to households with differing income levels, and for this we use the MSWW schedules derived in the previous section.

Note that this is not the only possible definition of the "cost of unequal growth". In particular, our measure will compute zero cost of inequality even if pre-tax income inequality grows, provided the tax system redistributes those gains to prevent inequality in disposable income from rising. An alternative question of interest asks how much economic growth society would be willing to sacrifice to obtain equal growth in wages—i.e., technological change that equally benefits all points of the ability distribution. Here

we focus on the former question, as it can be answered using the MSWW schedules derived in the previous section, whereas the latter requires a fuller structural model of labor supply choice.

To implement our procedure, we weight changes in average disposable income within each of the eight CBO quantiles by the set of corresponding MSWWs as derived in Section 1.3. Letting $\{g^i\}_{i=1}^8$ denote the vector of welfare weights across quantiles $i = 1 \dots 8$, and $c_t^i$ the mean disposable income of quantile $i$ in year $t$, the change in welfare from year $m$ to $n$, denoted $\Delta W_{m,n} \big| \{g^i\}_{i=1}^8$ is

$$\Delta W_{m,n} \big| \left\{g^i\right\}_{i=1}^8 = \sum_i \left(c_n^i - c_m^i\right) g^i f(y^i). \tag{4}$$

Next we calculate the counterfactual change in welfare $\Delta \tilde{W}_{m,n}(\rho)$ that would result from a given equally-distributed annual growth rate $\rho$ from years $m$ to $n$ under a given vector of MSWWs:

$$\Delta \tilde{W}_{m,n}(\rho) \big| \left\{g^i\right\}_{i=1}^8 = \sum_i \left(c_m^i (1+\rho)^{n-m} - c_m^i\right) g^i f(y^i). \tag{5}$$

Note that this definition of "equally distributed growth" holds the growth rate of disposable income constant across quantiles, which would result in constant values for many metrics of income inequality, including inter-quantile spreads and the Gini index.

By solving for the $\rho$ such that $\Delta \tilde{W}_{m,n}(\rho) = \Delta W_{m,n}$, we can compute the constant-inequality growth rate that would yield the same gain in social welfare as that experienced in reality. Of course, this value is strongly dependent on the chosen vector of MSWWs, and thus will vary depending on both the explanation for policy changes (i.e., the path of ETIs) and which year's revealed MSWWs are chosen. Table 2 provides a view of this sensitivity, reporting the share of growth which would be sacrificed to prevent rising inequality over the period 1979–2010 using MSWWs computed in 1980, 1990, 2000, and 2010, for a range of ETIs.

| Cost of Rising Inequality, 1979–2010 | | | | |
|---|---|---|---|---|
| | ETI | | | |
| | 0.1 | 0.2 | 0.3 | 0.6 |
| 1980 MSWWs | 3.8% | 8.0% | 13% | 29% |
| | ($1420) | ($2930) | ($4550) | ($10,200) |
| 1990 MSWWs | 1.2% | 2.4% | 3.6% | 7.7% |
| | ($430) | ($880) | ($1340) | ($2810) |
| 2000 MSWWs | 1.6% | 3.2% | 4.9% | 11% |
| | ($580) | ($1180) | ($1810) | ($3840) |
| 2010 MSWWs | 1.4% | 2.8% | 4.3% | 9.0% |
| | ($510) | ($1040) | ($1580) | ($3300) |

**Table 2:** Cost of rising inequality, 1979–2010. This table shows the share of realized growth which would be willingly sacrificed to preserve 1979 level of inequality in disposable income, using MSWWs implicit in the tax code as computed in 1980, 1990, 2000, and 2010, for a range of elasticities of taxable income. Figures in parentheses represent the corresponding reduction in average income per capita, in 2010 dollars.

Table 2 shows that using an ETI of 0.3 (with 2010 MSWWs) generates a cost of unequal growth three times as large as using an ETI of 0.1. Intuitively, if the ETI is larger, the implicit MSWWs reflect a greater aversion to inequality, raising the cost of unequal growth. Similarly, under our baseline elasticity value of

0.3, the welfare cost of unequal growth is 13% under 1980 MSWWs, yet only 4.3% under 2010 MSWWs. This difference is due to the significantly greater inequality aversion embedded in the former weights.

This measure of the cost of unequal growth is closely related to the important work of Hendren (2014), as noted in the Introduction. Hendren also estimates a cost of the growth in inequality since 1980, finding that adjusting for increased inequality would offset 15–20% of growth, and it is informative to examine how our calculation differs from his.

First, the conceptual motivation behind Hendren's approach is different, leading him to focus on unequal changes in income levels rather than unequal rates of income growth.[15] In the context of rising inequality, Hendren's approach computes the number of dollars by which every individual's disposable income would have risen if all economic gains had been equally distributed—in dollar terms—over the period between 1980–2010. This calculation is clearly related to the question we posed at the beginning of this section: "How much growth would be willingly sacrificed to prevent rising inequality?" Note that equally distributed gains, however, would reduce inequality under most measures (such as interquantile spreads or the Gini index). Our calculation is thus complementary to Hendren's, as ours indicates the amount of growth that would be given up to achieve equal percentage changes in disposable income, rather than equal level changes.

A second reason our approach differs from Hendren's is that we use different data, employing pre- and post-tax measures of income, averaged across eight quantiles, from CBO, rather than the universe of tax returns, as Hendren does. Our data is less granular and is unable to fully capture heterogeneity in marginal tax rates conditional on income (due to differing family structures, deductions, etc.). At the same time, by employing data from the CPS, the CBO may be better able to capture in-kind transfers and data on households not filing tax returns.

In a calculation designed to replicate Hendren's, using 2010 MSWWs (computed with an ETI of 0.3) to weight the pre-tax gains in levels to each CBO quantile since 1980, we find an inequality cost of 16% compared to Hendren's estimate of 15–20%. The similarity of these estimates provides reassurance that the coarseness of our data is not of primary importance to our results.

## 2.2 Costs of Unequal Distribution of Recessions

A similar methodology can produce estimates of the cost of the unequal policy-inclusive distribution of business cycle downturns. These computations have important implications both for stabilization policy and economic research. If the economic costs of business cycles are small, as suggested by Lucas (2003), then efforts to understand and further mitigate them may be less necessary than if they are large. In particular, we are interested in quantifying the welfare costs from the unequal incidence of recessions, which are not captured by the representative agent approach employed by Lucas.

We start by calculating, for four recessions between 1979 and 2010, the loss in social welfare under a given set of MSWWs, relative to a counterfactual in which no recession occurred. A graphical representation of this counterfactual is shown in Figure 10, in which we compute a counterfactual assuming smooth income growth over the recessionary period.

[15] Hendren's approach is motivated by the classic Kaldor-Hicks compensation principle, which holds that one environment dominates another if the "winners" in the former could hypothetically compensate the "losers", leaving everyone better off. Hendren proposes an intuitive revision: since actually implementing such transfers through, say, reforms to the income tax would have distortionary effects, those distortions should be included when making welfare comparisons. This can be accomplished by weighting surplus to each individual by an "inequality deflator", representing the distortionary cost of transferring a dollar from the population at large to a specific point in the income distribution. This inequality deflator turns out to be isomorphic to $g(y)$ if a social welfare function is assumed to exist (though an innovation of Hendren's work is its applicability in contexts without a rationalizing social welfare function).

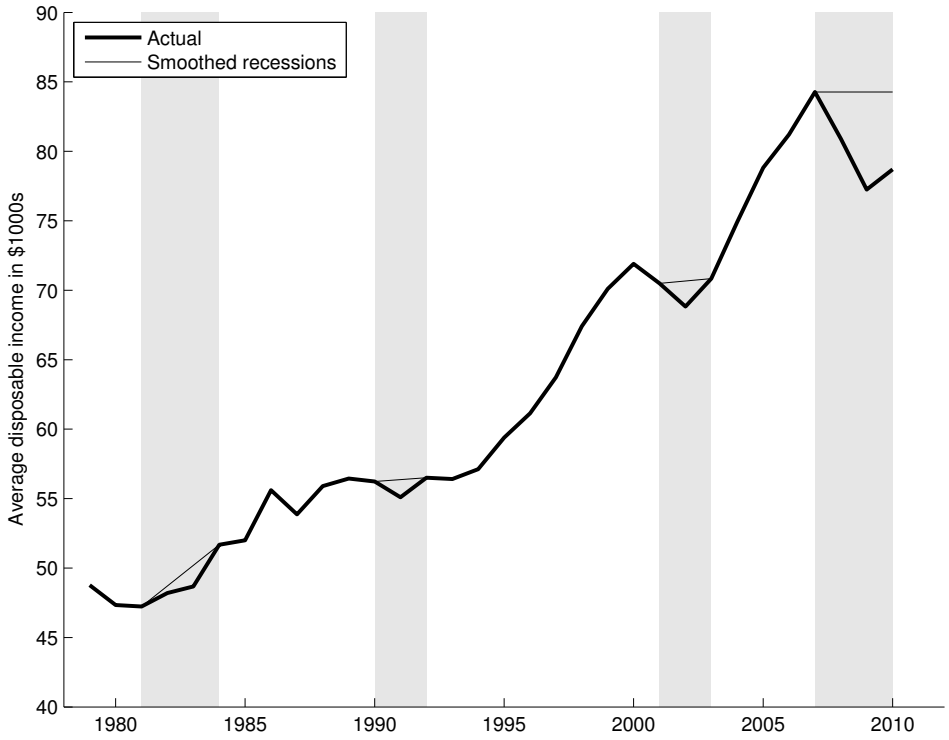



**Figure 10:** Recessions between 1979 and 2013. The bold line plots average income after federal taxes and transfers; recessionary periods are shaded. To compute welfare under a counterfactual without a recession in each case, we connect a straight line between the income levels at the beginning and end of each recession, as shown above. Because CBO data is only available through 2010, at which time income was still substantially lower than the pre-recession peak in 2007, in that case we simply extrapolate a flat level of mean income. (Income is reported in real 2010 dollars.)

We perform similar smoothing for the income path within each of the eight CBO quantiles. We then weight the lost income in recessions (relative to the counterfactual) by each quantile's MSWW to compute the welfare loss of recessions. Formally, we calculate the social welfare cost of those lower incomes using the MSWW vector $\left\{g^i\right\}_{i=1}^8$ during, say, the 2001–2003 recession, as:

$$\Delta\tilde{W}_m|\left\{g^i\right\}_{i=1}^8 = \sum_{m=2001}^{2003}\sum_i \left(\tilde{c}_m^i - c_m^i\right) g^i f(y_m^i), \tag{6}$$

where $c_m^i$ and $\tilde{c}_m^i$ denote the actual and counterfactual (non-recession) disposable incomes for individual $i$ in year $m$.

Next, we find a hypothetical recession that yields the same loss in social welfare but in which all types of individuals suffer a common proportional loss of income off of the constant-growth trends $\tilde{c}_m^i$. Formally, for a recession that lies between years $m$ and $n$, we solve for the value of $\alpha$ that satisfies the following:

$$\sum_{t=m}^{n}\sum_i \alpha\tilde{c}_m^i g^i f(y_m^i) = \Delta\tilde{W}_m|\left\{g^i\right\}_{i=1}^8, \tag{7}$$

which has the same right-hand-side as expression (6) . The value of $\alpha$ that satisfies (7) can be interpreted as the uniform recessionary shock (as a percent of no-recession income), that would have yielded the same social welfare loss as the actual recession.

In Table 3, we show the results for four recessionary periods using MSWWs from 1980 and 2010. All of these calculations assume an ETI of 0.3.

| Cost of Recessions, 1979–2010 | | | | |
|---|---|---|---|---|
| | Disposable income loss | | Equally-distributed equivalent | |
| Recession | Average | Middle quintile | 1980 MSWWs | 2010 MSWWs |
| 1981–1984 | 4.3% | 4.4% | 2.1% | 2.1% |
| 1990–1992 | 2.3% | 1.5% | 1.4% | 2.0% |
| 2001–2003 | 2.6% | 1.4% | 1.8% | 2.3% |
| 2007–2010 | 12% | 1.4% | 2.8% | 5.1% |

**Table 3:** Cost of recessions, 1979–2010. The first column shows the share of annual pre-recession disposable income lost over the course of a recession, and the second shows the loss to the quintile around the median earner. The third and fourth columns show the equally-distributed income loss which would generate the same welfare cost, using 1980 and 2010 MSWWs. These results assume an elasticity of taxable income equal to 0.3.

Two features of these results stand out. First, the equally-distributed shock is generally smaller than the actual (unequal) recessionary shock, especially in recent recessions. This is driven by the fact that reductions in disposable (not necessarily market) income during recessions are concentrated on high earners. Indeed, as can be seen from Figure 9, the CBO reports that disposable income did not fall at all during the Great Recession for the two lowest quintiles. Since MSWWs from all years place a lower weight on individuals at the top of the income distribution than at the bottom, they imply a lower welfare cost of recessions than the average income loss, which places full weight on the large losses of top earners.

The second notable feature of Table 3 is the sensitivity of the results to the year in which MSWWs are computed. The Great Recession, for example, exhibits a welfare cost over 80% larger under 2010 MSWWs as under 1980 MSWWs. This sensitivity further highlights the importance of understanding the evolution of implicit MSWWs, in order to determine whether one year's MSWWs reflect "true" social preferences, or whether those preferences have indeed changed over time.

# 3 Empirical evidence on possible explanations for the evolution of policy

In this section we present some available evidence on the three possible explanations highlighted in Section 1. We emphasize that this evidence is far from conclusive and that we view this effort as a preliminary examination upon which we hope future work, with better data, will be able to build. In particular, a clear lesson from this section's analysis is that more data, especially over time, on the perceived costs and benefits of redistributive taxation would be invaluable to researchers interested in the questions we ask in this paper.

## 3.1 Explanation 1: Increases in the perceived distortionary costs of taxation

We begin by considering perhaps the simplest interpretation of our results: that perceptions of the costs of redistributive taxation have risen over time while MSWWs have remained stable and at conventional levels. To explore this possibility, we draw on a range of sources to gauge the perceived incentive effects of taxation and the extent to which complementarities across workers are deemed important.

### 3.1.1 The incentive effects of taxation (i.e., ETIs)

First we look for direct survey evidence on popular perceptions of the incentive effects of taxation. Unfortunately, direct data on the general voting public's perceptions of these costs are very limited, with the World Values Survey providing what little, highly-imperfect evidence we have found. The World Values Survey (WVS) asked respondents in the United States in 1990, 1995, 1999, 2006, and 2011 to indicate where on a 10-point scale they would "place their views," where 1 was "Incomes should be made more equal" and 10 was "We need larger income differences as incentives." This question is far from ideal for the purposes of gauging the perceived elasticity of taxable income, but it does prompt the respondent to consider the incentive effects of redistribution. To that extent, a higher numeric response reflects a greater perceived distortionary cost of taxation. Figure 11 shows the distribution of responses to this question, as well as the mean response, from 1990 through 2011.

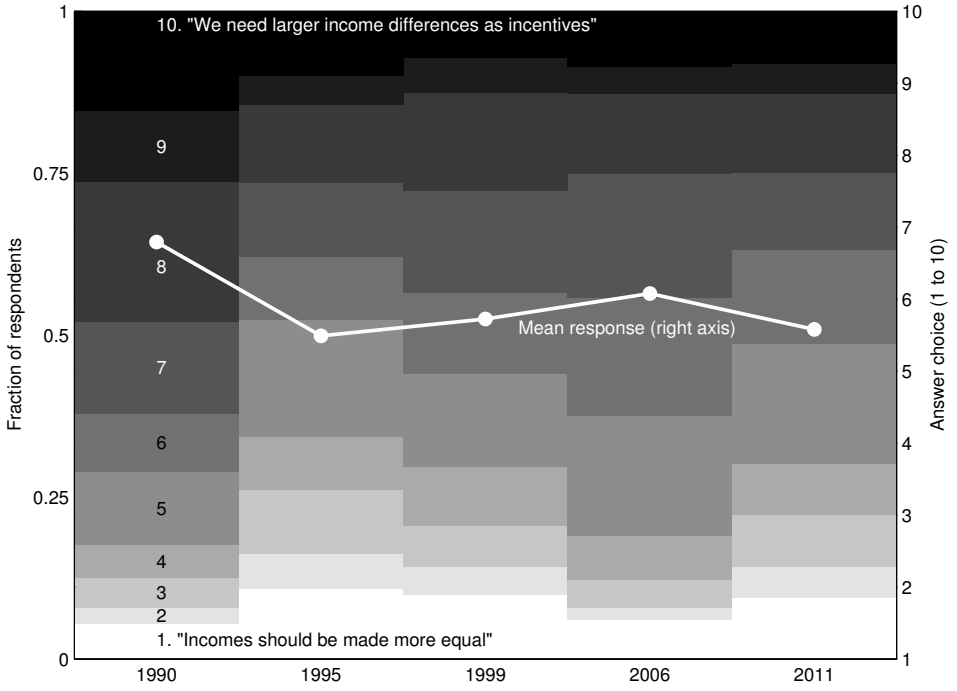



**Figure 11:** Responses to the World Values Survey question in which 1 is "Incomes should be made more equal" and 10 is "We need larger income differences as incentives." The solid line shows the mean answer in each year, as measured by the right axis.

The evolution of responses shown in this figure is qualitatively consistent with Figure 6 above, where the implied ETIs were shown. Of course, the key test of this interpretation would be whether it could match the more dramatic shifts from the 1980s to the period shown in Figure 11, and one implication of this paper is that the generation of such evidence would be valuable in the future.

Given the limited data available on public perceptions, we turn to the extensive academic literature estimating the true (not necessarily the perceived) incentive effects of taxation, with the implicit idea that public perceptions may track professional opinion. In particular, we focus on surveys of this literature by the CBO (1996), Blundell and MaCurdy (1999), Keane (2011), and Chetty (2012) to gauge the estimated compensated elasticity of labor supply (or taxable income) of prime-aged men over time.[16] The CBO's 1996

[16]Excluding estimates for female elasticities means that we ignore the dramatic increased participation of women in the

report provides a survey of surveys, giving the mean labor supply estimates from nine literature reviews spanning 1979 through 1993. Figure 12 plots the mean estimates reported by the survey papers in the CBO report and three other surveys, along with the point estimates from each of the papers reviewed in the very recent Keane and Chetty papers.

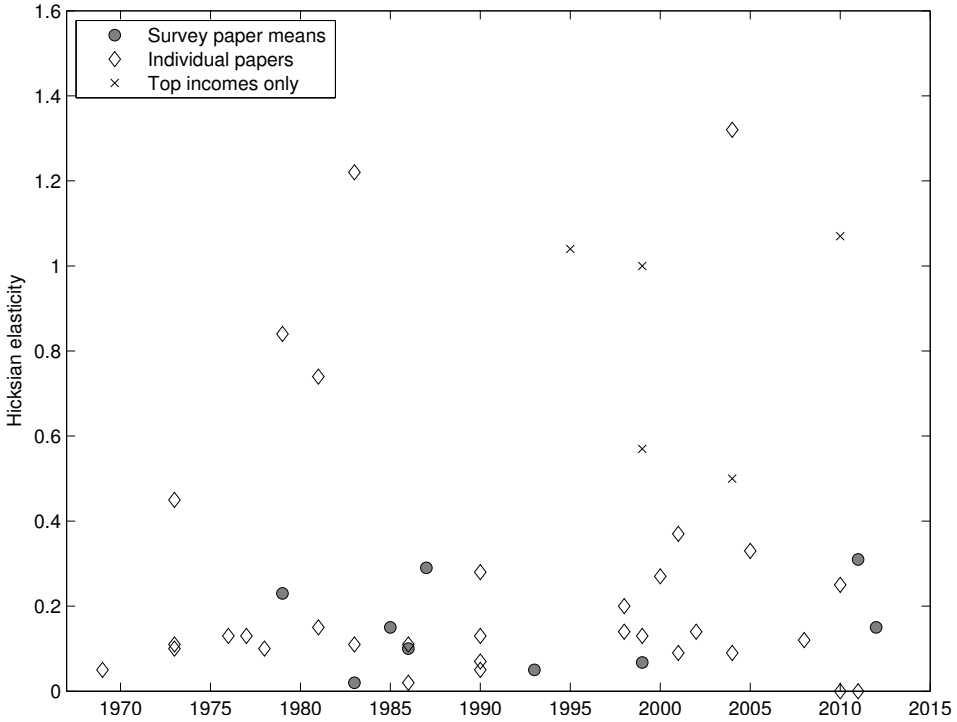



**Figure 12:** Elasticity estimates in the economics literature over time, drawn from the following survey papers: CBO (1996), Blundell and MaCurdy (1999), Keane (2011), and Chetty (2012).

The evidence in Figure 12 provides little support for a dramatic increase in the distortionary costs of taxation over this period. First, if we focus on the survey articles (the solid circles in the figure), there is no sharp rise in estimates between the early 1980s and later years.[17] The apparent consensus in early years, namely that elasticities were quite low, appears still to hold today: as Saez, Slemrod and Giertz (2012) wrote just a few years ago, "With some exceptions, the profession has settled on a value for this elasticity [of labor supply] close to zero."

One possibility raised by the data in Figure 12 relates to the estimated elasticities for top incomes, shown as "x" symbols in the figure, which arose relatively recently. (The first such estimate displayed in the figure is Feldstein, 1995.[18]) Though consensus on the true elasticity of top earners remains elusive, the estimates shown in the figure are well above almost all recent estimates of the overall average elasticity. In principle, the public may have revised its beliefs in line with (and perhaps due in part to) economic research on these

workforce and the subsequent decline in their elasticity: see Heim (2007) and Blau and Kahn (2007).

[17]For the Chetty (2012) mean, we use his calculation of the mean for his Panels A and B of Table 1, which includes hours and taxable income elasticity estimates. We do not adjust for the optimization frictions that are the focus of his important paper, so as to maintain consistency with how the other surveys' evidence is presented. He reports a value of 0.33 when that adjustment is made, bringing him closer to Keane's (2011) estimate.

[18]Feldstein (1995) also introduced to the literature a shift in emphasis from labor supply elasticities to the elasticity of table income. To the extent that estimates of the ETI are higher than conventional estimates of labor supply elasticity, this shift would be consistent with an increase in the perceived distortionary costs of taxation.

elasticities. However, there is no evidence of an increase in the perceived ETI from 1990 to 2010 in Figure 11, leaning against this idea. Indeed, that lack of evidence is consistent with some evidence suggesting that the long-run elasticity among high-earners—as opposed to the short-run elasticity, which may include temporal shifting of incomes—is not much higher than for the average worker (Goolsbee, 2000).

Finally, we consider a third source of data on perceived elasticities, in this case focusing on data that likely informs policymakers' beliefs. Policy in the United States is heavily influenced by the analyses and data provided by non-partisan (and partisan) groups in Washington, D.C. Foremost among these is the Congressional Budget Office, the government's official provider of budgetary and economic analyses for debates over legislation. In its economic analyses, the CBO uses values for the elasticities at issue in this paper, so evidence on the CBO's chosen values over time provides one way to gauge the perceived costs of distortionary taxation that informed policymaking.

We obtained several documents produced by the CBO that summarize its own views of the literature on the responses of economic activity to taxation.[19] The earliest relevant document was "An Analysis of the Roth-Kemp Tax Cut Proposal" from 1978. While no explicit elasticities are mentioned, the text contains the following statement: "For the labor force as a whole, a 10 percent increase in the disposable wage may lead to a 1 to 3 percent increase in hours worked." We infer from this that the CBO believed in an elasticity of 0.1 to 0.3 (it is unclear whether this was an uncompensated or compensated thought experiment).[20] The CBO's analysis of TRA-86 references Hausman (1985) as providing evidence of secondary-earner responsiveness to tax changes but downplays the response of labor supply to the reform overall. The CBO's 1994 analysis of President Bill Clinton's health insurance proposal states: "...the supply of labor is relatively insensitive to changes in take-home wages. Recent empirical studies suggest that the total hours supplied by U.S. workers would decline on 0.1 percent to 0.2 percent for each 1 percent reduction in their take-home wage," i.e., an elasticity of 0.1 to 0.2. In the 1996 review paper cited above, the CBO stated the range 0.2 to 0.4 from its reading of the literature. In 2012, the CBO reported that it had revised down that range to 0.1 to 0.3. Therefore, the CBO appears to have used roughly the same range for the elasticity controlling the distortionary costs of income taxation since 1978, and not surprisingly that range has been consistent with the academic consensus apparent in Figure 12.

#### 3.1.2 Beliefs in complementarities across skill levels

"Trickle-down economics" was the pejorative term applied in the 1980s to the idea that stimulating economic activity by high-earners would benefit low-earners, as well. Stiglitz (1982) is the best-known formalization of the idea in the optimal tax literature, and Rothschild and Scheuer (2013) recently expanded on his work. The basic idea of these analyses is that workers of different skill levels are complementary in production, so that an increase in effort by high-wage earners will raise the marginal productivities, and thus wages, of low-wage earners. That general-equilibrium dynamic is absent from our analysis thus far, and it may provide an alternative explanation of our findings. That is, if Americans more strongly believed in the idea of trickle down economics after the 1980s, they would have voted for policies much as if their perceptions of the distortionary costs of taxation had risen.

We can again look to the GSS for some (limited) evidence on this question. The GSS asked respondents in 1987 and 1996 whether they strongly agreed (5), agreed, felt neither way, disagreed, or strongly disagreed (1)

[19] Thanks to Ed Harris of the CBO for his guidance on which documents would have the relevant information.

[20] The report also states "It is generally agreed that the labor supply of all adult males is largely unaffected by changes in marginal tax rates. In most studies, both the substitution and the income effects are very close to zero."

with the statement: "Allowing business to make good profits is the best way to improve everyone's standard of living." The mean response decreased a bit from 1987 to 1996, from 2.73 to 2.66. Though of course only suggestive, these results are not consistent with a sharp increase in the belief in complementarities of the sort at work in Stiglitz (1982) over this period.

## 3.2 Explanation 2: Decreased political support for redistribution of high incomes

A second explanation for our findings could be that political support for redistribution decreased substantially over this period in the United States. We explore evidence related to two versions of this explanation: namely that support for redistribution among Americans truly declined, perhaps due to increased preference heterogeneity in the population, or that aspects of the political economy caused implicit policy preferences to falsely indicate such a decline, for example because of a shift of political influence toward higher-income Americans.

### 3.2.1 True preferences for redistribution

The General Social Survey (GSS) has asked U.S. respondents in either every or every-other year since 1978 for their opinions on a 7-point scale, ranging from "The government ought to reduce income differences" (1) to "Government should not concern itself with reducing income differences" (7). This question has been widely used in the literature on preferences for redistribution, such as Singhal (2013). Note that this question is well-suited for our purposes due to the phrasing of the extreme choices, emphasizing the normative judgment of the respondent apart from his or her beliefs about how distortionary it would be to reduce income differences.

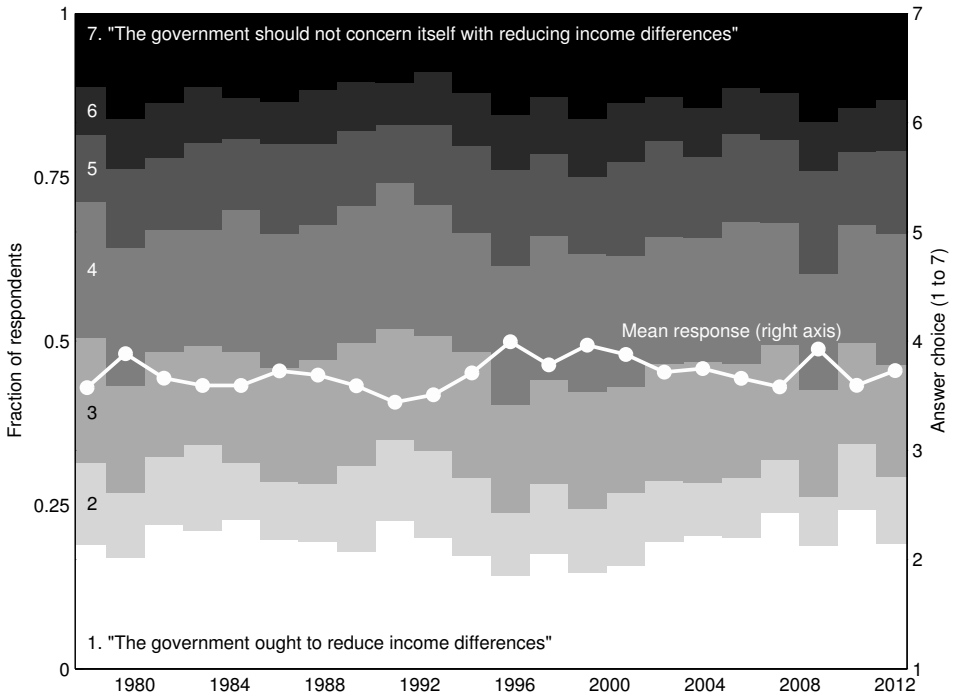



**Figure 13:** Responses to the General Social Survey on whether the government ought to reduce income differences. The solid line shows the mean answer in each year, as measured by the right axis.

Figure 13 presents the responses to this question following the format of Figure 11. It shows no trend in these attitudes toward redistribution. Over the last 35 years, Americans appear no more or less interested in having the government take responsibility for reducing income disparities.

Unfortunately, the proper interpretation of these data is unclear. One interpretation, which supports this explanation for our findings, views the respondents as answering with respect to contemporaneous policy. As long as policy accurately reflects their preferences, stability in their answers to this question would be consistent with decreasing support for redistribution over this period. An alternative interpretation views the respondents as answering with respect to their absolute preference for redistribution, divorced from current policy entirely. In that case, the stable trend in responses would indicate little change in true MSWWs over time.

**The extent of preference heterogeneity** One aspect of true preferences for redistribution involves the extent of heterogeneous preferences. Fleurbaey and Maniquet (2006) have developed in detail a distinction between dimensions of heterogeneity across individuals that likely merit redistribution, such as innate ability, and that likely do not, such as preferences for consumption relative to leisure. Lockwood and Weinzierl (2014), show that if such preferences vary across the population, optimal taxes will generally be less redistributive than in the conventional Mirrlees model. That paper also presents suggestive cross-sectional evidence that countries with greater preference heterogeneity (as reported in survey evidence) have less progressive tax codes. Increasing preference heterogeneity could therefore explain the trend toward less redistributive tax policy in the United States over this period.

Figure 14 plots survey responses pertaining to the relative preference for labor and leisure over time. We are most interested in changes to preference *heterogeneity*, and therefore the figure also plots the standard

deviation of responses over time.

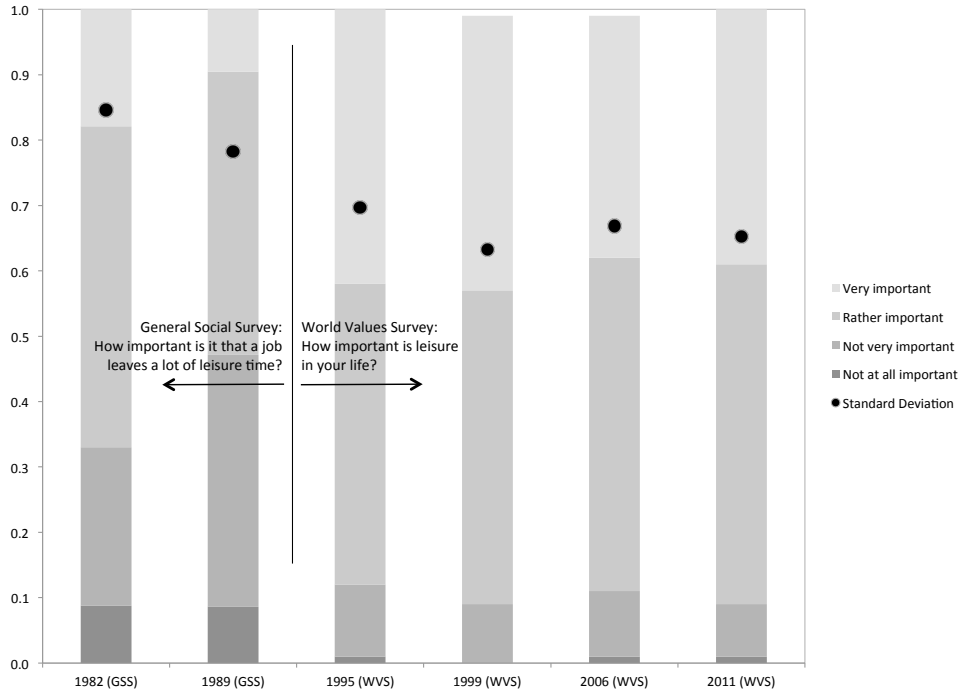



**Figure 14:** Responses to questions on the General Social Survey and the World Values Survey about preferences for leisure relative to consumption. Black dots plot the standard deviation of responses over time.

The evidence is hardly conclusive, and involves splicing together data from multiple surveys, but there appears to be little support for the hypothesis that preference heterogeneity rose substantially over time. On the contrary, the standard deviation of responses appears to decrease slightly over time.

### 3.2.2 Falsely indicated preferences due to political misrepresentation

As noted in Section 1.1, the inferred MSWWs may not reflect accurately the true redistributive preferences of voters if features of the political process have driven policy changes. A growing body of research in political science has explored the responsiveness of policy makers to public opinion. Of particular interest, the important work of Gilens (2012) and Gilens and Page (2014) presents evidence that policy is substantially more responsive to opinions of high-earners than to median-income constituents.[21]

To further explore this hypothesis, we present survey evidence on the perceived fairness of the political process over time. Figure 15 displays trends over time in the share of respondents agreeing with two relevant statements in the General Social Survey: "Public officials are interested in the problems of the average man" and "The average citizen has considerable influence on politics".

[21] See Gilens (2012) for a discussion of the literature on evolving political responsiveness more broadly.

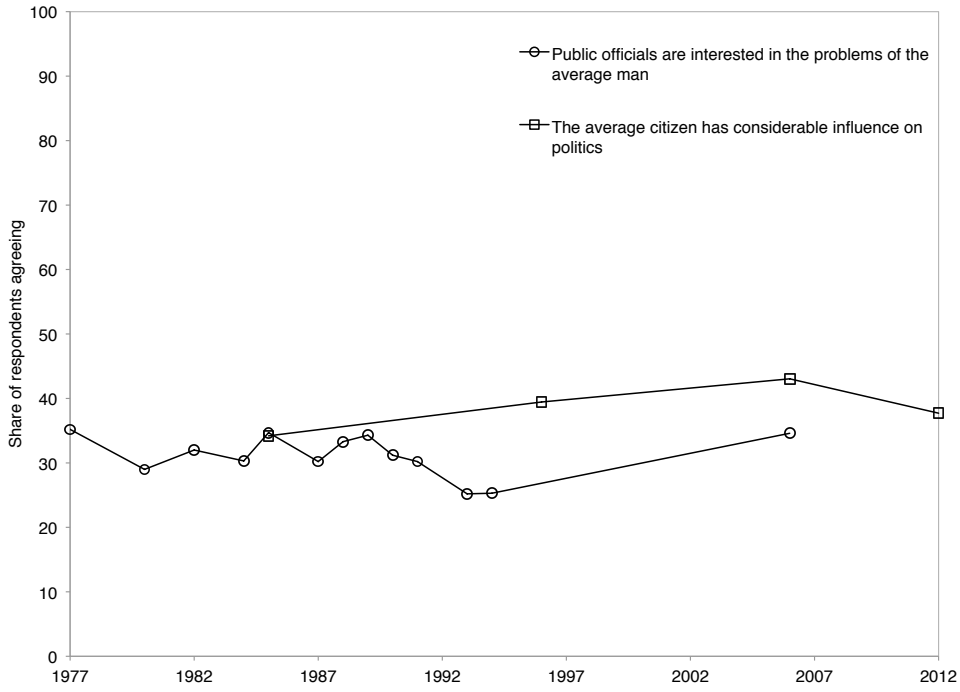



**Figure 15:** Responses to General Social Survey questions about political representation.

There is little evidence of a prolonged negative trend, as might be expected if the policy evolution toward less redistribution were at odds with true preferences, for example if policy were driven by growing political influence of high earners. Interestingly, the *level* of agreement among respondents indicates only a minority believe the political process reflects the interests of average citizens.

These features appear broadly consistent with findings in the political science literature. Gilens (2012) presents evidence that policy has become more responsive to public opinion generally and, more important for our purposes, that the influence of high earners does not appear to have risen relative to others, particularly during the 1980s when redistributive policy changed the most.

Other evidence on the possibility of evolving political representation comes from Pew survey results.[22] In 1997 and 2010, respondents were asked which of a number of groups—including "poor people", "the middle class", and "business leaders"—get "too much", the "right amount", or "too little attention" from government. Over that period, the share reporting that business leaders get too much attention actually fell from 50% to 45%, while the share believing the poor get too much attention rose from 10% to 17%. A stable share reported that the middle class gets too much attention; yet the share reporting that the middle gets too little attention rose from 54% to 66%. Although these results do not cover the 1980s—the period of the sharpest change in policy—they do not provide strong evidence of a shift in perceived political representation in favor of high earners over this period.

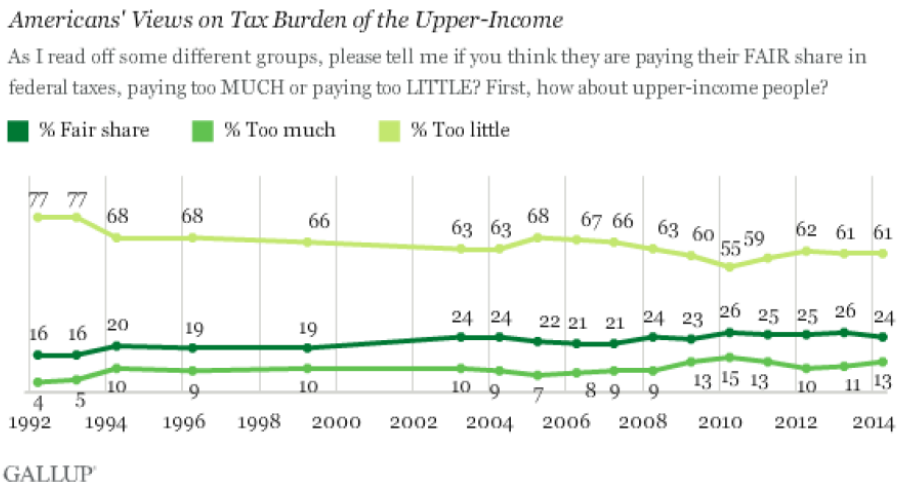
[22]See http://www.people-press.org/files/legacy-pdf/606.pdf.

### 3.3 Explanation 3: Stable but unconventionally low support for redistribution and limited increases in perceived ETIs

Finally, as discussed in Section 3.1, a higher top MSWW would allow for a much smaller (and perhaps plausible) change in perceived elasticities to rationalize observed policy evolution, e.g., from 0.17 in 1979 to 0.45 in 2010 under a top MSWW of 0.6 (see Figure 6). Moreover, if we assume that consensus estimates of the elasticity hold, substantially positive top MSWWs are unavoidable.

This explanation raises the question of why top MSWWs might be so high. The key question for our purposes is whether such high top MSWWs reflect Americans' true social preferences or whether the political system is catering to the rich–who likely support large top MSWWs–and failing to reflect the opinions of most Americans.

To make progress on this question, we examine existing data and generate some of our own. We start with survey evidence from Gallup, shown in the following figure:




In this figure, the share of respondents who say that "upper-income people are paying too little" is a substantial majority throughout the period. These Gallup data therefore strongly support the interpretation that the political system is failing to accurately reflect most Americans' (lower) preferred top MSWWs. Support for this interpretation can also be found in Gilens (2012) and–less directly–in Figure 15 above, where we saw that only a (constant) minority of respondents to the GSS believe the political system is responsive to the concerns of the average person.

One concern with any survey evidence, such as Gallup's, is its robustness to changing the framing or information that is salient to the respondents. In particular, it is unclear in this case what a respondent to Gallup's survey knows about the current tax burden of "upper-income people." As research using survey-based evidence on policy preferences has developed, increasing attention has been paid to such questions. For example, Kuziemko et al (2014) is an important recent study exploring how responsive reported preferences for redistribution (through income or estate taxes) are to the presentation of information.

In this paper, we try to address this uncertainty over the evidence on redistributive preferences by generating novel survey evidence in which we vary the information made salient to respondents. In addition to helping sort out the explanation for high top MSWWs, we hope that this evidence will be more broadly useful in the effort to understand Americans' redistributive preferences.

The details are as follows. Using Amazon MTurk, we randomize 200 respondents across four information treatments. All treatment groups face the same welcome and permissions screens at the start of the survey. Then, the control group #1 is asked the same questions as in the Gallup survey (except that we use the term "households" instead of "people" to better match the data we present the other groups). Group #2 is asked those same questions, but after being shown the distribution of pre-tax and -transfer income across household quintiles from the CBO. Group #3 is shown the distribution of federal income tax payments across these same quintiles, also from the CBO, while group #4 is shown both distributions before being asked the Gallup questions. These informational figures are as follows:

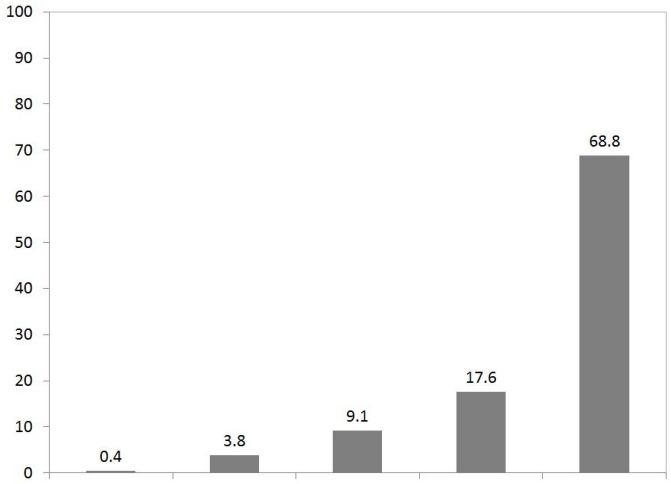



**Figure 16:** Distribution of market income from CBO (2013) by household income quintile.

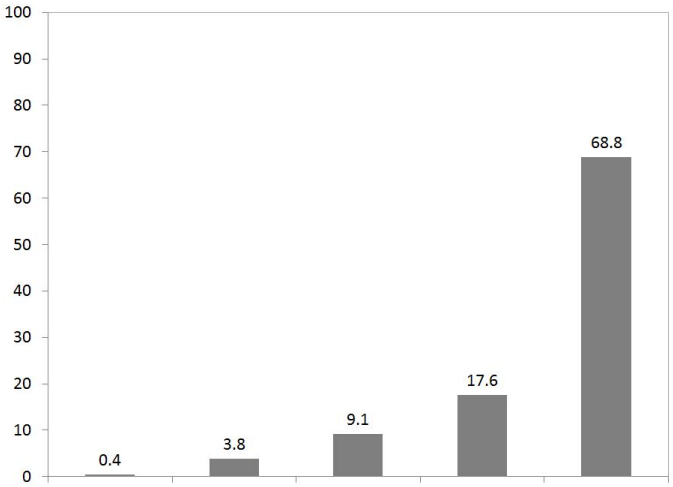



**Figure 17:** Distribution of federal tax payments from CBO (2013) by household income quintile.

The following table shows the distribution of responses to the Gallup question on upper-income people (households, in our survey). The unadjusted data are the raw results of the survey. The adjusted data reflect the fact that the sample from MTurk has a substantially greater share of respondents self-identifying as "Liberal" or "Left-leaning" than holds for Gallup's representative sample. We take Gallup's results on

the share of the population that self-identifies as Liberal, Moderate, or Conservative–a question we ask of our respondents as well–and rescale our results to match the same population shares. This gives us a result for the "No information" treatment that is much more similar to Gallup's than our unadjusted results.

Table 4: Results by treatment, raw and adjusted for political composition

**Unadjusted data**

| | Upper-income households pay... | | | |
|---|---|---|---|---|
| Treatment | Too little | Fair share | Too much | Obs |
| No information | 80% | 16% | 4% | 51 |
| Income shares | 70% | 22% | 7% | 54 |
| Tax shares | 48% | 33% | 19% | 39 |
| Both income and tax shares | 54% | 33% | 13% | 52 |

**Adjusted data (scaled by Gallup shares of political affiliation)**

| | Upper-income households pay... | | | |
|---|---|---|---|---|
| Treatment | Too little | Fair share | Too much | Obs |
| No information | 71% | 22% | 7% | 51 |
| Income shares | 59% | 29% | 12% | 54 |
| Tax shares | 34% | 40% | 26% | 39 |
| Both income and tax shares | 48% | 28% | 24% | 52 |

The table shows substantial sensitivity of respondents' preferences to information. Perhaps surprisingly, providing information on the concentration of income reduces the share of respondents believing that upper-income households pay too little. Most important for our purposes is the loss of the clear majority who say that upper-income people pay too little in taxes when information on tax shares is provided, even if that information is combined with information on income shares.

The sensitivity of preferences apparent in Table 4 is, in a sense, disappointing. If preferences had been unchanged across treatments, we would have strong evidence that most Americans robustly wish that taxes were more progressive and, in terms of what is at stake in this paper, that the large top MSWWs we inferred from policy are likely due to a political process that fails to reflect those wishes. Instead, a fundamental uncertainty remains: is policy not more progressive because most Americans are satisfied with policy as it is, or because they do not have the political influence to change it?

# 4 Conclusion

This paper infers the normative and positive judgments underlying recent U.S. tax policy from 1979 to 2010. We identify a puzzling divergence: conventional assumptions (or estimates) of the central parameters of modern optimal tax theory are inconsistent with these inferred judgments. We consider three possible explanations for this divergence: the perceived costs of redistribution grew unconventionally large; the perceived benefits of redistribution shrank unconventionally small; or the perceived costs grew somewhat while the perceived benefits were persistently and unconventionally small. Finally, we present some limited evidence on these possible explanations, including some new results on survey-based estimates of social preferences for redistribution and their sensitivity to the information provided to survey respondents.

Whatever turns out to be the true explanation for our findings, we believe that backing out the combinations of normative and positive judgments implicit in policy provides a promising path toward both a better understanding of policy priorities and more objective interpersonal comparisons for policy evaluation. With regard to the latter application, we demonstrate how these weights can be used to measure the welfare costs from unequal growth and recessions. We find that these costs vary greatly depending on which inferred results are used, a result that underlines the importance of better data and further analysis of the drivers of recent tax policy.

# 5 Appendix

**Derivation of expression (3).** Expression (2) can be written:

$$\int_y^\infty g(z)dF(z) = 1 - F(y) - \left(\frac{T'(y)}{1-T'(y)}\right)\varepsilon y f(y). \tag{8}$$

Using the Leibniz rule, differentiating both sides with respect to $y$ yields

$$-g(y)f(y) = \frac{d}{dy}\left[1 - F(y) - \left(\frac{T'(y)}{1-T'(y)}\right)\varepsilon y f(y)\right], \tag{9}$$

and dividing by $f(y)$ provides (3).

**Lorenz curves for calibrated income distributions.** Figure 18 plots the Lorenz curves for our calibrated distributions of market incomes across four decades. This progression demonstrates the well-documented concentration of income among high earners.

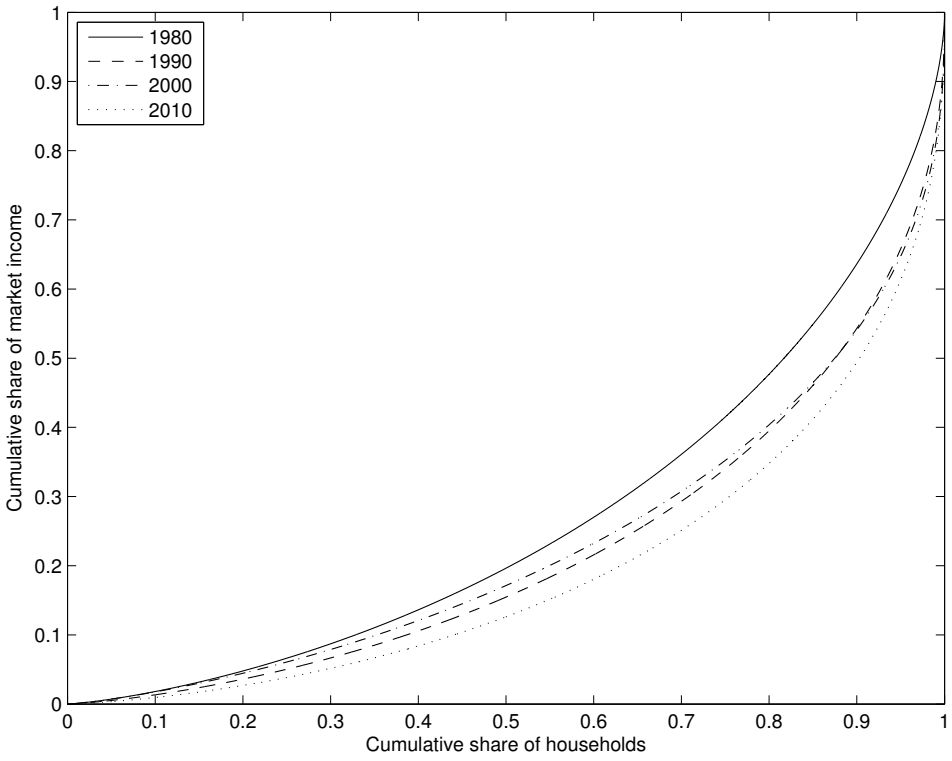



**Figure 18:** Lorenz curves for the calibrated distribution of market income in 1980, 1990, 2000, and 2010.

**Calculating MSWWs via numerical differentiation.** Figure 19 displays our calculated S weights and their smoothed version. This computation is sensitive to the assumed elasticity of taxable income—here we assume a value of 0.3. A technical complication is that discontinuities in the marginal tax rate schedule carry through to $\frac{\int_y^\infty g(z)dF(z)}{1-F(y)}$, generating points at which the schedule is not differentiable and the MSWWs are not defined. Since this feature is likely an artifact of the desire for a simple tax code, rather than a feature of underlying social preferences, we use a Gaussian kernel smoothing regression to smooth the schedule of S weights, rendering it differentiable, as shown in Figure 19. We use a bandwidth of 10% of the mean income in the 7th CBO bucket (the 95th to 99th percentile), which preserves the shape of the distribution while fully smoothing the kinks.

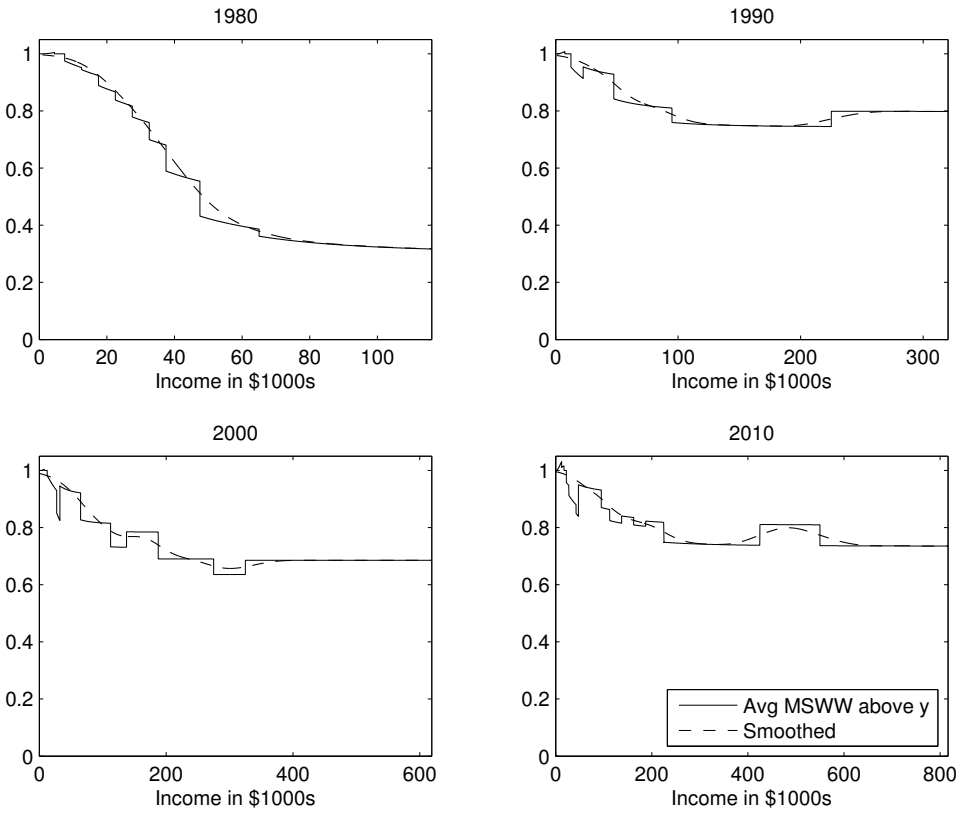



**Figure 19:** Average MSWWs above each income level, $\frac{\int_y^\infty g(z)dF(z)}{1-F(y)}$, in 1980, 1990, 2000, and 2010. Discontinuities are generated by kinks in the income tax schedule. Dashed lines are computed using Gaussian kernel smoothing regression with bandwidth equal to one tenth of the mean income in 95–99 percentile CBO bucket. Results are computed assuming an elasticity of taxable income of 0.3.

**Computing the evolution of MSWWs due solely to tax changes.** The evolution of MSWWs displayed in Figure 3 is driven both by rising inequality in market incomes and falling progressivity of the tax code. To isolate the effect of policy changes alone, Figure 20 shows what this evolution would have looked like if market income inequality had not risen over time—in this case the rising MSWW on top earners over this is driven entirely by reductions of top tax rates. This figure is constructed by scaling the market income distribution in 1980 by the change in average market incomes over time, so inequality (as measured by interquartile spreads or the Gini index) remains fixed. The schedule of MSWWs is plotted against the same vector of real disposable income as Figure 3 for comparability. Naturally MSWWs on high earners are somewhat lower when the effect of rising inequality is removed. Yet the qualitative resemblance between Figures 3 and 20 demonstrates the importance of tax reforms as the key driver of evolving implicit MSWWs.

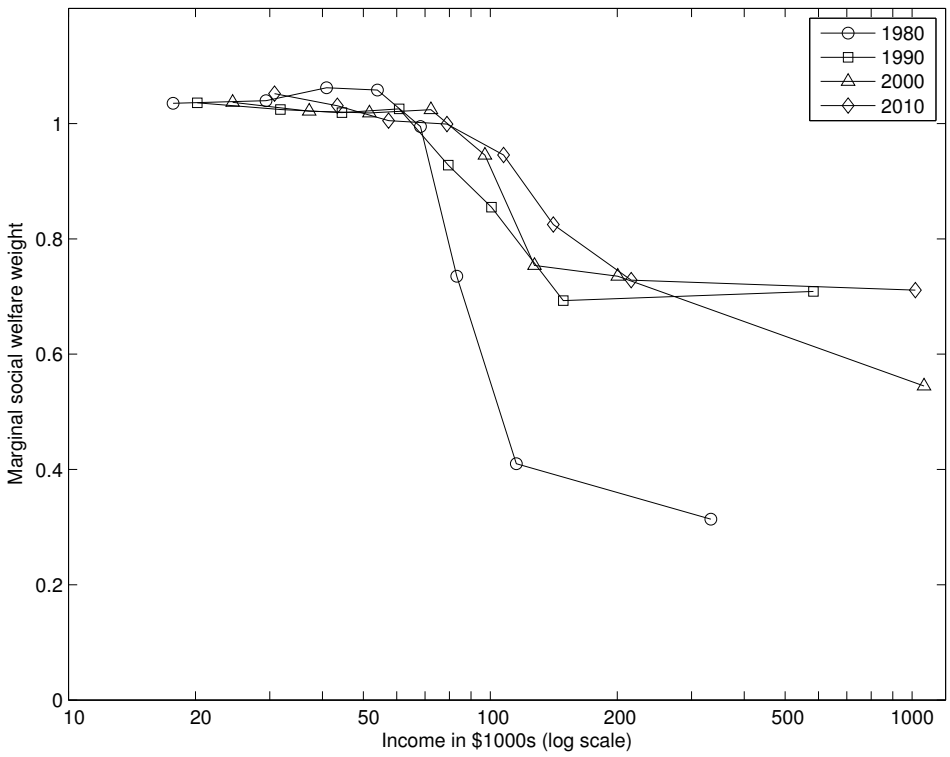



**Figure 20:** The evolution of MSWWs due to tax changes. This figure is identical to Figure 3, except the inequality of market income is held constant over time.

# References

**Bargain, Olivier, Dolls, Mathias, Immervoll, Herwig, Neumann, Dirk, Peichl, Andreas, Pestel, Nico and Siegloch, Sebastian**. (2013). 'Partisan Tax Policy and Income Inequality in the US, 1979-2007', Working paper.

**Bargain, Olivier, Dolls, Mathias, Neumann, Dirk, Peichl, Andreas and Siegloch, Sebastian**. (2011). 'Tax-benefit systems in Europe and the US: between equity and efficiency', Working paper.

**Bartels, Larry M**. (2008), *Unequal democracy: The political economy of the new gilded age*, Princeton University Press.

**Best, Michael and Kleven, Henrik Jacobsen**. (2014). 'Optimal income taxation with career effects of work effort', Working paper.

**Blau, Francine D and Kahn, Lawrence M**. (2007). 'Changes in the Labor Supply Behavior of Married Women: 1980–2000', *Journal of Labor Economics* 25(3).

**Blundell, Richard and MaCurdy, Thomas**. (1999), Labor supply: A review of alternative approaches, *in* 'Handbook of labor economics', Vol. 3, Elsevier, pp. 1559–1695.

**Bourguignon, François and Spadaro, Amedeo**. (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.

**Colombi, Roberto**. (1990), A new model of income distribution: The Pareto-lognormal distribution, *in* 'Income and Wealth Distribution, Inequality and Poverty', Springer, pp. 18–32.

**Congressional Budget Office**. (1978). 'An Analysis of the Roth-Kemp Tax Cut Proposal'.

**Congressional Budget Office**. (1987). 'Economic Impacts of the Tax Reform Act of 1986: Short-Run and Long-Run Perspectives'.

**Congressional Budget Office**. (1994). 'An Analysis of the Administration's Health Proposal'.

**Congressional Budget Office**. (1996). 'CBO Memorandum: Labor Supply and Taxes'.

**Congressional Budget Office**. (2012). 'A Review of Recent Research on Labor Supply Elasticities'.

**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.

**Diamond, Peter and Saez, Emmanuel**. (2011). 'The case for a progressive tax: from basic research to policy recommendations', *The Journal of Economic Perspectives* 25(4), 165–190.

**Farhi, Emmanuel and Werning, Ivan**. (2007). 'Inequality and social discounting', *Journal of Political Economy* 115(3), 365–402.

**Feldstein, Martin**. (1995). 'The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act', *Journal of Political Economy* 103(3).

**Fleurbaey, Marc and Maniquet, Francois**. (2006). 'Fair income tax', *The Review of Economic Studies* 73(1), 55–83.

**Gilens, Martin**. (2012), *Affluence and influence: Economic inequality and political power in America*, Princeton University Press.

**Gilens, Martin and Page, Benjamin I**. (2014). 'Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens', *Perspectives on Politics, Forthcoming* .

**Golosov, Mikhail, Kocherlakota, Narayana and Tsyvinski, Aleh**. (2003). 'Optimal indirect and capital taxation', *The Review of Economic Studies* 70(3), 569–587.

**Goolsbee, Austan**. (2000). 'What Happens When You Tax the Rich? Evidence from Executive Compensation', *The Journal of Political Economy* 108(2), 352–378.

**Hausman, J.A.** (1985), Taxes and labor supply, *in* 'Handbook of public economics', Elsevier, pp. 213–263.

**Heim, Bradley T**. (2007). 'The Incredible Shrinking Elasticities Married Female Labor Supply, 1978–2002', *Journal of Human Resources* 42(4), 881–918.

**Hendren, Nathaniel**. (2014). 'The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function', Working paper.

**Keane, Michael P**. (2011). 'Labor supply and taxes: A survey', *Journal of Economic Literature* pp. 961–1075.

**Kuziemko, Ilyana, Norton, Michael I, Saez, Emmanuel and Stantcheva, Stefanie**. (2013), How elastic are preferences for redistribution? Evidence from randomized survey experiments, Technical report, National Bureau of Economic Research.

**Lockwood, Benjamin B and Weinzierl, Matthew**. (2014). 'De Gustibus non est Taxandum: Heterogeneity in Preferences and Optimal Redistribution', Working paper.

**Lucas, Robert E**. (2003). 'Macroeconomic priorities', *American Economic Review* 93(1), 1–14.

**Mirrlees, James A**. (1971). 'An exploration in the theory of optimum income taxation', *The review of economic studies* pp. 175–208.

**Piketty, T and Saez, E**. (2007), Income and wage inequality in the United States, 1923-2002, *in* **Anthony B Atkinson and Thomas Piketty**., eds, 'Top Incomes Over the Twentieth Century: A Contrast Between Continental European and English-Speaking Countries', Oxford University Press, chapter 5, p. 141. Series updated by same authors.

**Reed, William J and Jorgensen, Murray**. (2004). 'The double Pareto-lognormal distribution–a new parametric model for size distributions', *Communications in Statistics-Theory and Methods* 33(8), 1733–1753.

**Rothschild, Casey and Scheuer, Florian**. (2013). 'Redistributive taxation in the roy model', *The Quarterly Journal of Economics* 128(2), 623–668.

**Saez, Emmanuel**. (2001). 'Using elasticities to derive optimal income tax rates', *The Review of Economic Studies* 68(1), 205–229.

**Saez, Emmanuel**. (2002). 'The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes', *Journal of Public Economics* 83(2), 217–230.

**Saez, Emmanuel, Slemrod, Joel and Giertz, Seth H**. (2012). 'The elasticity of taxable income with respect to marginal tax rates: A critical review', *Journal of Economic Literature* 50(1), 3–50.

**Saez, Emmanuel and Stantcheva, Stefanie**. (2013). 'Generalized social marginal welfare weights for optimal tax theory'.

**Singhal, Monica**. (2013). 'Quantifying Preferences for Redistribution', Working paper.

**Smith, Adam**. (1776), *An inquiry into the wealth of nations*, Strahan and Cadell, London.

**Spadaro, Amedeo, Mangiavacchi, Lucia and Piccoli, Luca**. (2012). 'Optimal taxation, social contract, and the four worlds of welfare capitalism', *DEA Working Papers* (51).

**Stantcheva, Stefanie**. (2014). 'Optimal Taxation and Human Capital Policies Over the Lifecycle', Working paper.

**Stiglitz, Joseph E**. (1982). 'Self-selection and Pareto efficient taxation', *Journal of Public Economics* 17(2), 213–240.

**Weinzierl, Matthew**. (2011). 'The surprising power of age-dependent taxes', *The Review of Economic Studies* 78(4), 1490–1518.

**Zoutman, Floris T, Jacobs, Bas and Jongen, Egbert LW**. (2013*a*). 'Optimal redistributive taxes and redistributive preferences in the Netherlands', Working paper.

**Zoutman, Floris T, Jacobs, Bas and Jongen, Egbert LW**. (2013*b*). 'Revealed Social Preferences of Dutch Political Parties', Working paper.