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

Citation

Lockwood, Benjamin B., and Matthew Weinzierl. "Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions." Harvard Business School Working Paper, No. 14-119, June 2014. [Revised October 2014.]

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:13479242>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

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

**Benjamin B. Lockwood
Matthew Weinzierl**

Working Paper

14-119

October 20, 2014

Copyright © 2014 by Benjamin B. Lockwood and Matthew Weinzierl

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

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

Benjamin B. Lockwood and Matthew Weinzierl*

October 20, 2014

Abstract

Calculating the welfare implications of changes to economic policy or shocks to the economy requires economists to decide on a normative criterion. One way to make that decision is to elicit the relevant moral criteria from real-world policy choices, converting a normative decision into a positive inference exercise as in, for example, the recent surge of so-called “inverse-optimum” research. We find that capitalizing on the potential of this approach is not as straightforward as we might hope. We perform the inverse-optimum inference on U.S. tax policy from 1979 through 2010 and identify two broad explanations for its evolution. These explanations, however, either undermine the reliability of the inference exercise’s conclusions or challenge conventional assumptions upon which economists routinely rely when performing welfare evaluations. We emphasize the need for better evidence on society’s positive and normative judgments in order to resolve the questions these findings raise.

Introduction

Economists are put in an awkward position when asked to calculate the welfare consequences of changes to economic policy or of shocks to the economy: we are asked to act as moral philosophers. Though we have largely converged on a standard approach to that task—i.e., by using a generalized form of utilitarianism—we have left room for a wide range of normative perspectives within that approach. For example, in optimal tax models we have remained agnostic about the values of the so-called marginal social welfare weights that determine the value of transferring resources across individuals (see Emmanuel Saez, 2001). Choosing a more specific normative perspective, for example choosing the values of the marginal social welfare weights, remains an uncertain and basically unwelcome task.

A potential solution to this awkwardness is to take our normative cues from society, eliciting the relevant moral criteria from real-world policy choices.¹ Progress along these lines has been facilitated by the recent “inverse-optimum” research that, following Francois Bourguignon and Amedeo Spadaro (2012), uses analytical results from optimal tax theory and assumptions on economic parameters to infer the marginal social welfare

*We are grateful to Alex Gelber, Andreas Peichl, and Danny Yagan for in-depth comments. Thanks also to Nathan Hendren, Louis Kaplow, Greg Mankiw, and Glen Weyl for insightful discussions. Lockwood: Harvard University, lockwood@fas.harvard.edu; Weinzierl: Harvard University and NBER, mweinzierl@hbs.edu.

¹See Saez and Stantcheva (2013), Weinzierl (2014a), and Weinzierl (2014b) for analyses that look for underlying principles that inform the inferred marginal social welfare weights.

weights (MSWWs) currently prevailing in a number of developed economies.² While that literature has largely refrained from using the inferred weights for welfare evaluation, the potential for doing so is clear. For example, related work by Nathan Hendren (2014) uses a similar inference exercise to calculate the implicit welfare costs of rising inequality in the modern United States, given mainstream estimates of the distortionary costs of taxation.³ In principle, these inference exercises would allow us convert the selection of a normative perspective into an empirical question, exempting economists from some difficult choices.

In this paper, we find that capitalizing on this potential solution is not as straightforward as we might hope. The problems start when we extend the previous literature by performing the inverse-optimum inference analysis intertemporally, using 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 from 1979 through 2010. These combinations fall roughly into two categories, each of which offers a different explanation for the evolution of U.S. policy. These explanations, however, raise their own questions to which the profession does not yet have good answers. Until they are answered, these questions either undermine the reliability of the inference exercise’s conclusions or challenge conventional assumptions upon which economists routinely rely when performing welfare evaluations. Answering these questions is thus, we argue, a prerequisite to fulfilling the promise of this empirical approach to normative questions and, more generally, an important task for economists interested in welfare analysis.

One explanation for the evolution of U.S. tax policy over the last several decades is as follows: if we assume that the perceived elasticity of taxable income (ETI) was stable at or above conventional levels, then the implicit schedule of MSWWs changed substantially during the 1980s. By the end of that decade, Americans appeared to place much less value on redistribution from high-income to low-income households than they did at the start. Smaller changes to the MSWW schedule continue to be implied after this major shift, and even larger changes are apparent if we extend our analysis back to the early twentieth century (where our data is more limited). This variation over time means that welfare calculations using inferred MSWWs may be, unfortunately, sensitive to the year from which one infers them. To illustrate this, we show that the calculated costs of rising inequality over this period fall by a factor of three if one uses the MSWWs implied by 2010 policy rather than those implied by 1980 policy. In what we believe is a novel application of this approach, we show similar results for the welfare costs of business cycles during this period. These findings naturally raise the question of whether these inferred MSWWs represent society’s “true” opinions or, instead, a biased outcome due to manipulation of the political system.⁴ In the latter case, the normative

²Francois Bourguignon 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 (2013; 2014a; 2014b) include the United States (as well as European countries) in their analyses. 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. In Bargain et al. (2014a) and Bargain et al. (2014b) 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. Floris Zoutman, Bas Jacobs, and Egbert L.W. Jongen (2013a; 2013b) 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. (2013a) they find that top MSWWs are negative unless the elasticity of labor supply is small, consistent with the Bourguignon and Spadaro results for France and ours for the United States. In Zoutman et al. (2013b) they show (among other things) that the tax proposals by Dutch political parties have the same feature.

³Specifically, 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).

⁴This bias includes the possibility that policymakers with tax policy preferences are elected based on non-tax issues, so that

relevance of such inference exercises is fundamentally called into question.

An alternative explanation is as follows: if we assume that MSWWs were stable over this period, then the perceived ETI more than doubled between 1980 and 2010. Assuming those stable MSWWs took conventional values (i.e., very small at high incomes), the perceived ETIs rose well above conventional empirical estimates. Such large and rising perceived ETIs have uncomfortable implications. If they reflect voters' awareness of real changes to the economy (i.e., that are not well-captured by conventional measures of the ETI), then optimal policy would be substantially less progressive than is conventionally thought and inference exercises ought to reconsider the value for the ETI they assume. If, instead, they reflect a divergence of voters' perceptions from the true values of the distortionary costs of taxation, the reliability of exercises that employ empirical ETI estimates to infer normative preferences is, again, called into question. These uncomfortable implications can be avoided, and the inferred ETIs can remain at or below conventional levels over the entire period, if stable MSWWs are allowed to take unconventionally large values at high incomes. Of course, large high-income MSWWs raise the challenging prospect that most economists' default approach to welfare calculations—unweighted utilitarianism—is inconsistent with revealed social preferences of Americans.⁵

Deciding between these explanations is difficult given the unobservable nature of the underlying data we require. As a first step, we present suggestive empirical evidence—both existing and new—on the perceived benefits of redistributive taxation (e.g., MSWWs) and its perceived distortionary costs (e.g., ETIs). On MSWWs, we show evidence that the level of support for government actions to reduce inequality has been stable throughout this period, though that stability may be consistent with either rising or stable high-income MSWWs. On whether the inferred MSWWs reflect society's true preferences, we find that beliefs about how well the political process represents average citizen's interests are stable over this period, and the extent of popular support for increasing the extent of redistribution above prevailing levels actually declines somewhat since the early 1990s. While the same evidence has shown that a substantial majority of Americans consistently want high earners to pay more in taxes, supporting the hypothesis that implicit top MSWWs are biased upward, we present (preliminary) new survey evidence that the level of that support is highly sensitive to the information provided to the respondent. On ETIs, some imperfect data suggests fluctuations in the perceived incentive effects of taxation over the latter part of this period that qualitatively fit with policy, but we lack data for the key transition in the 1980s. Academic research over this period consistently found that the average elasticity of labor supply, and later the average elasticity of taxable income (ETI), were consistently small, though some larger estimates of the latter elasticity among high earners were obtained during the 1990s.⁶ Finally, official government estimates of the ETI are small and stable throughout the period.

All together, we interpret these data as being inconsistent with at least one of the explanations for the evolution of U.S. policy: that the bias of the political system has caused inferred top-income MSWWs to rise farther above society's true preferences over time. Beyond that, existing data are simply insufficient to identify the true explanation. Given the remaining explanations' substantial implications for policy design

tax policy changes do not reflect the “median” voter's preferences. The inverse-optimum approach assumes, instead, that tax policy decisions are prominent and separable enough that this possibility can be ignored.

⁵This discussion is consistent with the point-in-time inverse-optimum calculations of Peter Diamond and Emmanuel Saez (2011), who note that conventional MSWWs imply that “The current [U.S. top marginal tax] rate, $\tau = 42.5$ percent, would be optimal only if the elasticity e were extremely high, equal to 0.9. [...] Alternatively, if the elasticity is $e = 0.25$, then $\tau = 42.5$ percent is optimal only if the marginal consumption of very high-income earners is highly valued, with $g = 0.72$.”

⁶See, for example, Feldstein (1995). Although these large estimates may have been driven by time shifting of earnings among high earners rather than a change in labor supply (Goolsbee, 2000), or by avoidance activities in which the social cost of avoidance is less than the private cost (Chetty, 2009), the high ETI estimates may have influenced popular perceptions of the distortionary costs of redistribution.

and evaluation, this indeterminacy means that a main lesson of this paper is that better data on the perceived costs and benefits of redistribution will be essential for progress along the revealed preference approach to normative policy analysis.

Fortunately, not all of our contributions in this paper are negative. We help develop the inverse-optimum 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. Finally, as far as we are aware, the application of this approach to the costs of recessions is new to this paper.

In addition, our paper’s novel survey evidence 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 the inversion analysis and our calibration of 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 and ETIs, using more limited evidence, back to 1920. We also distill from these results the two broad explanations for policy noted above. Section 2 uses the inference results of Section 1 to compute—and demonstrate the sensitivity of—the revealed preference costs of rising inequality and recessions. In Section 3 we discuss and introduce suggestive evidence pertaining to MSWWs and ETIs as a first step toward answering the difficult questions posed by our intertemporal inference results. 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.⁷ 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.

⁷In addition to the concerns with this inverse-optimum approach that are the focus of this paper, we may be concerned that actual tax policy reflects priorities not well-captured in this standard model. For instance, Smith (1776) famously recommended a “simple” tax system, a feature neglected by modern theory. One way to interpret our findings that policy differs from conventional recommendations is that we are measuring the impact of such non-standard factors. That possibility raises arguably even greater problems for the inverse-optimum approach, and for modern optimal tax theory in general, than a story in which parameters of the model take unexpected values.

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).⁸ 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 ε :

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

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.⁹

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.¹⁰ 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:

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

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

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

where the expression in brackets can be computed using numerical differentiation.¹¹ This provides an estimate of the MSWW $g(y)$ as a function of income.

⁸Income effects typically raise optimal marginal tax rates, especially at high incomes (e.g., see Saez, 2001, Table 1 and Figure 5). Including them in our analysis would therefore be likely to make even greater the departures from conventional assumptions required to explain the evolution of policy.

⁹A 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.

¹⁰Strictly 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.

¹¹The algebra behind this derivation is shown in the appendix.

In all of these expressions, the elasticity ε is written as a constant, but in fact it need not be. For example, if the elasticity varies with income, we might replace the constant ε with $\varepsilon(y)$, so that the results above apply when using the local elasticity of taxable income at each income y .

Concerns about model misspecification The optimal tax model on which this inversion exercise is based is clearly a stark simplification of reality, raising concerns that we may make incorrect 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.¹² 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.¹³ 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.¹⁴

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.¹⁵ 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

¹²Consistent with this, Bourguignon and Spadaro (2012) demonstrate that incorporating an extensive margin yields qualitatively similar results, particularly at the top, and Hendren (2014) ignores the extensive margin for those with incomes above EITC eligibility cutoffs.

¹³The 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.

¹⁴See Golosov, Kocherlakota and Tsyvinski (2003), Farhi and Werning (2007), Weinzierl (2011), and Best and Kleven (2014).

¹⁵We 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—see the appendix for further discussion of such taxes.

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.¹⁶ 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.¹⁷ We also use this definition of income 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 capital gains. Table 1 compares 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						
Year	99-100 share		95-99 share		90-95 share	
	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).

¹⁶For 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>)

¹⁷In 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 to factors discussed in Richard Burkhauser, Shuaizhang Feng, Stephen P. Jenkins, and Jeff Larrimore (2012), for example the focus on households in the CBO data rather than tax-paying units (also see related calculations 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.¹⁸

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 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

¹⁸Given the relatively low income at which the top marginal tax rate applies in the United States, we assign the same marginal tax rate to these households even if we somewhat understate their incomes.

1 displays averages reported by CBO and our best-fit calibration over the range of years for which we have data.

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.¹⁹ 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.

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.

Marginal tax rate schedule We use the NBER’s TAXSIM tool to obtain marginal tax rates on earned income 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

¹⁹Outside 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.

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. 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%.

Focusing on federal income tax rates ignores a number of important components of the tax system, including the phase-out of certain transfers and in-kind benefits such as SNAP (food stamps) and housing vouchers, and state and local taxes. This simplification allows us to make better comparisons across time, for which we lack fine-grained data on transfers and local tax policies. Nevertheless, we present alternative calibrations in the appendix to demonstrate that our main findings are likely robust to these considerations.

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). Diamond and Saez (2011) give a standard illustration of the latter feature: “For example, if the social value of utility is logarithmic in consumption, then social marginal welfare weights are inversely proportional to consumption. In that case, the social marginal utility at the \$1,364,000 average income of the top 1 percent in 2007 (Piketty and Saez, 2003) is only 3.9 percent of the social marginal utility of the median family, with income \$52,700 (U.S. Census Bureau, 2009).” Note that the same setup would imply the median family’s MSWW is only a fifth that of a family with a very low level of income (\$10,000), so that top MSWWs are close to zero in this familiar case.

Similarly, we will compare the implicit ETIs to the “conventional” levels of the elasticity of labor supply or, as available, the elasticity of taxable income estimated in the empirical labor literature. The Congressional Budget Office (2012*b*) for example, reviews that literature and finds “substitution elasticities [of labor supply] for the total population that range from 0.1 to 0.3.” Similarly, the survey article by Saez, Slemrod, and Giertz (2012) concludes “While there are no truly convincing estimates of the long-run elasticity [of taxable income], the best available estimates range from 0.12 to 0.40...[and]...there is no compelling evidence to date of *real* economic responses to tax rates...at the top of the income distribution.”

1.3.1 MSWW schedules and ETIs

First, we characterize the general features of the inferred MSWWs and ETIs. We begin by showing the MSWW schedules consistent with policy over time assuming the ETI equals 0.3, a conventional value. Figure 3 shows the averages of these MSWWs within each of the eight CBO income groups for 1980, 1990, 2000, and 2010.²⁰

²⁰To 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 $S(y) =$

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.

Figure 3 demonstrates that implicit average MSWWs are generally positive and decreasing with income, as conventionally assumed.²¹ Figure 3 also shows that there is substantial variation over time in the shape of the implicit average MSWW schedules at high incomes. Corresponding to the Tax Reform Act of 1986 (TRA86), there is a marked increase in the average MSWW for high earners in the late 1980s that persists over time, with smaller differences between the schedules for 1990, 2000, and 2010.²² These changes reflect both rising inequality in pre-tax income and changes to tax rates, though in the appendix we demonstrate

$\int_y^\infty g(z)dF(z)/(1 - F(y))$. 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). 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.

²¹One feature of these MSWW schedules that we do not explore in this paper, but which may be a fruitful topic for future examination, is the unconventional flatness of the MSWW schedules through median household income, which was approximately \$53,000 in 2010.

²²TRA86 included substantial base-broadening reforms, potentially prompting the concern that the effective marginal tax rate prior to this reform was lower than the 50 percent statutory rate, implying a smaller shift in preferences than we infer. Leaning against this concern are several considerations. Much of the broadening in TRA86 was done through reforms to taxation of corporate profits and capital income, not the earned income that we focus on. Those related to earned income were often inframarginal to high earners (e.g. changes to retirement savings deductibility) or, in the case of the expanded home mortgage interest deduction, actually narrowed the base for them. In fact, TRA86 was revenue-reducing on the personal income tax side, reflecting a strong desire by the policymakers to reduce the distortion and taxation of high earners (see Birnbaum and Murray (1987)). Consistent with the idea that statutory rates are the most appropriate measure of effective marginal tax rates on high earners in this period, Feldstein (1995) uses them in his analysis of the ETI around TRA86.

that the larger movements are largely explained by tax reforms alone.²³

Figure 4 shows these results for different ETI values. Average MSWWs 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 variation in MSWWs at the top of the income distribution is especially sensitive to the assumed ETI. At larger values of the ETI, the shift in high-income MSWWs corresponding to TRA86 is especially dramatic. If the ETI is 0.6, for example, the 1980 tax schedule is not Pareto efficient, placing negative average MSWWs (i.e., -0.4) on earners in the top percentile, while those MSWWs implicit in the 2010 tax policy are close to $+0.5$ —well *above* conventional assumptions. In contrast, if the ETI is 0.1, average MSWWs on high incomes vary much less, from 0.75 in the 1980 policy to 0.92 in the 2010 policy.

1.3.2 High-income MSWWs and ETIs

Figures 3 and 4 show that high-income MSWWs provide a simple way to distinguish MSWW schedules, so we use them to describe our results for the remainder of the paper. We denote by g_t^* the average MSWW on the top one percent of income-earners in year t , corresponding to the right-most point on the average MSWW schedules in Figures 3 and 4. Formally, let y_t^* 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.

²³Thanks to Diana Moreira for suggesting this decomposition.

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 Two explanations

We can now use Figure 5 as a simple summary graphic, allowing us to distill the results of our inference exercise into two broad options for explaining recent U.S. tax policy.

Stable elasticities, falling redistributive preferences The first possible explanation is that the perceived distortionary costs of taxation were stable at or above conventional levels over this period, while political support for redistribution among Americans decreased sharply after the 1980s. This explanation can be seen by looking vertically up and down Figure 5 at a given ETI, and in particular at ETIs at or above 0.3. Doing so reveals the variation in high-income MSWWs that is required to explain policy over this period at each stable ETI value.

Figure 6 shows the evolution of high-income MSWWs g_t^* under this explanation for each year of the 1979-2010 period and three ETI values at or above the conventional range of empirical estimates.

Figure 6: Implied social welfare weight on the top 1% for a range of elasticities of taxable income.

While the inferred values of g_t^* have been relatively stable over the last 20 years, they increase substantially in the late 1980s at the passage of TRA86 and then decrease sharply with the addition of the top 39.6 percent rate in the 1993 tax reform. Though such jumps likely reflect, at least under this explanation, a more gradual evolution of underlying social preferences, the figure strongly suggests that the perceived benefit of redistribution sharply decreased from the late 1970s to the early 1990s. This figure also shows a smaller rise in the implicit MSWWs from the 2001-2003 tax changes.

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 (on earned income), the same set of ETI values from Figure 6, and the simplified formula for the top marginal tax rate from Saez (2001) to back out the implicit g_t^* over this nearly 100-year period. Figure 7 shows the results.

Figure 7: Implied social welfare weight on the top 1% for a range of elasticities of taxable income, for the full history of the US income tax. This figure is constructed using the highest marginal tax rate on wages and other earned income, as reported at www.ctj.org/pdf/regcg.pdf.

As with the now-familiar figures from Thomas Piketty and Emmanuel Saez (2003) showing the U-shaped evolution of income inequality over this time period, Figure 7 demonstrates that the recent implicit values of g_t^* are higher than any since the early 1930s.

These figures raise the concern that any inverse-optimum exercise assuming a conventional (or higher) value for the ETI will give results on MSWW schedules that are highly contingent on the year from which it makes the inference. This concern remains even if we limit ourselves to variation in the post-1964 era, as we will argue below may be a reasonable restriction. An obvious example, given Figure 6, is that analyses assuming an ETI of 0.4 would have concluded that Americans' normative preferences in 1980 matched the conventional assumption of near-zero top MSWWs but were dramatically less redistributive just ten years later. This uncertainty over MSWWs translates directly into uncertainty over the welfare evaluations of policy changes or changes to the economy for which such inverse-optimum exercises are potentially useful.

Stable redistributive preferences, rising elasticities The second possible explanation is that American's preferences for redistribution were stable over this period, while the perceived elasticity of taxable income rose sharply, perhaps especially for high earners. This explanation can be seen by looking horizontally across Figure 5 at a given value of g^* . Doing so reveals the variation in the perceived ETI of high earners required to explain U.S. policy over this period when high-income MSWWs take stable values.

Figure 8 shows the evolution of required ETIs under this explanation for each year of the 1979–2010 period and four g^* values (note that the smallest of these values is the “conventional” assumption). Like the inferred values of g_t^* from the previous explanation, the implicit perceived ETIs have been relatively stable after jumps corresponding to the 1986 and 1993 tax changes.

Figure 8: 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, 0.6, or 0.9.

Using the same approach as with the previous explanation, we can extend this analysis back over the last century. Figure 9 shows the results.

Figure 9: Implied elasticities of taxable income for various welfare weights on the top 1% for the history of US income tax. This figure is constructed using the highest marginal tax rate on wages and other earned income, as reported at www.ctj.org/pdf/regcg.pdf.

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. In fact, we will argue below that a reasonable restriction may be to limit analysis to the post-1964 era, where the distortions of the Great Depression and World War II exerted less influence on policy.

As an examination of these figures reveals, within this explanation are two distinct scenarios, each of which has challenging implications. The first scenario applies for conventional values of the high-income MSWWs, where the required ETIs rise well above conventional ranges. For example, the benchmark $g^* = 0.01$ value would imply an ETI that rose from approximately 0.4 in 1980 to 1.5 in 1990 before settling in at 1.1 in 2010. The second scenario applies when high-income MSWWs are stable at levels above approximately 0.6, in which case rising ETIs remain within the conventional range of empirical estimates as suggested by Saez, Slemrod, and Giertz (2012): i.e., between 0.1 and 0.4.

The first scenario—with conventional values for high-income MSWWs—implies that voters’ perceptions of the distortionary costs of taxation far exceed professional estimates. Do those perceptions reflect the collective wisdom of the population, perhaps taking into account dimensions of the response to taxation that lie outside the scope of traditional estimates? If so, economists’ optimal policy recommendations and, perhaps more important their evaluations of policy proposals, ought to be based on far greater costs of redistribution than they typically are. Another possibility is that public perceptions of the ETI are misperceptions—perhaps due to the influence of propaganda from those opposed to progressive policies. In that case, we will infer incorrect values for the MSWWs from inverse-optimum exercises that use professionals’ estimates of the ETI, since those estimates are not what determines voters’ policy preferences.

The second scenario—with large high-income MSWWs—has the unique advantage of being the only explanation in which both ETIs and MSWWs are stable over time. Of course, the price of that stability is that the extreme values for high-income MSWWs would suggest that most economists’ default approach to welfare calculations—i.e., unweighted utilitarianism—is inconsistent with revealed social preferences of Americans.

1.3.4 1960s as a case study

As the longer time series figures above show, the two decades following World War II were characterized by extremely low perceived ETIs or highly negative implicit values for g^* . We might reasonably suspect that war-time tax policy, which persisted largely unchanged until the early 1960s, was influenced by factors poorly captured by the standard optimal tax model, so that inverse-optimum exercises are inappropriate for our data prior to the reform in 1964. A similar logic may make us skeptical of applying the model to the years prior to World War Two, including as they did the upheavals of the Great Depression and World War I.

If the reforms of 1964 marked the start of an era of income taxation to which we are more comfortable applying the inverse-optimum approach, it may be useful to understand the motivation behind those reforms. In fact, further examination of the reforms in 1964 provides a useful illustration of the ambiguity at the heart of this paper. The most important features of the 1964 reforms, for the purposes of this paper, were its substantial reductions in high-income marginal tax rates, for example from a top rate of 91 percent to a top rate of 70 percent on incomes over \$200,000.

President John F. Kennedy gave an argument for the 1964 reforms that stressed the distortionary costs of high marginal tax rates:

“Our present tax system, developed as it was, in good part, during World War II to restrain

growth, exerts too heavy a drag on growth in peace time; that it siphons out of the private economy too large a share of personal and business purchasing power; that it reduces the financial incentives for personal effort, investment, and risk-taking.”

Kennedy’s arguments are consistent with the second of our explanations above—namely he expresses a perception that the ETI was too high for a top marginal tax rate of 91 percent to be optimal in peacetime, while during the war it had another—unconventional—justification.

At nearly the same time, Ronald Reagan, soon to be the Governor of California and then President of the United States, gave a speech supporting Barry Goldwater, the Republican nominee for President in 1964. In it, he made a very different argument for a flatter marginal tax rate structure:

“Have we the courage and the will to face up to the immorality and discrimination of the progressive tax, and demand a return to traditional proportionate taxation? Today in our country the tax collector’s share is 37 cents of every dollar earned. Freedom has never been so fragile, so close to slipping from our grasp.”

Reagan’s arguments are consistent with the first of our explanations above—namely that the welfare weights on those paying the top marginal tax rate were relatively too large for a 91 percent rate to be optimal.

These arguments—and their opposites—have thus been in use throughout the last half-century of income taxation in the United States. Sorting out which of them, or which combination of them, determined policy is difficult. But as we show in the next section, it is also essential for fulfilling the potential of a positive approach to the welfare evaluations of policy.

2 Applications: welfare costs of unequal growth and recessions

Ordinarily, welfare calculations of changes to policy or the economy are controversial because they are sensitive to the modeler’s assumptions about the social preferences for redistribution (i.e., the 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 (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 the year from which MSWWs are derived (if we adopt the first explanation above that perceived ETIs were constant over this period) or to the level of stable MSWWs that are assumed to hold (if we adopt the second explanation that perceived ETIs varied over this period). These findings thereby warn against reliance on any single such welfare calculation and make clear the imperative of better evidence on society’s preferences and beliefs about the fundamental parameters that enter into policy evaluation.

2.1 Costs of unequal growth

From 1979 to 2010, U.S. 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 10 provides a graphical view of this evolution, plotting average income after federal taxes and transfers for the eight quantiles reported by the CBO.

Figure 10: 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 a set of corresponding MSWWs from 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 $\Delta W_{m,n} | \{g^i\}_{i=1}^8$ is

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

Next we calculate the counterfactual change in welfare $\tilde{\Delta W}_{m,n}(\rho)$ that would result from a given equally-

distributed annual growth rate ρ from years m to n under a given vector of MSWWs:

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

By solving for the ρ 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. 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.

The difference between the actual aggregate growth rate and ρ , as a share of the former, is our measure of the costs of unequal growth. Table 2 shows these costs using MSWWs computed in 1980, 1990, 2000, and 2010 for four values of the ETI. It turns out that these costs are strongly dependent on the vector of MSWWs and, therefore, on the explanation one adopts for the evolution of U.S. tax policy since 1980. This exercise therefore illustrates the challenge to researchers interested in using this positive approach to welfare evaluation.

Cost of Rising Inequality, 1979–2010				
	ETI			
	0.1	0.3	0.4	0.6
1980 MSWWs	3.8%	13%	17.6%	29%
	(\$1420)	(\$4560)	(\$6300)	(\$10,200)
1990 MSWWs	1.2%	3.6%	4.9%	7.7%
	(\$430)	(\$1340)	(\$1810)	(\$2810)
2000 MSWWs	1.6%	4.9%	6.7%	11%
	(\$580)	(\$1810)	(\$2450)	(\$3840)
2010 MSWWs	1.4%	4.3%	5.8%	9.0%
	(\$510)	(\$1580)	(\$2140)	(\$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.

Consider the first explanation for the evolution of policy—i.e., that the perceived ETI is stable at or above conventional levels. Starting with an ETI of 0.3, a preferred value in the empirical literature, we see that the costs of unequal growth are only 4.3 percent of realized growth using 2010 MSWWs but rise to 13 percent under 1980 MSWWs. This difference is due to the significantly greater inequality aversion embedded in the earlier weights. Higher values for the stable perceived ETI make that difference even greater. These differences complicate any discussion of the importance—from the perspective of welfare—of addressing unequal growth.

Table 2 also allows us to see the challenges posed by the second explanation for the evolution of policy—i.e., that the MSWW schedules are stable. Under this explanation, the costs of unequal growth ought to be the same regardless of the year from which MSWWs are inferred. Start by looking at the costs using 1980 MSWWs calculated with the empirically-central ETI of 0.3, where g^* also equals approximately 0.3. In this

case, the cost of unequal growth is 13 percent. To obtain the same implicit MSWWs, and therefore cost, in 1990, the table shows that not even an ETI of 0.6 is sufficient (i.e., the cost is only 7.7 percent, not 13). The same result holds for 2000 and 2010, implying that exceptionally high ETIs are required to explain policy if the 1980 MSWWs given a standard ETI are assumed. Even greater variation in ETIs would be required for MSWWs with more conventional values for g^* , raising important questions about the implied gap between perceived ETIs and professional economists’ estimates of them. An alternative apparent in Table 2 is to start with the 1980 MSWWs calculated with an ETI of 0.1. In that case, g^* is nearly 0.8, and the cost of unequal growth is less than four percent. Similar values for that cost can then be found in all future decades for ETIs below 0.4, keeping the perceived ETIs within conventional ranges. Of course, to achieve this stability in ETIs along with stable MSWWs requires exceptionally high values for top MSWWs, as suggested by the very small costs of unequal growth that this scenario implies.

Before moving on to our second application of this approach, we discuss how our calculations of the cost of unequal growth compare to the important work of Hendren (2014), who estimates that adjusting for increased inequality since 1980 would offset 15–20% of growth. 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.²⁴ 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. Unlike under our complementary definition based on relative changes, under his definition equally distributed gains *reduce* inequality as measured by interquantile spreads or the Gini index. 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. Reassuringly, 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 suggests 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.

²⁴Hendren’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).

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, showing how we smooth income growth over the recessionary period, is shown in Figure 11,

Figure 11: 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 $\{g^i\}_{i=1}^8$ during, say, the 2001–2003 recession, as:

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

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 α that satisfies the following:

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

which has the same right-hand-side as expression (6). The value of α 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. We dub α the “equally distributed equivalent loss”.

In Table 3, we show these costs for four recessionary periods using MSWWs from 1980 and 2010. All of these calculations assume an ETI of 0.3, so they best illustrate the challenges resulting from the first explanation for the evolution of U.S. tax policy over this period.

Cost of Unequal Distribution of Recessions, 1979–2010					
Recession	Average change in disposable income	Equally-distributed equivalent (α)		Ratio of α to average change	
		1980 MSWWs	2010 MSWWs	1980 MSWWs	2010 MSWWs
1981–1984	–4.3%	–2.1%	–2.1%	48%	48%
1990–1992	–2.3%	–1.4%	–2.0%	63%	87%
2001–2003	–2.6%	–1.8%	–2.3%	69%	89%
2007–2010	–12.3%	–2.8%	–5.1%	23%	41%

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. The second and third columns show the equally-distributed income loss, α , which would generate the same welfare cost, using 1980 and 2010 MSWWs. The final two columns show the ratio of the equally distributed equivalent to the average loss —this is the amount by which the recessionary shock must be adjusted to account for its unequal impact. 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 smaller than the actual (unequal) recessionary shock in all four of these recent recessions. This is driven by the fact that reductions in disposable (not necessarily market) income during these recessions are concentrated on high earners. For example, as can be seen from Figure 10, 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 echoes our results on the costs of unequal growth, i.e., the sensitivity of the results to the year in which MSWWs are computed. The Great Recession, for example, exhibits an equally distributed equivalent loss that is over 80% larger under 2010 MSWWs as under 1980 MSWWs. This feature is pronounced when costs are concentrated in portions of the income distribution over which different years’ MSWW schedules disagree. Thus the 1981–1984 recession, for example, has the same equally-distributed equivalent loss under 1980 and 2010 welfare weights, indicating that losses in this recession spread throughout the lower portion of the income distribution, where the MSWW schedules largely overlap.

The two applications in this section make clear the perils of relying on any single welfare estimate of changes to the economy or to economic policy unless we are confident that the inputs to that calculation—i.e., the welfare weights—are robust indicators of society’s true normative preferences. As we have argued above, acquiring that confidence for American normative preferences will require deciding between two competing explanations for the evolution of U.S. tax policy over the last several decades. In the next section, we offer some suggestive evidence as a first step toward making that decision.

3 Suggestive evidence on possible explanations for the evolution of policy

In this section we present some suggestive evidence on the key parameters, the MSWWs and the ETIs, at issue in the two 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. We believe both that obtaining such data is feasible and that doing so will require careful survey design. Our new survey evidence on attitudes toward current levels of redistribution illustrates, we hope, both of those beliefs.

3.1 Evidence on perceived ETIs

We begin by focusing on whether changes to the perceived distortionary costs of redistributive taxation offer a plausible explanation for the path of U.S. tax policy over the last several decades.²⁵ Because the data are limited, we draw on three sources of evidence: popular opinion, academic research, and official government estimates.

Public opinion 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 12 shows the distribution of responses to this question, as well as the mean response, from 1990 through 2011.

²⁵In the Appendix, we discuss and find no evidence for a possibility not directly related to ETIs, namely that the perceived importance of complementarities across workers in the spirit of Stiglitz (1982) increased over this period.

Figure 12: 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 8 above, where the implied ETIs were shown to dip between 1990 and 1995 and then gradually rise. Of course, the key test of this interpretation would be whether it could match the more dramatic shift during the 1980s shown in Figure 12. Moreover, the data in Figure 9 cannot tell us whether the high *levels* of the ETIs shown in Figure 6 match the public’s perceptions.

Academic research 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 taxable income (or labor supply, as available) of prime-aged men over time.²⁶ We include estimates of the elasticities of both taxable income and labor supply because the latter was the main empirical target for researchers during the early years of our sample. While the ETI includes margins of adjustment in addition to labor supply, the most recent surveys of the literature suggest that the quantitative difference is very small between the portions of these elasticities representing lost economic activity, rather than temporary shifts or strategic reporting of economic activity (see the Saez et al. (2012) summary in section 1.3, and see Chetty (2009) for a discussion of the components of the ETI).

The CBO’s 1996 report provides a survey of surveys, giving the mean labor supply estimates from nine literature reviews spanning 1979 through 1993. Figure 13 plots the mean estimates reported by the survey papers in the CBO report and four other surveys, along with the point estimates from each of the papers

²⁶Excluding estimates for female elasticities means that we ignore the dramatic increased participation of women in the workforce and the subsequent decline in their elasticity: see Heim (2007) and Blau and Kahn (2007). Another survey of the literature, showing consistent results, is Bargain and Peichl (2013).

reviewed in the very recent Keane and Chetty papers.

Figure 13: Elasticity estimates in the economics literature over time, drawn from the following survey papers: CBO (1996), Blundell and MaCurdy (1999), Keane (2011), Saez et al. (2012), and Chetty (2012).

The evidence in Figure 13 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.²⁷ 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 13 relates to the estimated elasticities for top incomes, shown as “x” symbols in the figure, which arose relatively recently. Lawrence Lindsey (1987) provided an influential analysis of the effects of top marginal tax rate reductions in the 1981 tax reform, estimating large ETIs—above 1.5—around the time of the 1986 tax reform. Such high estimates were found for the 1986 reforms in the well-known analysis of Feldstein (1995). These large estimates of high-earner elasticities were followed by smaller estimates (e.g., between 0.4 and 0.6) in subsequent work that took into account intertemporal shifting of incomes among high earners (Goolsbee, 2000) and econometric concerns (Gruber and Saez, 2002).²⁸ In fact, Goolsbee (1999) argued that the empirical results from the 1990s (based on 1980s) tax reforms were aberrations, as similar analyses applied to other periods’ tax changes yielded much smaller

²⁷For 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.

²⁸Diamond and Saez (2011) comment on the latter of these as follows: “The paper Gruber and Saez (2002) is often cited for its substantial taxable income elasticity estimate ($e = 0.57$) at the top of the distribution. However, its authors also found a small elasticity ($e = 0.17$) for income before any deductions, even at the top of the distribution.”

elasticity estimates. Summarizing the evidence, CBO (2012*b*) writes: “There is little compelling evidence that high-income taxpayers have substantially higher elasticities with respect to their labor input than lower-income taxpayers. Higher estimates of the elasticity of broad income among high-income taxpayers appear to reflect their greater ability to time their income rather than greater changes in their labor supply.” Therefore, the observed evolution of policy would be consistent with conventional values for high-income MSWWs if the public revised its beliefs on the distortionary costs of redistributive taxation in line with the early, but not later, results in this literature.²⁹ That is, for $g^* = 0.01$ the ETI would have had to rise well above 0.6 to explain the evolution of policy.

Official estimates Finally, we consider a third source of data on perceived elasticities, in this case focusing on data that likely inform policymakers’ beliefs. The Congressional Budget Office is the U.S. 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 were salient to policymakers.

We obtained several documents produced by the CBO that summarize its own views of the literature on the responses of economic activity to taxation.³⁰ 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).³¹ 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*b*, 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 13.

3.2 Evidence on MSWWs

Now we turn to evidence on whether the evolution of preferences for redistribution among Americans explains the evolution of U.S. tax policy.³² Recall that U.S. policy is consistent with either rising high-income MSWWs (if the perceived ETI is at or above conventional levels) or a stable but unconventionally large high-income MSWWs (if the ETI is smaller). We use evidence on “true” social preferences for redistribution as well as on whether bias in the political system caused a divergence between those true social preferences and

²⁹A related possibility is that the public began to perceive the full ETI, rather than only a portion of it, as reflecting the response of economic activity to taxation.

³⁰Thanks to Ed Harris of the CBO for his guidance on which documents would have the relevant information.

³¹The 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.”

³²In the Appendix, we discuss and show evidence against the hypothesis that Americans perceived greater heterogeneity in preferences for consumption relative to leisure over this period.

the MSWWs implicit in policy, for example because of a shift of political influence toward higher-income Americans.

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 14: 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 14 presents the responses to this question following the format of Figure 12. 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 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. A second 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.³³

³³Contemporaneous research by Kuziemko and Washington (2014) has documented in detail the overall stability of redistributive preferences among Americans over this period, though they also highlight important variation in trends across demographic groups. One finding of particular relevance is that redistributive preferences have risen among the young but fallen among those over 65 years of age over this period.

One way to differentiate between these interpretations is that, if the latter scenario were true, the divergence between stable absolute preferences for redistribution and decreasingly redistributive policy should be reflected in greater dissatisfaction with policy over time. 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”.

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

The trend and the levels in Figure 15 tell different stories. 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. At the same time, the *level* of the agreement with this GSS question indicates that only a minority of respondents 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. With respect to levels, however, Gilens (2012) and Gilens and Page (2014) present evidence that policy is substantially more responsive to opinions of high-earners than to median-income constituents.³⁴

Pew survey results over the most recent decade of this period tell a broadly consistent story.³⁵ 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.

³⁴See Gilens (2012) for a discussion of the literature on evolving political responsiveness more broadly.

³⁵See <http://www.people-press.org/files/legacy-pdf/606.pdf>.

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 provide at most weak evidence of a shift in perceived political representation in favor of high earners and stronger evidence that the level of representativeness is perceived to be persistently low.

The GSS and Pew data would therefore be consistent with implicit high-income MSWWs being large because the political system is persistently biased toward the interests of the rich.

To further probe that possibility, we turn to survey evidence from Gallup in which respondents were asked whether “upper income people” pay “too much”, “too little,” or their “fair share” in taxes from 1992 through 2014. Unlike the GSS and Pew questions, this Gallup question is directly related to taxation. Moreover, it asks the respondent to compare the current state of tax policy with his or her ideal, a suitable question for gauging whether a biased political system is neglecting popular opinion. The results are shown in the following figure:

These Gallup data support the interpretation that the political system is failing to accurately reflect most Americans’ (lower) preferred top MSWWs. In this figure, the share of respondents who say that upper-income people are paying too little is a substantial majority throughout the period 1992 to 2014.³⁶ Note that, as with the previous evidence, these data do not suggest a worsening of the dissatisfaction with policy.

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 and Washington (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

³⁶Changes to this share in response to policy changes are as would be expected. For example, the increase in top marginal tax rates in 1993 appears to match precisely the timing of a drop in the share of respondents who believe the rich pay too little in the Gallup data. The reverse happens after the 2001 and 2003 tax cuts lowered top rates on high earners.

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 (left panel) and distribution of federal tax payments from CBO (2013) by household income quintile (right panel).

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.

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. This finding suggests that the mismatch between policy and voter preferences may not be as great as the raw Gallup data indicated.

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 to match 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
Gallup results	61%	24%	13%	1,026

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

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 extends the recent surge of so-called inverse-optimum research, which seeks to use data on prevailing policy to infer society’s normative preferences, across time to explain U.S. tax policy from 1979 to 2010. We characterize the set of positive and normative judgments that are implied by the joint evolution of tax policy and the income distribution over this period and, for reference, back to 1920.

We classify the results into two broad explanations for the evolution of policy, each of which has troubling implications. In particular, in some cases the explanation requires key parameters of the underlying model—marginal social welfare weights on different households and the elasticity of taxable income—to vary dramatically over time. This variation undermines the reliability of this approach and casts doubt on any single welfare estimate for a change to the economy or economic policy based on such inferences. We illustrate this problem by applying this approach to the costs of unequal growth and unequal recessions in the United States over this period, which we find vary over a wide range depending on which (plausible) assumptions are used. This variation can be avoided by relying on other explanations, but those require the

key model parameters to take highly unconventional values that would have substantial, perhaps implausible implications for policy analysis and recommendations.

We attempt to determine which of these explanations is best supported by the limited relevant data, drawing on public opinion surveys, academic research, official government estimates, and our own novel survey evidence. We find strong evidence against the hypothesis that the bias of the political system has caused inferred preferences for redistribution to fall below society's true preferences over this period. But we have much weaker, or more mixed, evidence on the remaining explanations, including the possibility that the bias of the political system has consistently led to implicit preferences for redistribution that are below society's true preferences. Better data bearing on these possibilities are therefore essential if we hope to use this revealed preference approach to avoid putting key normative decisions in the hands of the economists doing welfare analysis.

5 Appendix

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

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

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

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

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

Lorenz curves for calibrated income distributions. Figure 17 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 17: Lorenz curves for the calibrated distribution of market income in 1980, 1990, 2000, and 2010.

Calculating MSWWs via numerical differentiation. Figure 18 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 18. 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 18: 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.

Accounting for transfers and state and local taxes As mentioned in the text, our measure of marginal tax rates excludes two important components of the overall tax burden. First, it does not include phase-outs of transfer programs such as SNAP (food stamps) and housing vouchers. This exclusion is only a concern for our main findings if such transfers have changed substantially over time. Moreover, the effect of such phase-outs are likely strongest at low incomes, whereas the primary evolution of MSWWs in Figure 3 is concentrated among high earners. Nevertheless, we explore the possible extent of rising transfers (and thus, rising implicit marginal tax rates from phase-outs) by recreating Figure 3 with an alternative tax schedule for 2010. Specifically, for those with annual earnings below \$30,000, we replace the marginal tax rate as reported by TAXSIM with a constant rate of 40%, intended to approximate the effect of incorporating phase-outs according to the Congressional Budget Office’s 2012*a* report on marginal tax rate’s among low and middle income households. The evolution of MSWWs under this alternative specification is displayed in Figure 19. This evolution can be considered an “upper bound” on the effect of incorporating evolving transfers, as it effectively assumes transfers were absent entirely in 1980 and fully present in 2010. As expected, this

modification lowers MSWWs for low and middle incomes in 2010, but leaves the qualitative message from Figure 3 intact.

Figure 19: Average MSWWs, plotted against disposable income in 1980, 1990, 2000, and 2010. The 2010 marginal tax rates are set to 40% for those earning less than \$30,000 annually, to approximate the effect of implicit marginal tax rates due to transfer phase-outs during that year.

Another important component of the total tax burden ignored in our baseline analysis is state and local taxes. As pointed out in Hendren (2014), to the extent that these taxes represent implicit fees for local amenities, such as school quality, it is appropriate to exclude this component of the tax burden. Nevertheless, to explore the effect of state taxes, we add a flat 8% marginal tax rate at all levels of income. In practice, state taxes appear mildly regressive (Institute on Taxation and Economic Policy, 2013), making this a conservative assumption with respect to our main results.

The effect of this modification is displayed in Figure 20. The higher marginal tax rates serve to make the schedule of MSWWs (calculated assuming an ETI of 0.3) decline more sharply with income, but preserve the sharp rise from 1980 to 1990. Top MSWWs remain above conventional values from 1990 to 2010.

Figure 20: Average MSWWs, plotted against disposable income in 1980, 1990, 2000, and 2010, with an additional 8 percentage points added to the marginal tax rate at each income level to approximate state and local taxes.

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 changes to the progressivity of the tax code. To isolate the effect of policy changes alone, Figure 21 shows what this evolution would have looked like if market income inequality had not risen over time. 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. As would be expected, MSWWs on high earners are somewhat lower when the effect of rising inequality is removed. Yet the qualitative resemblance between Figures 3 and 21 demonstrates the importance of tax reforms as the key driver of large changes to the implicit MSWWs—i.e., the large increase in the MSWWs on high earners from 1980 to 1990 is virtually identical in the two figures.

Figure 21: 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.

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) 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.

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 22 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 22: 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.

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, Mathias, Dolls, Neumann, Dirk, Peichl, Andreas and Siegloch, Sebastian. (2014a). ‘Comparing Inequality Aversion across Countries when Labor Supply Responses Differ’, *International Tax and Public Finance* (forthcoming).

- Bargain, Olivier, Mathias, Dolls, Neumann, Dirk, Peichl, Andreas and Siegloch, Sebastian.** (2014*b*). ‘Tax-Benefit Revealed Social Preferences in Europe and the US’, *Annals of Economics and Statistics* (forthcoming).
- Bargain, Olivier and Peichl, Andreas.** (2013). ‘Steady-state labor supply elasticities: A survey’, *ZEW Discussion Paper No. 13-084*.
- 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.
- Birnbaum, J.H. and Murray, A.S.** (1987), *Showdown at Gucci Gulch: Lawmakers, Lobbyists, and the Unlikely Triumph of Tax Reform*, Random House Publishing Group.
URL: <http://books.google.com/books?id=ApGRAAAAMAAJ>
- 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.
- Burkhauser, Richard V, Feng, Shuaizhang, Jenkins, Stephen P and Larrimore, Jeff.** (2012). ‘Recent trends in top income shares in the United States: reconciling estimates from March CPS and IRS tax return data’, *Review of Economics and Statistics* 94(2), 371–388.
- Chetty, Raj.** (2009). ‘Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance’, *American Economic Journal: Economic Policy* 1(2), 31–52.
- 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*). ‘Effective Marginal Tax Rates for Low- and Moderate-Income Workers’.
- Congressional Budget Office.** (2012*b*). ‘A Review of Recent Research on Labor Supply Elasticities’.

- Congressional Budget Office.** (2013). ‘The Distribution of Household Income and Federal Taxes, 2010’.
- 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.** (1999). ‘Evidence on the high-income Laffer curve from six decades of tax reform’, *Brookings Papers on Economic Activity* pp. 1–64.
- Goolsbee, Austan.** (2000). ‘What Happens When You Tax the Rich? Evidence from Executive Compensation’, *The Journal of Political Economy* 108(2), 352–378.
- Gruber, Jon and Saez, Emmanuel.** (2002). ‘The elasticity of taxable income: evidence and implications’, *Journal of Public Economics* 84(1), 1–32.
- 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.
- Institute on Taxation and Economic Policy.** (2013), *Who Pays?: A Distributional Analysis of the Tax Systems in All 50 States*, Institute on Taxation and Economic Policy.
- 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’, *NBER Working Paper No. 18865* .

- Kuziemko, Ilyana and Washington, Ebonya.** (2014). ‘Support for Redistribution in an Age of Rising Inequality’, Working paper.
- Lindsey, Lawrence B.** (1987). ‘Individual taxpayer response to tax cuts: 1982–1984: with implications for the revenue maximizing tax rate’, *Journal of Public Economics* 33(2), 173–206.
- 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, Thomas and Saez, Emmanuel.** (2003). ‘Income inequality in the United States, 1913–1998’, *The Quarterly Journal of Economics* 118(1), 1–41.
- 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.
- Weinzierl, Matthew.** (2014a). ‘The promise of positive optimal taxation: normative diversity and a role for equal sacrifice’, *Journal of Public Economics* 118, 128–142.
- Weinzierl, Matthew.** (2014b). ‘Revisiting the Classical View of Benefit-Based Taxation’, Working Paper 14-101.
- Zoutman, Floris T, Jacobs, Bas and Jongen, Egbert LW.** (2013a). ‘Optimal redistributive taxes and redistributive preferences in the Netherlands’, Working paper.
- Zoutman, Floris T, Jacobs, Bas and Jongen, Egbert LW.** (2013b). ‘Revealed Social Preferences of Dutch Political Parties’, Working paper.