

# The disutility of uncertainty: welfare costs of noisy tax perceptions

MAX GHENIS

February 23, 2026

# Contents

0.1	Introduction . . . . .	1
0.1.1	Introduction . . . . .	2
0.1.2	Literature review . . . . .	4
0.2	Model . . . . .	7
0.2.1	Theoretical framework . . . . .	8
0.2.2	Calibration . . . . .	13
0.3	Results . . . . .	16
0.3.1	Results . . . . .	17
0.3.2	Extensions and robustness . . . . .	20
0.4	Conclusion . . . . .	23
0.4.1	Conclusion . . . . .	24
0.5	Appendix . . . . .	26
0.5.1	Technical appendix . . . . .	27

**0.1 Introduction**

### 0.1.1 Introduction

Most American taxpayers cannot accurately state their marginal tax rate. Gideon [2017] finds that only 33.7% of respondents correctly understand that their marginal rate exceeds their average rate—a basic structural feature of a progressive tax system. Rees-Jones and Taubinsky [2020] document that taxpayers’ perceived marginal rates deviate from their true rates with a standard deviation on the order of 12 percentage points, and that these errors are not confined to the financially unsophisticated. The source of this confusion is not ignorance of tax law in the abstract but rather the complexity of the code itself: federal brackets interact with state income taxes, payroll tax phase-outs, the earned income tax credit (EITC), child tax credit phase-ins and phase-outs, and the alternative minimum tax (AMT) to produce effective marginal rates that no taxpayer can compute in their head.

This paper asks a simple question: how much does this misperception cost? If workers choose their labor supply based on a noisy signal of their true marginal rate, they will systematically work too much or too little relative to the optimum. The resulting utility loss—aggregated across the workforce—constitutes a deadweight loss from tax complexity that is distinct from, and additional to, compliance costs and the standard Harberger triangles of distortionary taxation.

I develop a model in which a worker has quasilinear-isoelastic preferences:

$$U(C, h) = C - \psi \frac{h^{1+1/\varepsilon}}{1 + 1/\varepsilon} \quad (1)$$

where  $C$  is consumption,  $h$  is hours of labor,  $\varepsilon$  is the Frisch elasticity of labor supply, and  $\psi$  is a scale parameter. Given a linear tax rate  $\tau$  and hourly wage  $w$ , the budget constraint is  $C = w(1 - \tau)h$ , and optimal hours are  $h^* = (w(1 - \tau)/\psi)^\varepsilon$ . Now suppose the worker does not observe  $\tau$  but instead receives a signal  $\hat{\tau} = \tau + \delta$ , where  $\delta \sim N(0, \sigma^2)$ . The worker optimizes against  $\hat{\tau}$  but realizes utility at the true rate. The individual deadweight loss is  $U(h^*) - U(\hat{h})$ , where  $\hat{h}$  denotes hours chosen under the perceived rate. A second-order Taylor expansion yields the expected per-worker deadweight loss:

$$\frac{E[\text{DWL}]}{\text{earnings}} \approx \frac{1}{2} \frac{\varepsilon \sigma^2}{1 - \tau} \quad (2)$$

This formula has an intuitive structure. Deadweight loss is proportional to the labor supply elasticity  $\varepsilon$ , because more elastic workers distort their behavior more in response to misperceived prices. It is proportional to  $\sigma^2$ , reflecting the variance of the misperception error—the noise in the signal. And it is inversely proportional to  $1 - \tau$ , the net-of-tax share, because a given absolute misperception error represents a larger proportional distortion when the net-of-tax wage is smaller.

I calibrate the model using three empirical inputs. First, the Frisch elasticity of labor supply is set to  $\varepsilon = 0.33$ , the central estimate from the meta-analysis of Chetty [2012]. Second, the misperception standard deviation is set to  $\sigma = 0.12$ , consistent with the distribution of errors documented in Rees-Jones and Taubinsky [2020]. Third, the mean marginal tax rate is  $\bar{\tau} = 0.30$ , drawing on Congressional Budget Office [2016] estimates of effective marginal rates that account for federal income taxes, payroll taxes, and means-tested transfers. Together with mean annual earnings of \$55,000 and 160 million workers, the central estimate of aggregate deadweight loss is \$30 billion per year, or 0.11% of GDP. Sensitivity analysis across a  $3 \times 3$  grid of elasticities ( $\varepsilon \in \{0.25, 0.33, 0.50\}$ ) and misperception levels ( $\sigma \in \{0.08, 0.12, 0.15\}$ ) generates a range of 0.04% to 0.25% of GDP, or \$10 to \$71 billion.

These magnitudes represent 4–25 percent of the conventional DWL of the income tax estimated by Feldstein [1999], who places the deadweight loss of the income tax at roughly 1% of GDP. Skinner [1988] showed that uncertainty about *future* tax policy imposes welfare costs of similar order; the present paper complements that finding by showing that uncertainty about *current* rates—driven by code complexity rather than policy instability—is also costly.

This paper makes three contributions. First, it derives a closed-form expression for the expected deadweight loss from tax misperception under standard preferences, making the welfare cost transparent and easy to calibrate. Second, it provides the first calibration of this cost using empirically grounded estimates of misperception dispersion from Rees-Jones and Taubinsky [2020] and Gideon [2017], combined with standard labor supply parameters. Third, it shows that a utilitarian social planner who accounts for misperception chooses a lower optimal linear tax rate—42.9% versus 44.5%

under perfect information—because the marginal welfare cost of taxation is amplified by the noise in taxpayers' perception of rates.

The remainder of the paper is organized as follows. The model chapter formalizes the preference structure, derives the deadweight loss formula, and characterizes optimal taxation under misperception. The calibration chapter describes the empirical inputs and sensitivity analysis. The results chapter presents the baseline estimates, distributional implications, and the planner's problem. The conclusion discusses policy implications—in particular, why reducing  $\sigma$  through simplification dominates reducing  $\tau$  through rate cuts—and outlines extensions to non-linear schedules and heterogeneous misperception.

## 0.1.2 Literature review

This paper draws on several intersecting literatures: tax salience and misperception, the welfare costs of tax uncertainty, optimal taxation with behavioral agents, the elasticity of taxable income, tax complexity and compliance costs, and the economics of information provision. I survey each in turn, highlighting how prior work motivates both the theoretical framework and the empirical calibration strategy of this paper.

### Tax salience and misperception

The foundational insight that the presentation and visibility of taxes shapes economic behavior comes from Chetty et al. [2009], who combined field experiments and quasi-experimental evidence to show that posting tax-inclusive prices at a grocery store reduced demand by approximately 8 percent, despite no change in actual tax liability. Their theoretical framework formalized the notion that agents respond not to actual tax rates but to *perceived* tax rates, and that the wedge between the two generates both behavioral distortions and welfare losses distinct from those arising under perfect information. This insight—that how taxes are communicated matters as much as what they are—provides the conceptual foundation for the present analysis of misperception-induced deadweight loss.

Finkelstein [2009] provides complementary evidence from electronic highway tolling: when tolls became invisible (deducted electronically rather than paid at a booth), toll rates rose, suggesting that reduced salience allows governments to levy higher taxes — a pattern consistent with the prediction that misperception of rates has welfare consequences.

Building directly on the “schmeduling” concept introduced by Liebman and Zeckhauser [2004], who first documented that taxpayers mentally linearize progressive tax schedules by treating their marginal rate as if it were their average rate, Rees-Jones and Taubinsky [2020] provided the first rigorous empirical measurement of such misperception. Using linked survey and administrative data, they found that approximately 43 percent of taxpayers employ an “ironing” heuristic, collapsing the graduated rate structure into a single perceived rate. For my calibration, their estimates imply a standard deviation of tax rate misperception of approximately  $\sigma \approx 0.12$  around the true marginal rate—the key parameter governing the magnitude of welfare losses in my model. The fact that this misperception is systematic rather than random (with ironing biasing perceived rates toward the average rate) means that approximately 43 percent of the working population makes labor supply decisions on the basis of incorrect price signals.

Gideon [2017] provides complementary survey evidence: only 33.7 percent of Americans correctly identify that their marginal tax rate exceeds their average tax rate. The remaining two-thirds either believe the rates are equal (consistent with the ironing heuristic) or report the marginal rate as lower—patterns that confirm the prevalence of misperception documented by Rees-Jones and Taubinsky in an independent sample. Feldman et al. [2016] offer additional evidence from the Child Tax Credit, showing that confusion about the structure of a single tax provision generates large behavioral distortions, reinforcing the view that misperception is not confined to the overall rate structure but pervades individual provisions as well. Ballard and Gupta [2018] find that 85 percent of Michigan taxpayers overstate their average federal income tax rate, with a mean overstatement of 11.6 percentage points, suggesting that the misperception problem may be even more severe than the ironing heuristic alone would imply.

### Tax uncertainty and welfare

The welfare cost of tax *policy* uncertainty—uncertainty about what future tax rates will be—was first quantified by Skinner [1988], who developed a two-period model in which agents must commit to savings decisions before the resolution of uncertainty about future tax rates. Skinner estimated welfare losses of approximately 0.4 percent of GDP annually, establishing that tax uncertainty imposes first-order costs comparable to major tax distortions. This paper complements Skinner’s analysis by studying a distinct source of uncertainty: agents’ imperfect knowledge of their *current* marginal tax rates. While Skinner’s costs arise from genuine randomness in the policy environment, mine arise from complexity-induced misperception of the existing tax code. The two sources of welfare loss are

additive: reducing misperception does not eliminate policy uncertainty, and stabilizing policy does not help agents who cannot parse the current schedule.

Alm [1988] extended the analysis of tax uncertainty to individual decision difficulty, demonstrating that when agents cannot easily determine their marginal tax rates, they make systematically suboptimal choices even absent policy uncertainty. Hassett and Metcalf [1999] examined how tax rate uncertainty affects investment decisions, finding that uncertainty can either stimulate or depress investment depending on the structure of adjustment costs and the degree of irreversibility—highlighting that uncertainty effects are not uniform across different margins of economic decision-making. More recently, Baker et al. [2016] constructed a comprehensive economic policy uncertainty index and showed that uncertainty shocks reduce investment and employment, with tax policy uncertainty accounting for a large share of the overall effect. Their macroeconomic evidence supports the theoretical prediction that tax uncertainty has first-order welfare consequences.

### Optimal taxation with behavioral agents

The classical optimal income tax literature, inaugurated by Mirrlees [1971], assumes that agents have perfect information about tax schedules and optimize accordingly. Diamond and Saez [2011] provide a modern synthesis of optimal income taxation that retains this full-information assumption while incorporating empirical estimates of behavioral elasticities. Saez [2001] formalized the sufficient-statistics approach to optimal income taxation, showing that the optimal tax schedule depends on the elasticity of taxable income and the shape of the income distribution. The present paper applies an analogous sufficient-statistics logic: the welfare cost of misperception reduces to three observable quantities. Yet a growing body of work has begun to relax the perfect-information assumption in substantive ways.

Most directly relevant, Farhi and Gabaix [2020] develop a general framework for optimal taxation when agents exhibit behavioral biases, including misperception of tax rates. Their sufficient-statistics approach modifies the standard optimal tax formula by introducing a “behavioral wedge” that accounts for the divergence between actual and perceived marginal rates. In their notation, the optimal tax formula includes a correction term proportional to the covariance between the behavioral wedge and the social marginal welfare weight. When misperception is uncorrelated with income (as in the baseline specification here), their framework predicts that the optimal tax rate falls—consistent with the results below. The present paper contributes a specific quantification of this channel: using the Rees-Jones and Taubinsky estimates of  $\sigma$ , I derive a closed-form expression for the additional deadweight loss from misperception and show how it alters the optimal tax rate under a simple linear tax. The linear-tax setting permits closed-form solutions and transparent comparative statics, complementing the greater generality but reduced tractability of the Farhi-Gabaix framework. Gerritsen [2016] reaches a complementary conclusion from a different modeling approach, showing that when agents do not maximize well-being (due to misperception or other behavioral frictions), the optimal nonlinear tax schedule differs systematically from the Mirrlees benchmark. Chetty et al. [2009] provide a behavioral framework for evaluating when policymakers should correct versus exploit misperceptions, concluding that correction is generally welfare-improving when misperception is unrelated to the policy objective.

### Elasticity of taxable income

The deadweight loss formula,  $\mathbb{E}[\text{DWL}]/\text{earnings} = \frac{1}{2}\varepsilon\sigma^2/(1 - \tau)$ , is structurally analogous to the Harberger triangle expression for the standard DWL of taxation, but applied to the variance of perceived rates rather than the level of the tax rate. The empirical content of the formula depends critically on the labor supply elasticity  $\varepsilon$ .

Feldstein [1999] pioneered the elasticity-of-taxable-income (ETI) approach to measuring the efficiency cost of taxation, arguing that the ETI captures all behavioral margins of response—not just hours worked but also effort, occupation choice, tax avoidance, and evasion. Saez et al. [2012] provide a comprehensive review of ETI estimates, concluding that a central value of approximately 0.25 is appropriate for the U.S. income tax, though estimates range widely depending on the population and identification strategy. For my calibration, I draw on Chetty [2012], who reconciles the apparent

conflict between small micro elasticities (from tax reform studies) and large macro elasticities (from cross-country comparisons) by showing that optimization frictions attenuate short-run responses. His preferred Frisch elasticity estimate of 0.5 on the extensive margin, combined with an intensive-margin elasticity near 0.33, informs my baseline  $\varepsilon = 0.33$ . Keane and Rogerson [2012] reach a similar conclusion through structural estimation, arguing that properly accounting for human capital, progressive taxation, and life-cycle dynamics yields macro-consistent elasticities from micro data.

### **Tax complexity and compliance costs**

The welfare costs of tax complexity extend beyond misperception of rates to include direct compliance costs and reduced take-up of beneficial provisions. Slemrod [2005] documents that U.S. state income tax systems vary widely in complexity and that this complexity imposes real resource costs on taxpayers and administrators—costs that function as a hidden tax. Goldin [2018] shows that complexity in the Earned Income Tax Credit reduces take-up among eligible households, meaning that the intended redistribution is partially undone by the very mechanism designed to deliver it. Benzarti [2020] estimates compliance costs using revealed preferences—observing that taxpayers forgo substantial tax savings rather than itemize deductions—and finds costs larger than survey-based estimates suggest.

Chetty and Saez [2013] provides the closest empirical analog to the mechanism in my model. In a field experiment, they provided EITC recipients with information about how the credit’s phase-in and phase-out affect their marginal incentives. Recipients who received the information adjusted their earnings in subsequent years, demonstrating that misperception of marginal rates has real behavioral consequences that can be corrected through information provision. This result directly supports my conclusion that reducing  $\sigma$  through better information generates welfare gains. Abeler and Jäger [2015] complement this finding with laboratory evidence showing that subjects respond less to tax incentives when the tax schedule is complex, even when they have access to all relevant information—suggesting that cognitive costs of processing complex schedules are themselves a source of effective misperception.

### **Information provision and economic decisions**

A broader literature has established that correcting misperceptions generates first-order welfare improvements across many domains, including retirement savings [Duflo and Saez, 2003], school choice [Hastings and Weinstein, 2008], and returns to education [Jensen, 2010]. My analysis extends this logic to the tax domain, where misperception of the marginal tax rate directly distorts the perceived return to labor supply.

**0.2 Model**

### 0.2.1 Theoretical framework

This section develops a model of labor supply under tax misperception. A worker who misperceives her marginal tax rate chooses a suboptimal quantity of labor, generating a welfare loss that depends on the elasticity of labor supply, the variance of the misperception, and the level of the tax rate itself. The central result is a closed-form expression for the expected deadweight loss as a fraction of earnings.

#### Setup

Consider a population of workers indexed by  $i$ . Worker  $i$  earns hourly wage  $w_i$  and faces a linear marginal tax rate  $\tau_i$ . Due to the complexity of the tax code—overlapping brackets, phase-ins, phase-outs, credits, and the interaction of federal, state, and payroll taxes—the worker does not observe  $\tau_i$  directly. Instead, she perceives a noisy signal

$$\hat{\tau}_i = \tau_i + \delta_i, \quad \delta_i \sim N(0, \sigma^2), \quad (3)$$

where  $\sigma$  measures the severity of misperception—the misperception error has a standard deviation of  $\sigma$  (e.g., 12 percentage points when  $\sigma = 0.12$ ). The noise  $\delta_i$  captures the gap between the worker’s mental model of her tax rate and the actual statutory-plus-effective rate she faces. This formulation nests perfect information as the special case  $\sigma = 0$ .

The normality assumption is a modeling convenience that delivers tractable closed-form results. The DWL formula depends only on  $E[\delta^2] = \sigma^2$ , so it holds for any symmetric distribution with variance  $\sigma^2$ , not just the Gaussian. Departures from normality (e.g., heavier tails or discreteness in actual misperception) would affect higher-order corrections to the Taylor approximation but not the leading term. The Monte Carlo validation in the approximation accuracy subsection confirms that the second-order formula is robust to the clipping of extreme draws at the  $[0, 1]$  boundary.

The mean-zero assumption ( $E[\delta_i] = 0$ ) is a simplification. Rees-Jones and Taubinsky [2020] document that misperception is partly systematic: “ironing” taxpayers tend to perceive their marginal rate as closer to their average rate, which in a progressive system means  $E[\delta_i] < 0$  (underestimation). Decomposing misperception into bias  $\mu_\delta = E[\delta_i]$  and noise  $\sigma^2 = \text{Var}(\delta_i)$ , the DWL formula generalizes to include a bias term:  $E[\text{DWL}] \approx \frac{1}{2}\varepsilon wh^*(\mu_\delta^2 + \sigma^2)/(1 - \tau)$ . If systematic bias is present, the total welfare cost would be larger than the mean-zero case studied here. The calibrated  $\sigma = 0.12$  already incorporates both systematic and idiosyncratic components of the cross-sectional dispersion in perceived rates, so it partially captures the bias channel. Specifically, the calibrated  $\sigma = 0.12$  is best interpreted as the root-mean-square error (RMSE) of the misperception distribution, since it is derived from the cross-sectional dispersion of (perceived - actual) rates which includes any systematic component. Under this interpretation,  $\sigma^2 = \mu_\delta^2 + \text{Var}(\delta_i)$ , and the mean-zero DWL formula  $\frac{1}{2}\varepsilon wh^* \sigma^2 / (1 - \tau)$  remains correct with  $\sigma$  interpreted as the RMSE rather than the standard deviation of mean-zero noise.

Preferences are quasilinear in consumption with isoelastic labor disutility:

$$U(C, h) = C - \psi \frac{h^{1+1/\varepsilon}}{1 + 1/\varepsilon}, \quad (4)$$

where  $C$  is consumption,  $h$  is hours of labor,  $\varepsilon > 0$  is the Frisch elasticity of labor supply, and  $\psi > 0$  is a scale parameter governing the disutility of work. Quasilinear preferences eliminate income effects on labor supply, so the Frisch elasticity coincides with the uncompensated and compensated elasticities. This tractability isolates the substitution-effect channel through which tax misperception distorts behavior, as discussed below.

The worker’s budget constraint is

$$C = w(1 - \tau)h + v, \quad (5)$$

where  $v$  is a lump-sum transfer (or non-labor income). Because utility is linear in  $C$ , the transfer  $v$  does not affect the labor supply decision—it shifts consumption one-for-one but leaves the first-order condition for hours unchanged. This separation is essential: it means that misperception of  $\tau$  affects only the hours margin, not the marginal utility of income.

### Optimal labor supply

A worker who correctly perceives her tax rate  $\tau$  maximizes  $U$  by setting the marginal benefit of an hour of work equal to its marginal cost:

$$w(1 - \tau) = \psi h^{1/\varepsilon}. \quad (6)$$

Solving for optimal hours:

$$h^*(\tau) = \left( \frac{w(1 - \tau)}{\psi} \right)^\varepsilon. \quad (7)$$

Hours are decreasing in the tax rate and increasing in the wage, with the Frisch elasticity  $\varepsilon$  governing the responsiveness of labor supply to the net-of-tax wage.

### Misperceived labor supply

A worker who perceives the tax rate to be  $\hat{\tau} = \tau + \delta$  solves the same first-order condition but evaluated at  $\hat{\tau}$ :

$$h(\hat{\tau}) = \left( \frac{w(1 - \hat{\tau})}{\psi} \right)^\varepsilon = \left( \frac{w(1 - \tau - \delta)}{\psi} \right)^\varepsilon. \quad (8)$$

When  $\delta > 0$  the worker overestimates her tax rate and undersupplies labor; when  $\delta < 0$  she underestimates the rate and oversupplies labor. In both cases, utility evaluated at the true tax rate  $\tau$  is lower than it would be at the optimum  $h^*(\tau)$ .

### Individual deadweight loss

Define the individual deadweight loss from a misperception of magnitude  $\delta$  as the utility gap between optimal and misperceived behavior, with consumption evaluated at the true tax rate in both cases:

$$\text{DWL}(\delta) = U(h^*(\tau)) - U(h(\tau + \delta)). \quad (9)$$

Here  $U(h)$  is shorthand for indirect utility at hours  $h$  given the true budget constraint:  $U(h) = w(1 - \tau)h + v - \psi h^{1+1/\varepsilon}/(1 + 1/\varepsilon)$ . Since  $h^*(\tau)$  maximizes this expression,  $\text{DWL}(\delta) \geq 0$  for all  $\delta$ , with equality only at  $\delta = 0$ .

### Second-order approximation

To obtain a closed-form expression for expected DWL, expand  $U(h)$  around the optimum  $h^*$ . Write  $\Delta h = h(\hat{\tau}) - h^*(\tau)$  for the hours distortion.

The first derivative of indirect utility with respect to hours is

$$U_h = w(1 - \tau) - \psi h^{1/\varepsilon}, \quad (10)$$

which equals zero at  $h = h^*$  by the first-order condition. The second derivative is

$$U_{hh} = -\frac{\psi}{\varepsilon} h^{*(1/\varepsilon - 1)} = -\frac{w(1 - \tau)}{\varepsilon h^*}, \quad (11)$$

where the second equality substitutes  $\psi = w(1 - \tau)/h^{*1/\varepsilon}$  from the first-order condition. Since  $U_h = 0$  at the optimum, the Taylor expansion gives

$$\text{DWL}(\delta) \approx \frac{1}{2} |U_{hh}| (\Delta h)^2. \quad (12)$$

Next, compute  $\Delta h$ . Since  $h(\hat{\tau}) = (w(1 - \tau - \delta)/\psi)^\varepsilon$ , a first-order expansion around  $\delta = 0$  yields

$$\Delta h \approx -\varepsilon h^* \frac{\delta}{1 - \tau}. \quad (13)$$

Substituting both expressions:

$$\text{DWL}(\delta) \approx \frac{1}{2} \cdot \frac{w(1-\tau)}{\varepsilon h^*} \cdot \varepsilon^2 h^{*2} \frac{\delta^2}{(1-\tau)^2} = \frac{1}{2} \varepsilon w h^* \frac{\delta^2}{1-\tau}. \quad (14)$$

The deadweight loss is quadratic in the misperception  $\delta$ : small errors are relatively harmless, but large errors are disproportionately costly. This convexity is central to the welfare analysis.

### Expected deadweight loss

Taking expectations over  $\delta \sim N(0, \sigma^2)$  and using  $E[\delta^2] = \sigma^2$ :

$$E[\text{DWL}] = \frac{1}{2} \varepsilon \cdot w \cdot h^* \cdot \frac{\sigma^2}{1-\tau}. \quad (15)$$

Since  $w \cdot h^*$  equals earnings, this can be expressed as a fraction of earnings:

$$\boxed{\frac{E[\text{DWL}]}{\text{earnings}} \approx \frac{1}{2} \varepsilon \cdot \frac{\sigma^2}{1-\tau}}. \quad (16)$$

This is the paper's core formula. It gives the welfare cost of tax misperception as a share of labor income, depending on three key parameters: the Frisch elasticity, the variance of misperception, and the tax rate.

**Approximation accuracy** Because the formula is a second-order Taylor expansion of utility around the optimum, it is exact only for infinitesimal  $\delta$  and may diverge for large misperceptions. To assess its accuracy over the parameter ranges used in this paper, I compare the analytical approximation to a Monte Carlo estimate (500,000 draws of  $\delta \sim N(0, \sigma^2)$ , evaluating exact utility at each draw). At the baseline calibration ( $\varepsilon = 0.33$ ,  $\sigma = 0.12$ ,  $\tau = 0.30$ ), the analytical formula understates the Monte Carlo estimate by 1.1%. At the most extreme corner of the sensitivity grid ( $\varepsilon = 0.50$ ,  $\sigma = 0.15$ ,  $\tau = 0.43$ ), where  $\sigma/(1-\tau) = 0.26$ , the approximation understates the true expected DWL by 4.7%. The closed-form expression is therefore reliable across the entire parameter space considered here.

### Discussion of the formula

The formula reveals several comparative statics.

*Quadratic in  $\sigma$ .* Doubling the standard deviation of misperception quadruples the expected welfare cost. This convexity implies that reducing misperception from  $\sigma = 0.15$  to  $\sigma = 0.10$  (a 56% reduction in  $\sigma^2$ ) lowers DWL by more than reducing from  $\sigma = 0.10$  to  $\sigma = 0.05$  (a 75% reduction in  $\sigma^2$  but from a lower base).

*Linear in  $\varepsilon$ .* More elastic labor supply amplifies the cost of misperception because workers with elastic supply distort their hours more in response to a given perceived price change. If labor supply were perfectly inelastic ( $\varepsilon = 0$ ), misperception would have no welfare cost at all—workers would supply the same hours regardless of their beliefs about tax rates. The empirical magnitude of  $\varepsilon$  is therefore a key input to the calibration.

A note on elasticity concepts is warranted. The Frisch elasticity  $\varepsilon$  in the formula governs the responsiveness of hours to the net-of-tax wage, holding the marginal utility of wealth constant. Under quasilinear preferences, the Frisch, compensated (Hicksian), and uncompensated (Marshallian) elasticities all coincide because there is no income effect. The elasticity of taxable income (ETI) is a broader concept that captures responses on all margins—hours, effort, avoidance, and evasion—and is therefore typically larger than the hours elasticity. The baseline calibration uses  $\varepsilon = 0.33$ , which corresponds to the hours-only Frisch elasticity from Chetty [2012]. If misperception also distorts avoidance and reporting decisions, the relevant elasticity would be closer to the ETI (approximately 0.25–0.50 per Saez et al. [2012]), and the welfare costs could be correspondingly larger or smaller depending on the specific margin. The present analysis isolates the labor supply channel for transparency.

*Denominator  $(1-\tau)$ .* Higher tax rates amplify the welfare cost of misperception. This occurs because the same absolute misperception  $\delta$  represents a larger proportional change in the net-of-tax

wage when  $\tau$  is high. A 10-percentage-point misperception shifts the net-of-tax rate from 70% to 60% (a 14% change) when  $\tau = 0.30$ , but from 40% to 30% (a 25% change) when  $\tau = 0.60$ .

*Connection to Harberger triangles.* The formula is a direct application of the Harberger approximation to the deadweight loss of taxation, except that the distortionary “wedge” is not the tax itself but the misperception  $\delta$ . Just as the traditional DWL of a tax  $\tau$  is approximately  $\frac{1}{2}\varepsilon\tau^2/(1-\tau)$  times earnings, the DWL of misperception replaces  $\tau^2$  with  $\sigma^2$ . The structural parallel clarifies that misperception generates a welfare loss with the same second-order form as a tax distortion, though unlike a tax, misperception generates no offsetting revenue.

### Social planner extension

Consider a utilitarian social planner who chooses a linear tax rate  $\tau$  and finances a uniform lump-sum transfer (demogrant)  $v$  from the resulting revenue. The government budget constraint requires

$$v = \frac{\tau}{N} \sum_{i=1}^N w_i h_i. \quad (17)$$

The social welfare function uses inverse-wage weights  $\omega_i = \bar{w}/w_i$  to create a redistribution motive, so  $W = (1/N) \sum_i \omega_i U_i$ . Workers with below-average wages receive higher weight, reflecting the standard utilitarian motive for progressive taxation. Under perfect information, these weights generate an interior optimal tax rate that balances redistribution gains against efficiency losses.

The disutility scale  $\psi$  is normalized to one. Since rescaling  $\psi$  by any constant  $\lambda$  multiplies all hours, all revenues, and all utility levels by the same factor  $(1/\lambda)^\varepsilon$ , the optimal tax rate  $\tau^*$  is invariant to the choice of  $\psi$ . The calibrated optimal tax therefore depends only on the wage distribution and the elasticity, not on the normalization of hours.

Under perfect information ( $\sigma = 0$ ), each worker chooses  $h_i = h^*(\tau_i)$  and the planner selects  $\tau$  to balance the equity gains from redistribution against the efficiency losses from labor supply distortion.

When workers misperceive the tax rate ( $\sigma > 0$ ), the planner faces an additional cost. Workers choose  $h(\hat{\tau}_i)$  instead of  $h^*(\tau_i)$ , reducing both private welfare and government revenue (since misallocated labor reduces the tax base). The expected per-worker welfare loss from misperception,  $\frac{1}{2}\varepsilon w h^* \sigma^2 / (1-\tau)$ , is increasing in  $\tau$  through the  $(1-\tau)$  denominator. This means that at the margin, raising the tax rate imposes a larger misperception cost when rates are already high.

The optimal tax rate  $\tau^*$  is therefore decreasing in  $\sigma$ . As misperception worsens, the planner optimally reduces the tax rate, accepting less redistribution in exchange for lower misperception-induced deadweight loss. In the extreme of perfect information ( $\sigma = 0$ ), the planner chooses the standard optimal tax; as  $\sigma$  grows large, the optimal rate converges toward zero because the misperception cost eventually dominates the redistribution benefit.

The model predicts that a reduction in  $\sigma$  permits higher  $\tau$  at constant total DWL. Whether real-world measures to reduce misperception are compatible with maintaining or increasing progressivity is an empirical question beyond the scope of this model.

### Why quasilinear preferences, not Cobb-Douglas?

A natural alternative specification is Cobb-Douglas utility over leisure  $L$  and consumption  $C$ :

$$U(L, C) = L^\alpha C^\beta, \quad (18)$$

with time constraint  $L + h = T$ . Under this specification, optimal leisure is  $L^* = \alpha T / (\alpha + \beta)$  when the worker receives no transfers, and this expression is independent of the net-of-tax wage  $w(1-\tau)$ . The income and substitution effects of a wage change exactly cancel, producing a knife-edge result: hours are invariant to the perceived tax rate, so misperception has zero welfare cost.

This invariance is a well-known property of Cobb-Douglas preferences [Block et al., 1980], and while it provides a useful theoretical benchmark, it is empirically implausible as a description of labor supply responses to taxation. A large body of evidence reviewed by Chetty [2012] documents positive compensated labor supply elasticities, meaning that the substitution effect dominates the income effect

on the relevant margins. The quasilinear specification isolates this substitution effect by eliminating the income effect entirely, making it the natural workhorse for analyzing how tax misperception distorts labor supply.

More precisely, the welfare cost of misperception operates through the substitution effect: a worker who overestimates her tax rate reduces labor supply because she perceives a lower return to work, not because she feels wealthier. By shutting down income effects, the quasilinear model isolates this channel cleanly. Any utility specification with positive compensated elasticity would generate qualitatively similar results; the quasilinear form simply yields the most transparent closed-form expressions.

A further limitation of the model is its focus on the intensive margin of labor supply. Most workers cannot continuously adjust their hours; the majority of employees work fixed schedules set by employers [Keane and Rogerson, 2012]. However, the intensive-margin framework can be reinterpreted as capturing responses at the extensive margin (employment entry and exit) or along the effort margin, both of which are governed by similar elasticities. Moreover, the minority of workers with flexible hours—the self-employed, gig workers, and those choosing between part-time and full-time positions—may have substantially higher effective elasticities than the population average, partially offsetting the lower responsiveness of fixed-hours workers. The calibrated  $\varepsilon = 0.33$  represents a blend of intensive and extensive margin responses, as discussed by Chetty [2012].

### Misperception versus inattention

The model assumes that workers actively optimize against a perceived tax rate  $\hat{\tau}$ , which may differ from the true rate. An alternative behavioral model posits that workers are *inattentive* to the tax rate—they assign zero or near-zero weight to the marginal rate in their labor supply decisions and instead follow heuristics or social norms (e.g., working 40 hours because the employer expects it). Under the inattention model, the welfare cost depends on the degree of under-reaction to tax incentives (the “attention parameter”  $\theta$  in the framework of Chetty et al. [2009]), not on the variance of perceived rates.

The distinction matters empirically. Survey evidence on perceived rates [Rees-Jones and Taubinsky, 2020], [Gideon, 2017] reveals that taxpayers do form beliefs about their marginal rates—they are not ignoring the rate but rather estimating it with error. The behavioral responses documented by Chetty and Saez [2013], in which EITC recipients adjusted their earnings after receiving information about marginal incentives, suggest that workers do optimize against perceived rates rather than ignoring them entirely. The misperception framework adopted here is therefore the more appropriate model for the population segments that form and act on tax rate beliefs, while the inattention framework may apply to workers who delegate all tax decisions to accountants or employers. The two models are complements rather than substitutes, and a complete welfare accounting would include both channels.

## 0.2.2 Calibration

The expected DWL formula derived in the previous section— $E[\text{DWL}]/\text{earnings} \approx \frac{1}{2}\varepsilon\sigma^2/(1-\tau)$ —requires three inputs: the Frisch elasticity of labor supply  $\varepsilon$ , the standard deviation of tax misperception  $\sigma$ , and the average marginal tax rate  $\tau$ . This section draws on the empirical public finance literature to calibrate each parameter, then computes the implied welfare costs under a range of assumptions.

### Frisch elasticity ( $\varepsilon$ )

The Frisch (marginal-utility-constant) elasticity governs how strongly workers adjust their hours in response to changes in the perceived net-of-tax wage. Chetty [2012] provides the most comprehensive meta-analysis of this parameter, reconciling the long-standing gap between micro and macro estimates. Micro studies using individual-level panel data (tax reforms, lottery winners) typically find Frisch elasticities of 0.25–0.50 on the intensive margin. Macro estimates, which incorporate extensive-margin responses and general equilibrium adjustments, center around 0.50. Chetty argues that optimization frictions—precisely the kind of frictions this paper studies—attenuate micro estimates, so the true structural elasticity lies between the micro and macro figures.

I adopt a central estimate of  $\varepsilon = 0.33$ , which represents a conservative macro-consistent value that accounts for both intensive and extensive margin responses while remaining below the highest macro estimates. The sensitivity analysis varies  $\varepsilon$  over the range  $[0.25, 0.50]$ , spanning the micro consensus to the macro midpoint. Values below 0.25 would imply that labor supply barely responds to tax incentives (inconsistent with the macro evidence), while values above 0.50 would exceed most structural estimates.

### Misperception standard deviation ( $\sigma$ )

The standard deviation of tax misperception  $\sigma$  is the parameter most distinctive to this paper, and multiple lines of evidence inform its calibration.

Rees-Jones and Taubinsky [2020] provide the most detailed evidence on how taxpayers mentally approximate their tax schedules. They find that 43% of taxpayers use an “ironing” heuristic, mentally replacing the piecewise-linear statutory schedule with a linearized version centered on their average tax rate. Because the average rate is below the marginal rate for all taxpayers in a progressive system, this heuristic systematically underestimates the marginal rate. The typical ironer perceives a marginal rate roughly equal to her average rate, generating errors on the order of 5–15 percentage points depending on income.

Gideon [2017] provides direct survey evidence on tax rate knowledge. Only 33.7% of respondents correctly identify that their marginal tax rate exceeds their average tax rate—the basic qualitative feature of a progressive tax system. The remaining two-thirds either believe the rates are equal (consistent with the ironing heuristic) or believe the marginal rate is lower than the average rate (an error of both sign and magnitude). Gideon estimates mean absolute errors of approximately 15 percentage points relative to the true marginal rate.

The cross-sectional root-mean-square error of perceived minus actual marginal rates in Rees-Jones and Taubinsky [2020] is approximately 0.12, which I adopt directly as the calibrated  $\sigma$ . I use  $\sigma = 0.12$  as the central estimate and explore the range  $[0.08, 0.15]$ . The lower bound ( $\sigma = 0.08$ ) corresponds to a population with moderate confusion about bracket boundaries but correct understanding of the progressive structure. The upper bound ( $\sigma = 0.15$ ) reflects the raw survey evidence without adjustment for any respondent sophistication.

### Mean marginal tax rate ( $\bar{\tau}$ )

The Congressional Budget Office (CBO) has published detailed estimates of effective marginal tax rates (MTRs) faced by U.S. workers. Congressional Budget Office [2012] and Congressional Budget Office [2016] compute effective rates that incorporate federal income taxes, payroll taxes (employer and employee shares), state income taxes, and the phase-out of major transfer programs. These reports

Table 1: Sensitivity of deadweight loss estimates to Frisch elasticity and misperception standard deviation

$\varepsilon$	$\sigma$	Per-worker DWL	Total DWL (\$B)	% of GDP
0.25	0.08	\$63	\$10.1	0.04
0.25	0.12	\$141	\$22.6	0.08
0.25	0.15	\$221	\$35.4	0.13
0.33	0.08	\$83	\$13.3	0.05
0.33	0.12	\$187	\$29.9	0.11
0.33	0.15	\$292	\$46.7	0.17
0.50	0.08	\$126	\$20.1	0.07
0.50	0.12	\$283	\$45.3	0.16
0.50	0.15	\$442	\$70.7	0.25

find effective marginal rates of 25–32% for households in the middle three income quintiles, rising to 43% for high-income households (who face the top federal bracket plus the Medicare surtax).

For the representative-worker calculation, I use  $\bar{\tau} = 0.30$ , which reflects the earnings-weighted average across the working population. This figure includes the combined federal income tax rate (averaging roughly 18% across brackets), the employee share of payroll taxes (7.65%), and an average state income tax contribution of approximately 4%. The sensitivity to  $\bar{\tau}$  enters through the  $(1 - \tau)$  denominator: at  $\tau = 0.25$  the amplification factor is  $1/0.75 = 1.33$ , while at  $\tau = 0.43$  it rises to  $1/0.57 = 1.75$ .

### Labor market parameters

I calibrate the representative worker using standard Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS) data [Bureau of Labor Statistics, 2023]. Mean hourly wages across all occupations are approximately \$27.50 (2023 dollars). At 2,000 annual hours (50 weeks at 40 hours), this implies mean annual earnings of \$55,000. The total number of employed workers in the U.S. civilian labor force is approximately 160 million [Bureau of Labor Statistics, 2024]. U.S. GDP is approximately \$28 trillion [Bureau of Economic Analysis, 2024], used as the denominator for expressing aggregate costs as a share of national output.

For the optimal tax computation, I draw wages from a lognormal distribution with log-standard-deviation 0.5, calibrated to match the mean hourly wage of \$27.50. This produces a wage distribution with a Gini coefficient of approximately 0.28, which is conservative relative to the full U.S. hourly wage distribution (Gini  $\sim$  0.40).

### Baseline and sensitivity results

Applying the formula  $DWL \text{ per worker} = \frac{1}{2}\varepsilon \cdot \text{earnings} \cdot \sigma^2 / (1 - \tau)$  with earnings = \$55,000 and  $\tau = 0.30$ , and aggregating over 160 million workers:

The central estimate (row with  $\varepsilon = 0.33$ ,  $\sigma = 0.12$ ): the expected deadweight loss from tax misperception is approximately \$187 per worker, or \$29.9 billion in aggregate—0.11% of GDP. This is a cost that arises purely from the complexity of the tax code, independent of any distortionary effect of the tax rate itself.

Several features of the table merit emphasis. First, the quadratic dependence on  $\sigma$  is visible: moving from  $\sigma = 0.08$  to  $\sigma = 0.15$  (roughly doubling the misperception) nearly quadruples the per-worker DWL within each elasticity row. Second, the range across the full grid is wide—from \$10.1 billion to \$70.7 billion—reflecting genuine uncertainty about the structural parameters. The most conservative cell (\$10.1 billion at  $\varepsilon = 0.25$ ,  $\sigma = 0.08$ ) and the most aggressive (\$70.7 billion at  $\varepsilon = 0.50$ ,  $\sigma = 0.15$ ) differ by a factor of seven.

The central estimate of 0.11% of GDP is smaller than Skinner [1988]’s estimate of 0.4% from future tax policy uncertainty, which is appropriate since Skinner’s model includes both risk aversion and intertemporal channels that are absent here. The present paper estimates that the static misperception

channel—holding tax policy constant and considering only current-period confusion about marginal rates—generates welfare losses of 0.04–0.25% of GDP.

These calibrated values align with the code implementation in the `Calibration` class, which uses identical parameter values and the same analytical formula to produce the sensitivity grid. All results are reproducible from the accompanying code repository.

### Empirical marginal tax rate distribution

The representative-agent calculations above assume a single mean marginal rate  $\bar{\tau} = 0.30$  applied uniformly across workers. To assess the distributional implications of MTR heterogeneity, I compute household-level comprehensive marginal tax rates using PolicyEngine-US [PolicyEngine, 2024], a microsimulation model built on the Enhanced Current Population Survey. The simulation covers approximately 149 million working-age adults (18–64) with positive employment income in 2024.

Marginal tax rates are clipped to the interval  $[0, 0.99]$  before analysis: benefit phase-outs can produce MTRs above 100% (benefit cliffs), and refundable credits can produce negative MTRs. I clip rather than exclude these observations to avoid selection bias, though the choice has minimal impact on weighted means because extreme MTRs affect a small share of the weighted population.

The weighted mean comprehensive MTR is 0.25, somewhat below the CBO estimate of 0.30 used in the stylized calibration. This difference reflects the large mass of low-income workers who face zero or low federal income tax rates, partially offset by payroll taxes. The weighted standard deviation is 0.16, and the distribution spans from 0.00 at the 10th percentile to 0.45 at the 90th percentile. The median MTR is 0.26, with an interquartile range of  $[0.20, 0.36]$ .

These empirical rates are *comprehensive*: they include federal income taxes, employee payroll taxes (Social Security and Medicare), state income taxes, and the phase-out of transfer programs (EITC, CTC, SNAP, Medicaid). This comprehensiveness is precisely the source of a critical scope mismatch with the misperception parameter  $\sigma$ .

**The  $\sigma$  scope mismatch** The calibrated  $\sigma = 0.12$  is drawn from Rees-Jones and Taubinsky [2020], who measured misperception of *federal income tax rates only*. Their survey instrument asked respondents to estimate federal marginal tax rates; it did not cover payroll taxes, state taxes, or benefit phase-outs. However, the marginal tax rates in both the stylized calibration ( $\bar{\tau} = 0.30$ ) and the microsimulation include all of these components.

This creates an asymmetry:  $\tau$  is comprehensive but  $\sigma$  captures only one component of the tax system. Workers who face EITC phase-outs, state income surtaxes, or benefit cliffs confront additional sources of complexity beyond the federal income tax schedule. If the misperception of each additional tax component is even partially independent of federal income tax misperception, then the true comprehensive  $\sigma$  exceeds 0.12. This means the calibrated  $\sigma = 0.12$  may understate the standard deviation of comprehensive MTR misperception, though the magnitude of this understatement depends on the correlation structure of errors across tax components.

The magnitude of this understatement depends on how correlated the various misperception components are. If workers who misperceive their federal rate also misperceive payroll and state taxes in the same direction (correlated errors), the comprehensive  $\sigma$  may be only modestly larger than 0.12. If the errors are largely independent—which is plausible given the distinct administrative origins of each tax—the comprehensive  $\sigma$  could be substantially larger. A rough calculation: if federal income tax misperception contributes  $\sigma_{\text{fed}} = 0.12$  and the combined payroll-plus-state-plus-benefit component contributes an independent  $\sigma_{\text{other}} = 0.06$ , the comprehensive  $\sigma = \sqrt{0.12^2 + 0.06^2} \approx 0.13$ . With a larger non-federal component ( $\sigma_{\text{other}} = 0.10$ ), the comprehensive  $\sigma$  would be approximately 0.16.

For illustrative purposes, if comprehensive  $\sigma$  were 0.16 rather than 0.12, the welfare cost would be 78% larger ( $0.16^2/0.12^2 = 1.78$ ). However, this figure is speculative, as no direct measurement of comprehensive  $\sigma$  exists. The empirical microsimulation results in the next chapter provide a complementary perspective by using actual MTR heterogeneity rather than a single representative rate.

**0.3 Results**

### 0.3.1 Results

This chapter presents the main quantitative findings on the welfare costs of tax rate misperception, drawing on the analytical framework developed earlier and calibrated to U.S. data. All results derive from the closed-form approximation for deadweight loss (DWL) under quasilinear isoelastic preferences:

$$\frac{E[\text{DWL}]}{\text{earnings}} \approx \frac{1}{2} \frac{\varepsilon \sigma^2}{1 - \bar{\tau}} \quad (19)$$

where  $\varepsilon$  is the Frisch elasticity of labor supply,  $\sigma$  is the standard deviation of tax rate misperception, and  $\bar{\tau}$  is the mean marginal tax rate. The formula captures the welfare cost of workers choosing suboptimal hours because they do not know their true marginal rate.

#### Baseline welfare cost

Under the central calibration ( $\varepsilon = 0.33$ ,  $\sigma = 0.12$ ,  $\bar{\tau} = 0.30$ ), the expected deadweight loss is \$187 per worker per year, or \$29.9 billion nationally across 160 million workers. This amounts to 0.11% of GDP.

It is worth emphasizing what this quantity represents and what it does not. The \$187 figure is the expected utility loss, expressed in dollar terms, from a worker who optimizes labor supply against a perceived marginal rate that differs from the true rate by a random draw with standard deviation 0.12. When the perceived rate is too high, the worker supplies too little labor; when it is too low, too much. In both cases the worker ends up on a lower indifference curve than under perfect information. This is a pure behavioral distortion cost – distinct from compliance costs, which represent resources spent on filing and record-keeping rather than misallocated time.

For context, Kopczuk [2013] reviews estimates suggesting the deadweight loss of the federal estate tax is on the order of \$20–30 billion per year. The misperception estimate is similar in aggregate magnitude, though it differs in incidence: misperception-induced DWL is spread across all workers, whereas the estate tax DWL is concentrated among a small number of decedents.

#### Sensitivity analysis

The baseline estimate rests on point values for  $\varepsilon$  and  $\sigma$  that are themselves uncertain. The full sensitivity grid (Table Table 1) spans per-worker DWL from \$63 to \$442, or 0.04% to 0.25% of GDP. DWL is far more sensitive to  $\sigma$  than to  $\varepsilon$ : moving across a row (doubling  $\sigma$  from 0.08 to 0.15) roughly quadruples the cost, whereas moving down a column (doubling  $\varepsilon$  from 0.25 to 0.50) only doubles it. This follows directly from the formula, in which  $\sigma$  enters quadratically and  $\varepsilon$  enters linearly. The most conservative cell (low elasticity and low misperception) implies a \$10 billion annual cost; the most aggressive implies \$71 billion.

#### Distributional pattern

The per-worker DWL formula can be written as a fraction of that worker’s earnings:

$$\frac{\text{DWL}_i}{\text{earnings}_i} = \frac{1}{2} \frac{\varepsilon \sigma_i^2}{1 - \tau_i} \quad (20)$$

Two forces shape the distributional pattern under homogeneous  $\sigma$ . First, the multiplicative interaction of earnings and  $1/(1 - \tau_i)$  concentrates absolute DWL among high earners. Second, the denominator  $(1 - \tau_i)$  means that workers facing higher marginal rates lose a larger *share* of their earnings to misperception, all else equal. The microsimulation results below confirm that the first force dominates: the top income quintile bears the majority of total DWL.

If  $\sigma$  varies across the income distribution—with middle-income workers facing higher misperception due to earned income tax credit (EITC) and Child Tax Credit (CTC) phase-outs, and high-income workers facing lower misperception due to use of tax professionals—the distributional pattern could

shift toward the middle of the distribution. This possibility motivates the heterogeneous- $\sigma$  extension discussed in the conclusion.

### Optimal tax rate under misperception

A utilitarian planner choosing a linear tax rate to maximize social welfare faces a modified tradeoff when workers misperceive their marginal rate. Higher tax rates finance larger transfers and reduce inequality, but they also amplify the welfare cost of misperception because the DWL formula has  $(1 - \tau)$  in the denominator: as  $\tau$  rises, each unit of misperception becomes more costly.

Under the baseline calibration, the optimal tax rate in the absence of misperception is 44.5%. When the planner accounts for  $\sigma = 0.12$  misperception, the optimum falls to 42.9%, a decline of 1.6 percentage points. The planner accepts somewhat less redistribution because the marginal cost of public funds is higher when workers cannot perceive their tax rate accurately.

Within the model, this 1.6 percentage-point reduction corresponds to approximately \$141 billion in annual revenue in the linear-tax framework. This figure reflects the model's tradeoff between redistribution and misperception costs and should not be interpreted as a recommendation for the actual tax code.

### Value of information

Because DWL is quadratic in  $\sigma$ , partial reductions in misperception yield welfare gains that are proportionally larger than the reduction in  $\sigma$  itself. Reducing  $\sigma$  from 0.12 to 0.07 lowers total DWL from \$29.9 billion to \$10.2 billion, a reduction of approximately two-thirds (\$19.7 billion per year). A reduction to  $\sigma = 0.08$  would lower DWL by \$16.6 billion annually. The model is agnostic about *how* such reductions are achieved—better withholding tables, clearer pay-stub reporting, interactive tax tools, or legislative simplification could all contribute.

For comparison, the National Taxpayers Union Foundation estimates that U.S. taxpayers spend approximately \$260 billion annually on tax compliance costs, including both time costs and out-of-pocket expenses [National Taxpayers Union Foundation, 2023]. The misperception-induced DWL documented here is an *additional* cost, distinct from compliance costs: it arises not from time spent filing but from the residual errors that persist even after filing.

### Microsimulation results

The stylized calculations above use a single representative marginal tax rate ( $\bar{\tau} = 0.30$ ) and a log-normal wage distribution. To assess the distributional incidence of misperception-induced DWL across the actual income distribution, I apply the per-worker DWL formula to household-level data from PolicyEngine-US [PolicyEngine, 2024], a microsimulation model that computes comprehensive marginal tax rates for approximately 149 million working-age adults with positive employment income.

**Marginal tax rate distribution** The weighted mean comprehensive marginal tax rate (MTR) is 0.25, with a standard deviation of 0.16. The distribution is right-skewed: the median (0.26) is close to the mean, but the 90th percentile (0.45) is far above it, reflecting the steep statutory rates faced by high earners. At the bottom, 10% of workers face an MTR at or near zero, typically because they earn below the payroll tax threshold or receive offsetting credits.

**DWL by income quintile** Applying the per-worker DWL formula  $\frac{1}{2}\varepsilon \cdot \text{earnings}_i \cdot \sigma^2 / (1 - \tau_i)$  with  $\varepsilon = 0.33$  and  $\sigma = 0.12$  to each worker in the microsimulation yields the distributional breakdown in Table 2.

These results assume homogeneous  $\sigma = 0.12$  across all quintiles. If misperception varies with income—as is likely, given that lower-income workers navigate denser phase-out regions while higher-income workers more frequently use professional tax preparers—the distributional pattern could shift. I explore heterogeneous  $\sigma_i$  as a direction for future work in the conclusion.

Table 2: Deadweight loss from tax misperception by income quintile

Quintile	Mean earnings	Mean MTR	Per-worker DWL	Share of total
1 (lowest)	\$7,355	0.11	\$20	1.8%
2	\$23,550	0.23	\$73	6.8%
3	\$40,742	0.27	\$132	12.3%
4	\$64,301	0.31	\$220	20.1%
5 (highest)	\$171,542	0.34	\$614	58.9%

Three findings stand out. First, the top quintile bears 59% of total DWL despite constituting only 20% of workers. This reflects the multiplicative interaction of higher earnings and higher marginal rates in the DWL formula: workers earning \$172k at a 34% MTR lose \$614/year to misperception, compared to \$20/year for workers earning \$7k at an 11% MTR.

Second, as a *fraction of earnings*, the distributional pattern is more uniform. The DWL-to-earnings ratio is  $\frac{1}{2}\varepsilon\sigma^2/(1-\tau_i)$ , which depends only on  $\varepsilon$ ,  $\sigma$ , and the individual’s MTR. Workers in the top quintile lose 0.36% of earnings, versus 0.27% for bottom-quintile workers—a narrower gap than the absolute dollar figures suggest. The  $(1-\tau)$  denominator drives this mild progressivity: higher-MTR workers face a larger amplification of the same misperception variance.

Third, the aggregate DWL from the microsimulation (\$36.6 billion) exceeds the comparable stylized estimate (\$27.8 billion) by 32%. The stylized comparison applies the DWL formula at the microsimulation population’s weighted-mean earnings and MTR (rather than the macro-calibrated \$55,000 and  $\tau = 0.30$  used in the sensitivity grid, which yield \$29.9 billion over 160 million workers). This divergence arises because the DWL formula is convex in both earnings and  $1/(1-\tau)$ : the actual distribution of incomes and tax rates has heavier tails than the representative-agent approximation, and Jensen’s inequality ensures that the population-level DWL exceeds the DWL evaluated at population means.

In sum, the welfare costs of tax rate misperception range from 0.04% to 0.25% of GDP under every calibration considered, and the microsimulation analysis yields an aggregate estimate of \$37 billion—32% above the stylized baseline. The costs are concentrated in absolute terms among higher earners but are moderately progressive as a share of earnings. Within the optimal tax framework, accounting for misperception reduces the optimal linear tax rate by 1.6 percentage points.

### 0.3.2 Extensions and robustness

The baseline results rely on quasilinear isoelastic preferences and a static, representative-agent calibration. This chapter examines how the findings change under alternative preference specifications, dynamic considerations, richer heterogeneity, and endogenous information acquisition. Together, these extensions bracket the plausible range of welfare costs and clarify where the baseline estimates sit within that range.

#### Cobb-Douglas benchmark

A natural starting point is the Cobb-Douglas utility function  $U = L^\alpha C^\beta$ , which the earliest version of this project employed. Under Cobb-Douglas preferences with no government transfers, the optimal labor supply is

$$h^* = \frac{\beta T}{\alpha + \beta} \quad (21)$$

where  $T$  is the time endowment. This expression contains no tax rate: the income effect of a higher tax (which makes leisure cheaper in opportunity-cost terms) exactly offsets the substitution effect (which makes leisure more attractive). Because the worker supplies the same hours regardless of the perceived tax rate, misperception has no effect on behavior, and the DWL from uncertainty is exactly zero.

This is a knife-edge result. It holds only when (i) preferences are exactly Cobb-Douglas and (ii) the government rebates no revenue. Once the government finances a lump-sum transfer  $G$  from tax revenue, the budget constraint becomes  $C = w(1 - \tau)(T - L) + G$ , and the tax rate re-enters the labor supply decision through the transfer channel. Even under Cobb-Douglas preferences, the worker now adjusts hours in response to perceived  $\tau$  because the transfer makes the income and substitution effects no longer cancel. The resulting DWL is positive but smaller than under quasilinear preferences, because the substitution-effect channel – the dominant channel empirically – is shut down by assumption.

The quasilinear model captures the substitution-effect channel directly: workers unambiguously reduce labor supply when they perceive a higher marginal rate. Empirical evidence from Chetty et al. [2009] and Feldstein [1999] consistently finds that labor supply responds to perceived net-of-tax rates on the substitution margin, confirming that the quasilinear specification is the more empirically relevant of the two.

#### Constant elasticity of substitution (CES) generalization

Both Cobb-Douglas and quasilinear preferences are special cases of the CES family

$$U(L, C) = [\alpha L^\rho + (1 - \alpha)C^\rho]^{1/\rho} \quad (22)$$

with elasticity of substitution  $\sigma_u = 1/(1 - \rho)$  (where  $\sigma_u$  denotes the elasticity of substitution between leisure and consumption, distinct from the misperception standard deviation  $\sigma$ ). Cobb-Douglas corresponds to  $\sigma_u = 1$  (where income and substitution effects cancel), and the quasilinear specification captures the limit in which the substitution effect dominates the income effect on the relevant margin. As  $\sigma_u$  rises above one, the substitution effect increasingly outweighs the income effect, and the DWL from misperception grows.

The quasilinear calibration therefore provides an estimate toward the upper end of the CES family for a given Frisch elasticity. Conversely, preferences with  $\sigma_u < 1$  would produce smaller welfare costs. The empirical labor supply literature, which consistently finds that compensated elasticities exceed uncompensated elasticities, implies  $\sigma_u > 1$  for most workers, suggesting that  $\sigma_u > 1$  is the empirically relevant range. Whether the quasilinear estimates are above or below a more general CES specification depends on the exact value of  $\sigma_u$ , which varies across populations and estimation methods.

## Dynamic considerations

The static framework ignores at least four channels through which tax misperception could affect welfare beyond the intensive-margin adjustment modeled above.

First, savings and investment decisions depend on perceived after-tax returns. A worker who overestimates her marginal rate may also overestimate the tax on capital income, leading to under-saving. This adds a DWL term on the capital margin that is absent from the baseline model.

Second, human capital investment – decisions about education, training, and occupational choice – responds to perceived lifetime tax rates. If a prospective medical student overestimates the marginal rate on physician earnings, she may choose a lower-return career, generating a misallocation cost that compounds over decades.

Third, life-cycle labor supply involves timing decisions (when to enter the labor force, when to retire) that depend on perceived tax rates at different ages. Misperception of the Social Security earnings test or Medicare surtax, for example, could distort retirement timing.

Fourth, the model focuses on continuous hours adjustment, whereas most workers face fixed-hours contracts. The extensive margin — the decision of whether to work at all — may be more responsive to perceived tax rates than the intensive margin for many workers. A worker who overestimates her marginal rate by 10 percentage points may not reduce her weekly hours (which are contractually fixed) but may delay entry into the labor force, choose part-time over full-time work, or retire earlier. These discrete decisions are governed by participation elasticities, which Chetty [2012] estimates at approximately 0.25 on the extensive margin. The calibrated Frisch elasticity of 0.33, which blends intensive and extensive margin responses, implicitly captures some of this channel.

These channels could increase or decrease the baseline welfare cost. Savings and human capital distortions would add to the cost, while within-year learning about true rates could reduce it. The net direction is an empirical question that the static framework cannot resolve.

## Heterogeneous misperception

The baseline model assumes a uniform  $\sigma = 0.12$  for all workers. In practice, misperception almost certainly varies across the income distribution.

Low-income workers face relatively simple statutory schedules, but limited financial literacy and less experience with tax filing may raise their subjective uncertainty. The EITC phase-in and phase-out create steep effective marginal rates that many recipients do not understand, as documented by Chetty and Saez [2013].

Middle-income workers face the largest number of interacting provisions: CTC phase-outs, the EITC cliff, the transition between federal brackets, the standard-versus-itemized deduction choice, AMT exposure, and state income taxes that piggyback on federal definitions. For these households,  $\sigma$  may differ from the population average of 0.12, though direct measurement is lacking.

High-income workers face complex capital-gains rules, the net investment income tax, and state-specific provisions, but they also disproportionately hire certified public accountants (CPAs) and financial advisors. Professional tax planning likely compresses  $\sigma$  below the population mean for this group.

Incorporating a  $\sigma$  that rises and then falls with income would shift the aggregate welfare cost toward middle-income workers and could raise or lower the population total depending on the covariance between  $\sigma_i$  and  $\tau_i$ .

## Endogenous information acquisition

Workers are not passive recipients of tax confusion; they can invest in learning their marginal rate by hiring accountants, purchasing tax software, or spending time reading Internal Revenue Service (IRS) publications. The observed  $\sigma$  is therefore an equilibrium outcome in which each worker balances the cost of information against the utility gain from more accurate optimization.

The DWL formula implies that the private benefit of reducing one's own  $\sigma$  by a small amount  $d\sigma$  is approximately  $\varepsilon \cdot \sigma \cdot \text{earnings} \cdot d\sigma / (1 - \tau)$ . Workers with higher earnings, higher  $\tau$ , or higher  $\varepsilon$  gain

more from information and should invest more in acquiring it. This is consistent with the observation that high-income households are far more likely to hire professional tax preparers.

The National Taxpayers Union Foundation estimates that U.S. taxpayers spend approximately \$260 billion annually on tax compliance, including both time costs and out-of-pocket expenses [National Taxpayers Union Foundation, 2023]. This expenditure represents the private sector’s partial response to the misperception problem. The scale of compliance spending is consistent with the hypothesis that tax complexity imposes costs, though compliance spending serves multiple purposes beyond reducing marginal-rate misperception.

### Comparison to prior estimates

Skinner [1988] estimated the welfare cost of tax *policy* uncertainty – uncertainty about future legislative changes – at 0.4% of national income per year. That concept is distinct from the misperception studied here: Skinner’s agents know the current tax schedule perfectly but face uncertainty about what it will be next year. Whether the two sources of welfare cost are additive depends on modeling assumptions not addressed here (e.g., whether risk aversion amplifies or dampens the misperception channel). Direct comparison of the magnitudes is illustrative but should not be interpreted as a rigorous aggregation.

Chetty et al. [2009] find that salience manipulations shift consumer behavior by amounts equivalent to 1–4% of the tax revenue at stake. Translating their experimental magnitudes into a population-wide DWL is not straightforward, but the orders of magnitude are compatible: if misperception shifts effective labor supply responses by a few percent of the tax wedge, the resulting DWL as a share of GDP falls squarely in the estimated range.

Taken together, the extensions and comparisons provide context for the baseline estimate of 0.11% of GDP, which varies between 0.04% and 0.25% across the parameter grid. Some omitted channels (dynamic distortions, heterogeneous misperception) would increase the estimate, while others (within-year learning, inattention rather than misperception) could decrease it.

**0.4 Conclusion**

### 0.4.1 Conclusion

This paper has shown that the inability of workers to accurately perceive their marginal tax rates imposes a welfare cost on the order of 0.1% of GDP. Under quasilinear-isoelastic preferences, the expected deadweight loss per worker takes the closed form  $\frac{1}{2}\varepsilon\sigma^2/(1-\tau)$  times earnings. Calibrated to a Frisch elasticity of 0.33, a misperception standard deviation of 12 percentage points ( $\sigma = 0.12$ ), and a mean marginal rate of 0.30, the stylized central estimate is \$187 per worker and \$30 billion in aggregate—0.11% of GDP annually. Applying the formula to household-level marginal tax rates from PolicyEngine-US microsimulation yields a higher aggregate of \$37 billion, as the convexity of the DWL formula amplifies losses from the heavy tails of the income and tax rate distributions. The top income quintile bears 59% of total DWL, while the bottom quintile bears less than 2%. The full sensitivity range spans 0.04% to 0.25% of GDP (\$10 to \$71 billion), reflecting plausible variation in labor supply elasticities and misperception magnitudes. These losses arise not from future policy uncertainty but from the current opacity of the tax code.

### Policy implications

The quadratic dependence on  $\sigma$  means that proportional reductions in misperception variance have a larger effect on DWL than proportional reductions in the tax rate, holding all else equal. A 5 percentage point reduction in the misperception standard deviation (from  $\sigma = 0.12$  to  $\sigma = 0.07$ ) lowers  $\sigma^2$  by 66%, cutting per-worker deadweight loss by roughly two-thirds. By contrast, a 3 percentage point reduction in the statutory rate (from  $\tau = 0.30$  to  $\tau = 0.27$ ) reduces the  $1/(1-\tau)$  term by only about 4%. The relative cost-effectiveness of these approaches depends on implementation costs that are outside the scope of this model. The model is agnostic about *how* misperception is reduced. Several channels could lower  $\sigma$ , each with different costs and tradeoffs:

- **Better information provision:** More accurate withholding tables, clearer pay-stub reporting of marginal rates, or interactive tools that show workers their true rate could reduce  $\sigma$  without any change to the tax code itself.
- **Tax simplification:** Consolidating brackets, phase-outs, and overlapping credits would make the rate schedule more legible, but provisions typically exist for distributional or incentive reasons. The net welfare gain from simplification depends on what is lost in the process.
- **Professional tax preparation:** Wider access to competent preparation could reduce misperception, though the compliance cost literature suggests current preparation spending is already substantial [National Taxpayers Union Foundation, 2023].

The model does not rank these channels; it quantifies the welfare cost of a given  $\sigma$  regardless of its source. Because DWL is quadratic in  $\sigma$ , a reduction from 0.12 to 0.10 lowers DWL by 31%, and a reduction to 0.07 lowers it by 66%.

The optimal tax analysis reinforces this point from a different angle. A utilitarian planner who accounts for misperception chooses a lower optimal linear tax rate of 42.9%, compared with 44.5% under perfect information. The gap of 1.6 percentage points reflects the additional efficiency cost that misperception imposes at the margin: each dollar of revenue raised creates more deadweight loss when workers cannot accurately perceive the rate they face. Within the model, lower  $\sigma$  shifts the optimal linear tax rate upward, allowing more redistribution at a given level of total deadweight loss.

### Limitations

Several caveats apply. The model is static: workers choose labor supply once, facing a single marginal rate. In reality, labor supply decisions unfold over time, and workers may learn about their true rate as the year progresses. This suggests the static model may overstate the cost for workers with stable employment but understate it for those making discrete labor market transitions (entering or leaving the workforce, choosing between jobs with different hours).

The analysis assumes a linear tax schedule with a single marginal rate. The actual U.S. tax code is piecewise linear with multiple brackets, phase-outs, and cliffs. Misperception of a non-linear schedule

may differ qualitatively from misperception of a single rate, as workers may be uncertain not only about the level of their marginal rate but also about where bracket thresholds fall.

I have assumed that the misperception error  $\delta$  is normally distributed with mean zero and constant variance across the population. In practice, misperception may be correlated with income (lower-income workers may face more complex effective schedules due to benefit phase-outs), with financial sophistication, and with access to tax preparation services. The model treats  $\sigma$  as a population-level parameter and does not capture this heterogeneity. Moreover, the calibrated  $\sigma = 0.12$  is drawn from Rees-Jones and Taubinsky [2020], who measured misperception of federal income tax rates only. Since the comprehensive marginal tax rate includes payroll taxes, state taxes, and benefit phase-outs, the true comprehensive  $\sigma$  may differ from 0.12. If errors across tax components are partially independent, the comprehensive  $\sigma$  would exceed 0.12; if workers who misperceive federal rates also misperceive other components in the same direction, the difference would be smaller.

Finally, the analysis focuses exclusively on labor supply and ignores other margins of response—savings, portfolio allocation, tax avoidance, and organizational form—that may also be distorted by rate misperception. These omitted margins could increase or decrease the total welfare cost depending on the sign and magnitude of the distortions they introduce.

### Future work

Two extensions would strengthen and generalize these results. First, while this paper applies the DWL formula to heterogeneous  $\tau_i$  via PolicyEngine microsimulation, it maintains homogeneous  $\sigma$ . The assumption of uniform misperception can be relaxed by linking misperception variance to observable household characteristics (income, filing status, use of tax preparers) using the survey data from Gideon [2017] and Rees-Jones and Taubinsky [2020]. Heterogeneous  $\sigma_i$  would sharpen the aggregate estimates and identify which populations and provisions contribute most to misperception-induced welfare loss.

Second, measuring comprehensive misperception—not just federal income tax misperception—is essential for accurate welfare accounting. One approach would be to replicate the Rees-Jones & Taubinsky survey instrument but expand it to ask about total marginal rates (including payroll, state, and benefit phase-outs). An alternative approach would use large language models as a benchmark: presenting realistic household scenarios to AI systems and comparing their estimated MTRs to true rates from PolicyEngine could validate the structural misperception hypothesis while isolating the contribution of tax system complexity from individual cognitive limitations.

**0.5 Appendix**

### 0.5.1 Technical appendix

#### Derivation of optimal labor supply

Consider a worker with quasilinear isoelastic preferences:

$$U(C, h) = C - \frac{\psi h^{1+1/\varepsilon}}{1 + 1/\varepsilon} \quad (23)$$

subject to the budget constraint  $C = w(1 - \tau)h + v$ , where  $w$  is the wage,  $\tau$  is the marginal tax rate,  $h$  is hours worked, and  $v$  is a lump-sum transfer.

Substituting the budget constraint:

$$U = w(1 - \tau)h + v - \frac{\psi h^{1+1/\varepsilon}}{1 + 1/\varepsilon} \quad (24)$$

The first-order condition with respect to  $h$ :

$$U_h = w(1 - \tau) - \psi h^{1/\varepsilon} = 0 \quad (25)$$

Solving for  $h$ :

$$h^*(\tau) = \left( \frac{w(1 - \tau)}{\psi} \right)^\varepsilon \quad (26)$$

This confirms that labor supply is increasing in the net-of-tax wage  $w(1 - \tau)$  with constant elasticity  $\varepsilon$ .

#### Proof of the DWL formula

**Setup** A worker perceives  $\hat{\tau} = \tau + \delta$  where  $\delta$  is the misperception error. The worker chooses labor to maximize utility under the perceived rate:

$$h(\hat{\tau}) = \left( \frac{w(1 - \hat{\tau})}{\psi} \right)^\varepsilon \quad (27)$$

But utility is realized at the true rate  $\tau$ :

$$U(\tau, h(\hat{\tau})) = w(1 - \tau)h(\hat{\tau}) + v - \frac{\psi h(\hat{\tau})^{1+1/\varepsilon}}{1 + 1/\varepsilon} \quad (28)$$

**Second-order Taylor expansion** Define  $DWL(\delta) = U(\tau, h^*(\tau)) - U(\tau, h(\hat{\tau}))$ . Since  $h^*$  maximizes  $U(\tau, \cdot)$ , the first-order term vanishes by the envelope theorem, and:

$$DWL(\delta) \approx -\frac{1}{2} U_{hh} \Big|_{h=h^*} \cdot (\Delta h)^2 \quad (29)$$

Computing  $U_{hh}$ :

$$U_{hh} = -\frac{\psi}{\varepsilon} h^{1/\varepsilon - 1} \quad (30)$$

At the optimum,  $\psi(h^*)^{1/\varepsilon} = w(1 - \tau)$ , so  $\psi = w(1 - \tau) \cdot (h^*)^{-1/\varepsilon}$ . Substituting:

$$U_{hh} \Big|_{h^*} = -\frac{w(1 - \tau)}{\varepsilon h^*} \quad (31)$$

Computing  $\Delta h$ :

$$\Delta h = h(\hat{\tau}) - h^*(\tau) = \left( \frac{w(1 - \tau - \delta)}{\psi} \right)^\varepsilon - \left( \frac{w(1 - \tau)}{\psi} \right)^\varepsilon \quad (32)$$

For small  $\delta$ , linearizing:

$$\Delta h \approx -\frac{\varepsilon h^*}{1-\tau} \cdot \delta \tag{33}$$

Combining:

$$\text{DWL}(\delta) \approx \frac{1}{2} \cdot \frac{w(1-\tau)}{\varepsilon h^*} \cdot \left(\frac{\varepsilon h^*}{1-\tau}\right)^2 \delta^2 = \frac{1}{2} \varepsilon w h^* \cdot \frac{\delta^2}{1-\tau} \tag{34}$$

**Expected DWL** If  $\delta \sim N(0, \sigma^2)$ , then  $E[\delta^2] = \sigma^2$  and:

$$E[\text{DWL}] \approx \frac{1}{2} \varepsilon w h^* \cdot \frac{\sigma^2}{1-\tau} \tag{35}$$

Dividing by earnings  $wh^*$ :

$$\frac{E[\text{DWL}]}{\text{earnings}} \approx \frac{1}{2} \varepsilon \cdot \frac{\sigma^2}{1-\tau} \tag{36}$$

### Comparative statics

The formula  $E[\text{DWL}]/\text{earnings} = \frac{1}{2} \varepsilon \sigma^2 / (1-\tau)$  yields immediate comparative statics:

Parameter	Effect on DWL	Intuition
$\varepsilon$ (Frisch elasticity)	Linear, positive	More elastic workers make larger errors
$\sigma$ (misperception std dev)	Quadratic, positive	Larger errors cause disproportionately larger losses
$\tau$ (marginal rate)	Positive via $1/(1-\tau)$	Higher rates amplify errors in the net-of-tax wage
$wh^*$ (earnings)	Linear, positive	Higher earners lose more in absolute terms

### Optimal tax under misperception

Consider a utilitarian planner choosing  $\tau$  to maximize average welfare with a balanced-budget demogrant  $v = \tau \bar{w} \bar{h} / N$ . The planner's problem under perfect information yields the standard equity-efficiency tradeoff. Under misperception, the effective elasticity of the DWL with respect to  $\tau$  increases, since:

$$\frac{\partial E[\text{DWL}]}{\partial \tau} > 0 \tag{37}$$

The planner internalizes that a marginal increase in  $\tau$  raises the DWL from misperception (through the  $1/(1-\tau)$  term), leading to:

$$\frac{\partial \tau^*}{\partial \sigma} < 0 \tag{38}$$

The optimal tax rate decreases with misperception noise. This is confirmed numerically: the calibration shows  $\tau^*$  falls from 44.5% to 42.9% when  $\sigma$  increases from 0 to 0.12.

**Cobb-Douglas knife-edge result**

For Cobb-Douglas utility  $U(L, C) = L^\alpha C^\beta$  with leisure  $L = T - h$  and no transfers ( $v = 0$ ), the optimal leisure is:

$$L^* = \frac{\alpha T}{\alpha + \beta} \quad (39)$$

This is independent of the tax rate. The income and substitution effects of a tax change exactly cancel, so misperceiving  $\tau$  has no effect on labor supply and hence zero DWL. This knife-edge property motivates the use of quasilinear preferences, which isolate the substitution effect — the empirically relevant channel for tax rate misperception.

# Bibliography

- J. Abeler and S. Jäger. Complex tax incentives. *American Economic Journal: Economic Policy*, 7(3): 1–28, 2015. doi: 10.1257/pol.20130137. URL <https://doi.org/10.1257/pol.20130137>.
- J. Alm. Uncertain tax policies, individual behavior, and welfare. *American Economic Review*, 78(1): 237–245, 1988. URL <https://www.jstor.org/stable/1814710>.
- S. R. Baker, N. Bloom, and S. J. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016. doi: 10.1093/qje/qjw024. URL <https://doi.org/10.1093/qje/qjw024>. Working paper and data available at [https://www.policyuncertainty.com/media/EPU\\_BBD\\_Mar2016.pdf](https://www.policyuncertainty.com/media/EPU_BBD_Mar2016.pdf).
- C. L. Ballard and S. Gupta. Perceptions and realities of average tax rates in the federal income tax: Evidence from Michigan. *National Tax Journal*, 71(2):263–294, 2018. doi: 10.17310/ntj.2018.2.03. URL <https://doi.org/10.17310/ntj.2018.2.03>.
- Y. Benzarti. How taxing is tax filing? Using revealed preferences to estimate compliance costs. *American Economic Journal: Economic Policy*, 12(4):38–57, 2020. doi: 10.1257/pol.20180664. URL <https://doi.org/10.1257/pol.20180664>. Working paper available at <https://www.nber.org/papers/w23903>.
- M. K. Block, J. M. Kramer, and J. M. Wagner. Labor supply under wage uncertainty. *Economic Inquiry*, 18(2):261–269, 1980. doi: 10.1111/j.1465-7295.1980.tb00573.x. URL <https://doi.org/10.1111/j.1465-7295.1980.tb00573.x>.
- Bureau of Economic Analysis. Gross Domestic Product, 2024. URL <https://www.bea.gov/data/gdp/gross-domestic-product>.
- Bureau of Labor Statistics. Occupational Employment and Wage Statistics, 2023. URL <https://www.bls.gov/oes/>.
- Bureau of Labor Statistics. Labor Force Statistics from the Current Population Survey, 2024. URL <https://www.bls.gov/cps/>.
- R. Chetty. Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*, 80(3):969–1018, 2012. doi: 10.3982/ECTA9043. URL <https://doi.org/10.3982/ECTA9043>. Working paper available at <https://www.nber.org/papers/w15616>.
- R. Chetty and E. Saez. Teaching the tax code: Earnings responses to an experiment with EITC recipients. *American Economic Journal: Applied Economics*, 5(1):1–31, 2013. doi: 10.1257/app.5.1.1. URL <https://doi.org/10.1257/app.5.1.1>. Working paper available at <https://www.nber.org/papers/w14836>.
- R. Chetty, A. Looney, and K. Kroft. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009. doi: 10.1257/aer.99.4.1145. URL <https://doi.org/10.1257/aer.99.4.1145>. Working paper available at <https://www.nber.org/papers/w13330>.
- Congressional Budget Office. Effective Marginal Tax Rates for Low- and Moderate-Income Workers. Report, Congressional Budget Office, 2012. URL <https://www.cbo.gov/publication/43709>. Full

report available at <https://www.cbo.gov/sites/default/files/112th-congress-2011-2012/reports/11-15-2012MarginalTaxRates.pdf>.

Congressional Budget Office. Effective Marginal Tax Rates for Low- and Moderate-Income Workers in 2016. Report, Congressional Budget Office, 2016. URL <https://www.cbo.gov/publication/51923>. Full report available at <https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51923-MarginalTaxRates.pdf>.

P. Diamond and E. Saez. The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives*, 25(4):165–190, 2011. doi: 10.1257/jep.25.4.165. URL <https://doi.org/10.1257/jep.25.4.165>.

E. Duflo and E. Saez. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment. *Quarterly Journal of Economics*, 118(3):815–842, 2003. doi: 10.1162/00335530360698432. URL <https://doi.org/10.1162/00335530360698432>. Working paper available at <https://www.nber.org/papers/w8885>.

E. Farhi and X. Gabaix. Optimal taxation with behavioral agents. *American Economic Review*, 110(1):298–336, 2020. doi: 10.1257/aer.20151079. URL <https://doi.org/10.1257/aer.20151079>. Working paper available at <https://www.nber.org/papers/w21524>.

N. E. Feldman, P. Katuščák, and L. Kawano. Taxpayer confusion: Evidence from the Child Tax Credit. *American Economic Review*, 106(3):807–835, 2016. doi: 10.1257/aer.20131189. URL <https://doi.org/10.1257/aer.20131189>.

M. Feldstein. Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics*, 81(4):674–680, 1999. doi: 10.1162/003465399558391. URL <https://doi.org/10.1162/003465399558391>. Working paper available at <https://www.nber.org/papers/w5055>.

A. Finkelstein. E-ZTAX: Tax salience and tax rates. *Quarterly Journal of Economics*, 124(3):969–1010, 2009. doi: 10.1162/qjec.2009.124.3.969. URL <https://doi.org/10.1162/qjec.2009.124.3.969>.

A. Gerritsen. Optimal taxation when people do not maximize well-being. *Journal of Public Economics*, 144:122–139, 2016. doi: 10.1016/j.jpubeco.2016.10.006. URL <https://doi.org/10.1016/j.jpubeco.2016.10.006>.

M. Gideon. Do individuals perceive income tax rates correctly? *Public Finance Review*, 45(1):97–117, 2017. doi: 10.1177/1091142115615670. URL <https://doi.org/10.1177/1091142115615670>.

J. Goldin. Tax benefit complexity and take-up: Lessons from the earned income tax credit. *Tax Law Review*, 72:59–110, 2018. URL [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2711869](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2711869). Open access version available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2711869](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2711869).

K. A. Hassett and G. E. Metcalf. Investment with uncertain tax policy: Does random tax policy discourage investment? *Economic Journal*, 109(457):372–393, 1999. doi: 10.1111/1468-0297.00453. URL <https://doi.org/10.1111/1468-0297.00453>. Working paper available at <https://www.nber.org/papers/w4780>.

J. S. Hastings and J. M. Weinstein. Information, school choice, and academic achievement: Evidence from two experiments. *Quarterly Journal of Economics*, 123(4):1373–1414, 2008. doi: 10.1162/qjec.2008.123.4.1373. URL <https://doi.org/10.1162/qjec.2008.123.4.1373>. Working paper available at <https://www.nber.org/papers/w13623>.

R. Jensen. The (perceived) returns to education and the demand for schooling. *Quarterly Journal of Economics*, 125(2):515–548, 2010. doi: 10.1162/qjec.2010.125.2.515. URL <https://doi.org/10.1162/qjec.2010.125.2.515>.

- M. Keane and R. Rogerson. Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature*, 50(2):464–476, 2012. doi: 10.1257/jel.50.2.464. URL <https://doi.org/10.1257/jel.50.2.464>.
- W. Kopczuk. *Taxation of Intergenerational Transfers and Wealth*, volume 5, pages 329–390. Elsevier, 2013. doi: 10.1016/B978-0-444-53759-1.00006-6.
- J. B. Liebman and R. J. Zeckhauser. Schmeduling. Harvard University Working Paper. Available at [https://scholar.harvard.edu/files/jeffreyliebman/files/Schmeduling\\_WorkingPaper.pdf](https://scholar.harvard.edu/files/jeffreyliebman/files/Schmeduling_WorkingPaper.pdf), 2004.
- J. A. Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(2):175–208, 1971. doi: 10.2307/2296779. URL <https://doi.org/10.2307/2296779>.
- National Taxpayers Union Foundation. Complexity 2023: 6.5 Billion Hours, \$260 Billion: What Tax Complexity Costs Americans. Technical report, National Taxpayers Union Foundation, 2023. URL <https://www.ntu.org/foundation/tax-page/complexity-2023-65-billion-hours-260-billion-what-tax-complexity-costs-americans>.
- PolicyEngine. Policyengine US: Open-source tax-benefit microsimulation model. <https://github.com/PolicyEngine/policyengine-us>, 2024. URL <https://github.com/PolicyEngine/policyengine-us>. Version 1.45.1. Available at <https://policyengine.org>.
- A. Rees-Jones and D. Taubinsky. Measuring “Schmeduling”. *Review of Economic Studies*, 87(5): 2399–2438, 2020. doi: 10.1093/restud/rdz045. URL <https://doi.org/10.1093/restud/rdz045>. Working paper available at <https://www.nber.org/papers/w22884>.
- E. Saez. Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, 68(1):205–229, 2001. doi: 10.1111/1467-937X.00166. URL <https://doi.org/10.1111/1467-937X.00166>.
- E. Saez, J. Slemrod, and S. H. Giertz. The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature*, 50(1):3–50, 2012. doi: 10.1257/jel.50.1.3. URL <https://doi.org/10.1257/jel.50.1.3>. Working paper available at <https://www.nber.org/papers/w15012>.
- J. Skinner. The welfare cost of uncertain tax policy. *Journal of Public Economics*, 37(2):129–145, 1988. doi: 10.1016/0047-2727(88)90067-3. URL [https://doi.org/10.1016/0047-2727\(88\)90067-3](https://doi.org/10.1016/0047-2727(88)90067-3). Working paper available at <https://www.nber.org/papers/w2173>.
- J. Slemrod. The etiology of tax complexity: Evidence from U.S. state income tax systems. *Public Finance Review*, 33(3):279–299, 2005. doi: 10.1177/1091142105275003. URL <https://doi.org/10.1177/1091142105275003>.