Optimal Tax Information With Privacy Concerns*

Joel Slemrod Damián Vergara

Abstract

Having more information about taxpayers facilitates tax enforcement, but generates utility losses due to privacy concerns. We posit that individuals put negative value on two distinct aspects of privacy: revealing their level of affluence, and revealing details of their income conditional on affluence. We characterize the enforcement-privacy tradeoff for optimal income tax policy by introducing the granularity of the tax base as a key policy variable. The first-order conditions reveal the sufficient statistics needed to describe the optimal tax rate structure and the extent of granularity. We discuss how future empirical research could improve the characterization of optimal information.

1 Introduction

It is well known that governments benefit from taxpayers' information because it facilitates tax enforcement (Slemrod, 2019). Examples include the use of information to target tax audits and the advantages of third-party reporting. Less is known about the privacy costs that individuals incur when providing information to governments and their policy implications. Slemrod (2025a) surveys the literature and argues that taxpayers value privacy, limiting the social value of providing information to tax authorities. These privacy costs arise from taxpayers' desire to avoid disclosing their affluence, and to avoid revealing income or expenditure details (conditional on affluence) to the tax authority and, potentially via leaks, to the general public. This issue is likely to become more important in the future, given technological advances that facilitate access to personal information and growing concerns that some governments worldwide may use tax information against individuals for reasons other than tax compliance.

What are the implications of these private privacy costs for optimal tax policy? This paper proposes a simple framework for studying the tradeoffs embedded in the optimal amount of information that tax authorities should gather from taxpayers. We focus on the granularity of required income reporting as a novel policy parameter: providing more details in a given income report is helpful for tax enforcement, but generates larger privacy costs. The framework is purposely stylized to inform high-level reflections on

*First version: June 2025. This version: June 2025. Slemrod: University of Michigan, jslemrod@umich.edu. Vergara: University of Michigan, damiv@umich.edu. Usual disclaimers apply.

the optimal use of taxpayers' information, hopefully guiding empirical research to quantify this tradeoff and serving as a building block for more general theoretical analyses of tax privacy.

In the model, individuals make standard labor supply and reporting decisions given tax instruments. New to the analysis is the inclusion of two privacy costs. People experience *revealed affluence costs* because individuals dislike being revealed by the government, and potentially the public at large, as affluent. People also experience *granularity costs* because, conditional on affluence, individuals dislike it when the government knows details about their income composition, such as the income sources, the firms they work for, or the assets they invest in. Privacy costs generate mechanical labor supply distortions as they work as implicit income taxes. This applies to both real and reporting decisions: to hide the true affluence, individuals may engage in income underreporting even in the absence of income taxes.

We then model the decision of a government with a generalized utilitarian objective that chooses the labor income tax, the granularity of the tax system, and a universal lump-sum transfer to maximize social welfare. We first show that the optimal labor income tax rate resembles the standard formula in optimal linear income taxation, where, in a sufficient statistics spirit, the optimal tax only depends on welfare weights and the elasticity of taxable income (ETI) with respect to the net-of-tax rate (Feldstein, 1999; Piketty and Saez, 2013). However, two subtleties arise in this result. First, privacy costs implicitly affect the ETI as they alter effective tax rates from the taxpayer's perspective. Second, as in Slemrod and Kopczuk (2002), the ETI is a function of the granularity parameter, as a more granular tax system makes evasion more difficult, dampening the efficiency cost of income taxation.

We also characterize the optimal granularity of the tax system, which balances two opposing forces. On one hand, more granularity generates a positive fiscal externality: by increasing the cost of evasion, more granularity increases revenue in proportion to the income tax rate and the elasticity of taxable income with respect to the system's granularity, a novel reduced-form sufficient statistic to measure the welfare benefits from information provision. On the other hand, more granularity generates utility costs for taxpayers, both in terms of evasion costs and, more directly, due to privacy concerns. These welfare costs are proportional to individual social welfare weights, implying that privacy concerns will deter the government from gathering additional information especially when privacy costs are disproportionately borne by high-welfare-weight (i.e., low well-being) individuals.

As in Keen and Slemrod (2017), the tax rate and the granularity of the tax system are strategic complements. A higher tax rate mechanically increases the revenue gains from increased enforcement due to more granularity. Likewise, a higher degree of granularity decreases the ETI, making the income tax less distortionary. This insight highlights the importance of considering the joint decision of taxes and information requirements. Focusing on a single instrument may generate welfare losses for governments.

We end the paper by noting that the sufficient statistics required to calibrate the optimal granularity have not yet been estimated in the related literature. We provide a brief discussion of related empirical work that suggests these sufficient statistics are quantitatively important. However, further research is needed for converting this body of evidence into the relevant elasticities for the normative analysis.

2 Model

Individuals have quasi-linear utilities in consumption and are heterogeneous in earnings potential θ. For simplicity, assume that θ has two mass points, $\theta \in \{0, 1\}$, with exogenous shares s_0 and s_1.

Individuals of type $\theta = 0$ cannot work, so their consumption is funded only by a universal lump-sum transfer T. Individuals of type $\theta = 1$ can generate true income y with utility cost $v(y)$, with $v_y > 0$ and $v_{yy} > 0$. The government applies a linear income tax rate τ over total reported income $z = y - e$, where e denotes evaded income. Income must be reported by detailing N income categories that add up to the total. Think of N as the number of lines the tax form has: the government could ask individuals only to report total income, to split it between capital and labor, to split labor income by firm and capital income by asset, and so on.[1] We refer to N as the *granularity* of the tax system. Individuals can under-report true income by evading some amount $e \leq y$ with utility cost $h(e; N)$, with $h_e > 0$ and $h_{ee} > 0$ and, crucially, $h_N > 0$, as evading becomes more costly when the government has access to more details about the income composition (e.g., because of third-party reporting or better audits).

Individuals incur privacy costs of two distinct kinds, summarized by the function $\psi(z; N)$. The first concerns disclosing their level of income. Individuals dislike it when the government perceives their level of affluence, so we assume $\psi_z > 0$. We refer to this cost as the *revealed affluence cost*. Individuals also dislike it when the government (or the public) knows details about their income composition conditional on affluence (e.g., the specific firms they work for or the specific assets they invest in), so we assume $\psi_N > 0$. We refer to this cost as the *granularity cost*. By assuming a general function ψ, we allow for potential interactions between the two privacy costs, meaning that ψ_{zN} is left unrestricted.

Individuals of type $\theta = 0$ do not generate income and, therefore, have utility $T - \psi(0; N)$. Individuals of type $\theta = 1$ solve the following problem:

$$\max_{y,e} U(y, e) = y - \tau(y - e) + T - v(y) - h(e; N) - \psi(y - e; N). \tag{1}$$

The first-order conditions (FOCs) with respect to y and e yield:

$$y: \qquad 1 - \tau = v_y + \psi_z, \tag{2}$$

$$e: \qquad \tau + \psi_z = h_e. \tag{3}$$

[1] By not characterizing or differentiating exactly what each line item is, we are not able to address some central optimal tax questions such as the optimal tax base or the relative tax on labor versus capital. Future research should seek to generalize our stylized characterization of granularity.

These expressions extend intuitions from standard labor supply models. Equation (2) shows that individuals choose their income by setting the marginal benefit from an additional dollar of true earnings (after-tax consumption) equal to the marginal cost of generating that dollar. The non-standard element in equation (2) is that the marginal privacy cost affects optimal earnings because the individual is willing to sacrifice consumption not to appear as affluent to the government and, due to leaking of information, to society. Equation (3) shows that individuals choose optimal evasion by setting the marginal benefit of doing so equal to its marginal cost. The non-standard element in equation (3) is that individuals are willing to evade even if the tax rate is zero because income understatement provides privacy benefits.

Note that privacy costs work as (marginal and average) implicit taxes on individuals. The marginal net utility return to an additional dollar of true earnings is $1 - \tau - \psi_z$. As such, privacy costs exacerbate the elasticity of real earnings to income taxes. However, because the net utility benefit of an additional dollar of evaded income is $\tau + \psi_z$, privacy costs attenuate the evasion responsiveness to the net-of-tax rate, as they encourage evasion regardless of the tax. It follows that the elasticity of taxable income with respect to the net-of-tax rate (ETI), $\varepsilon = [dz/d(1-\tau)] \cdot [(1-\tau)/z]$ will be a function of privacy costs. The net effect relative to a counterfactual with no privacy costs is ambiguous, as it will depend on the relative curvatures of v and h.

Planner's problem Indirect utility of individuals of type $\theta = 0$ is denoted by $V_0(\tau, T, N) = T - \psi(0; N)$. Indirect utility of individuals of type $\theta = 1$ is denoted by $V_1(\tau, T, N) = U(y(\tau, T, N), e(\tau, T, N))$. The government chooses (τ, T, N) to maximize a generalized utilitarian objective subject to a government budget constraint as follows:

$$\max_{\tau,T,N} SWF = s_0 G(V_0(\tau, T, N)) + s_1 G(V_1(\tau, T, N)) \quad \text{subject to} \quad s_1 \tau z = T, \tag{4}$$

where G is increasing and concave, reflecting social preferences for redistribution. Let μ denote the budget constraint multiplier. We define the social marginal welfare weights (WWs) as $g_0 = G'(V_0(\tau, T, N))/\mu$ and $g_1 = G'(V_1(\tau, T, N))/\mu$.

Proposition 1. *At the social optimum:*

$$\tau^* = \frac{1 - g_1}{1 - g_1 + \varepsilon}, \qquad N^* = \frac{\tau z \chi}{g_0 \psi_N + g_1 (h_N + \psi_N)}, \tag{5}$$

where $\varepsilon = [dz/d(1-\tau)] \cdot [(1-\tau)/z]$ is the elasticity of taxable income z with respect to the net-of-tax rate $1-\tau$ (ETI) and $\chi = [dz/dN] \cdot [N/z]$ is the elasticity of taxable income z with respect to the granularity N. Also, T is chosen such that $s_0 g_0 + s_1 g_1 = 1$.

Proof. See Appendix.

The optimal tax rate τ^* is identical to the standard optimal linear tax rate expression in Mirrleesian frameworks (Piketty and Saez, 2013). The ETI ε remains a sufficient statistic for assessing the efficiency costs of the income tax (Feldstein, 1999; Saez et al., 2012; Kleven, 2021). But the ETI ε is now a function of the privacy costs and the policy parameter N, in the spirit of Slemrod and Kopczuk (2002). As such, the optimal tax will depend on how much individuals value privacy and also on how much the government is controlling evasion possibilities through the choice of N. It follows, as in Keen and Slemrod (2017), that, for given privacy costs, τ and N are strategic complements: increasing N pushes e down, which in turn makes a given τ^* less distortionary.[2]

The optimal degree of granularity N^* balances its marginal revenue benefits, which are governed by the elasticity of taxable income with respect to the system's granularity χ, with its marginal utility costs, including both the utility penalty from higher evasion costs h_N and the pure privacy cost of providing income details to the government ψ_N. Note that the utility costs are weighted by the WWs, meaning that the government will consider restricting the information requirements, especially when high-WW (i.e., low utility) individuals experience large privacy costs at the margin. If privacy costs are concentrated in affluent individuals, the utilitarian optimal policy will favor larger values of N^*. This consideration is analogous to the result in Craig and Slemrod (2024), where the utility costs of misunderstanding the tax system affect optimal policy more when borne by low utility individuals.

Beyond this intuitive characterization of the privacy-enforcement tradeoff, the expression for N^* highlights which empirical objects are essential for assessing the tradeoff in a sufficient statistics logic. The elasticity χ is a reduced-form statistic for the revenue gains (driven by lower evasion rates and, therefore, higher taxable incomes) from increased granularity (additional information). Likewise, ψ_N represents the dollar-valued cost of providing additional information in the form of income details to the government. These objects have not been directly estimated in related literature. As such, our analysis provides guidance for future empirical research. However, in Section 4 we briefly survey related empirical work that suggests these sufficient statistics are empirically relevant, providing support for the importance of our normative analysis.

3 Illustrative Example

Consider iso-elastic functional forms and additively separable linear privacy costs with respect to both types of privacy:

$$v(y) = \frac{\theta}{1+\frac{1}{\epsilon}} \left(\frac{y}{\theta}\right)^{1+\frac{1}{\epsilon}}, \quad h(e;N) = \frac{1}{N\left(1+\frac{1}{\nu}\right)} (Ne)^{1+\frac{1}{\nu}}, \quad \psi(z;N) = \kappa z + \gamma N. \tag{6}$$

[2] T being chosen to set the average WW equal to one is a standard result in optimal tax models with quasi-linear utilities.

The solution of the individual problem is characterized by:

$$y^* = \theta\,(1-\tau-\kappa)^{\epsilon}, \qquad e^* = \frac{(\tau+\kappa)^{\nu}}{N}. \tag{7}$$

It becomes clear that marginal privacy costs due to affluence κ depress true incomes even in the absence of taxes, the same costs encourage individuals to understate income even in the absence of taxes, and the amount evaded is decreasing in N. While privacy costs due to granularity γ do not directly affect optimal decisions, they indirectly affect them through N (and, therefore, τ) because policy is optimally chosen by the government following Proposition 1. We also have that:

$$\varepsilon = \frac{dz^*}{d(1-\tau)}\frac{(1-\tau)}{z^*} = \left(\frac{\epsilon y^*}{1-\tau-\kappa} + \frac{\nu e^*}{\tau+\kappa}\right)\frac{1-\tau}{y^*-e^*}, \tag{8}$$

so privacy costs κ exacerbate responses in real incomes while attenuating evasion responsiveness to the net-of-tax rate. The net effect on the ETI is, in principle, ambiguous and will depend, in this stylized parametric example, on the relative sizes of ϵ and ν. Finally:

$$\chi = \frac{dz^*}{dN}\frac{N}{z^*} = \frac{e^*}{y^*-e^*}, \tag{9}$$

so the elasticity with respect to the information requirement is proportional to the ratio of evaded income to taxable income and, as such, is endogenous to the granularity N and the privacy costs κ.

Simulations We provide simple numerical simulations to provide further insight into the optimal policy problem and the implied comparative statics at the optimum. Our baseline simulation uses $(\theta, \epsilon, \nu, \kappa, \gamma, s_0) = (10, 0.5, 0.5, 0.1, 0.1, 0.1)$. These values are purely illustrative and are not meant to match empirical moments, so the levels of the solutions below are not directly interpretable.

The first numerical exercise consists of exploring the strategic complementarity at the joint optimum. For that purpose, we solve for the optimal τ^* given a value of N and for the optimal N^* given a value of τ. Figure 1 shows the results. The x-axis accounts for values of N and the y-axis accounts for values of τ. The blue curve with squares shows the optimal income tax τ^* as a function of a given level of N. The red curve with triangles shows the optimal granularity N^* as a function of a given level of τ^*. The figure clearly displays the strategic complementarity: both τ^* and N^* are monotonically increasing in the other policy parameter. As N increases, the planner optimally sets a larger τ^* due to the decreased efficiency cost of labor income taxation. Likewise, as τ increases, the planner optimally sets a larger N^*, as the benefit from information gathering is proportional to the marginal tax rate.

The second numerical exercise assesses the effects of the privacy costs on the elasticities ε and χ and, consequently, on the optimal policy. For this purpose, we solve for the joint optimum (τ^*, N^*, T^*) for different values of revealed affluence costs κ and granularity costs γ. Figure 2 shows the results.

Panels (a) and (b) show comparative statics with respect to the revealed affluence cost κ, fixing γ at its baseline level. Panels (c) and (d) show comparative statics with respect to the granularity cost γ, fixing κ at its baseline level. Panels (a) and (c) show that privacy costs reduce the optimal values of all three policy instruments.[3] Panels (b) and (d) shed light on one of the key reasons for this result: the sufficient statistics ε and χ increase with the privacy costs, both directly and indirectly through the policy parameters. In this stylized example, revealed affluence costs have stronger effects, as they affect the marginal decisions of individuals and, therefore, have larger impacts on the efficiency costs of taxation. Granularity costs affect the optimum only indirectly through the choice of N, which is the only variable whose relationship with respect to the granularity cost is stronger than that of the affluence parameter. This latter behavior is driven by the inclusion of utility considerations in the determination of N^* which, in this stylized example, are also experienced by individuals with zero income and, therefore, high WW.

4 Existing Empirical Evidence

Proposition 1 shows that the elasticity of taxable income with respect to the net-of-tax rate ε is needed for calibrating the optimal labor income tax τ^*. This parameter has been extensively studied in the related literature, so we will not pursue it here (see, e.g., Saez et al., 2012; Kleven and Schultz, 2014). More novel in Proposition 1 are the sufficient statistics needed for calibrating the optimal granularity N^*, most importantly the elasticity of taxable income with respect to the granularity χ and the marginal utility costs of granularity ψ_N. These objects have not been explicitly targeted by empirical research, and we believe that future research should aim to identify them. In what follows, we briefly survey related empirical evidence that, while not providing estimates of the aforementioned elasticities, suggests these sufficient statistics may be quantitatively important and, therefore, worthy of more exploration.

Enforcement returns to information That information provision facilitates tax enforcement, which constrains tax evasion, is undeniable (Slemrod, 2019). The compelling statistic supporting this assertion comes from the IRS tax gap studies, which conclude that the noncompliance rate for income subject to information reporting is under 1 percent, but is over 50 percent for income not subject to third-party reporting. Other studies supporting the effectiveness of third-party reporting include Kleven et al. (2011), Pomeranz (2015), Slemrod et al. (2017), and Naritomi (2019).

Recent research has also highlighted the importance of taxpayers' information in optimizing tax audits, both for ex-ante deterrence and ex-post detection of non-compliance. Battaglini et al. (2024) finds that using machine learning techniques to predict non-compliance can raise substantial revenue. Caspi et al. (2024) model optimal tax audits and find that the welfare benefit of improving the information available

[3] In Panel (c), the curve for optimal N^* is truncated as its optimal value grows significantly as γ goes to zero.

to the tax authority is proportional to the reduction in the conditional variance in the non-compliance prediction driven by the new information which, in principle, can be estimated with data on predicted revenue effects and audit costs. Paradisi and Sartori (2024) also explore optimal audits, and posit that the welfare benefit from increasing the precision of the non-compliance signal can yield better results than budget-equivalent increases in non-targeted audits.

Privacy costs What is known and unknown, circa 2024, about several aspects of tax privacy and its interaction with tax policy is summarized in Slemrod (2025a). Some particularly relevant details follow.

That many people (Americans, in this case) care at all about tax privacy is clear from surveys. For example, a 2023 Pew survey found that 71 percent of people are concerned about how the government uses the data it collects about them, while 77 percent say they have little to no understanding about what the government does with the data it collects about them (McClain et al., 2023). Very few people are likely aware of under what conditions the IRS can share its data, even with other government agencies, but these restrictions can be rendered moot by illegal leaks and recent executive actions to undermine them, such as the use of tax records to ascertain the addresses of suspected illegal immigrants.

One can also infer information about taxpayers' valuation of tax privacy by observing their behavior. Taxpayers in countries with public disclosure of non-granular tax return information above an income threshold bunch just below the threshold, suggesting that they value anonymity with respect to revealed affluence (Hasegawa et al., 2013; Hoopes et al., 2018). Pitt and Slemrod (1989) and Benzarti (2020) infer the perceived cost incurred in the process of itemizing deductions for income tax purposes by estimating how much tax saving taxpayers forego by taking the standard deduction rather than itemizing deductions, which requires revealing potentially sensitive information to the IRS about charitable contributions, medical expenses, and other expenditures. Although these studies infer that this bunching behavior is informative about the compliance costs of itemizing, they are likely generating an estimate of the sum of compliance costs and the privacy costs of itemizing, rather than just the compliance costs alone; disentangling the two awaits further research attention.

There is also anecdotal evidence that in some situations potential leaks from income tax agencies affect the income reports of those who fear negative non-tax consequences of knowledge of their affluence. Londoño-Velez (2012) mentions that, in Colombia, affluent individuals understated their wealth to the tax authority with no clear associated tax savings in order to reduce the chance they would be kidnapped and held for ransom. There is also evidence from Norway and Pakistan that public disclosure of incomes directly reduces tax evasion among the self-employed, presumably largely because of the fear that private whistle blowers could provide information to the tax authority that contradicts the disclosed tax reports (Bø et al., 2015; Slemrod et al., 2022).

5 Conclusion

Many people care about preserving the privacy of the information made available to the government via the income tax system. Although this issue is often prominent in public policy discussions about taxation, precisely how it should affect policy has not been rigorously addressed. In this paper, we begin that effort by developing a stylized model of income taxation when people value privacy.

In the model, we augment a standard labor-supply framework with two distinct aspects of concern about privacy. We assume that people care about what is revealed about their level of affluence, and also with the level of detail revealed about their income, conditional on affluence. This ingredient gives rise to a novel policy parameter, the granularity of the tax system, which is chosen by the government in conjunction with the tax rate. Two direct implications emerge. First, the tax system reduces true earnings not only because of the statutory marginal tax rate, but also because of the revealed-affluence-related privacy cost of income. Second, individuals will choose to understate total income even if the tax rate is zero, because of the privacy utility gain from reducing the income revealed to the government. These conclusions imply that privacy concerns impact behavioral responses and, therefore, optimal policy.

Using the privacy-expanded model, we show that the optimal labor income tax rate still depends, as in the standard model, on the ETI. However, this elasticity is structurally modified by privacy concerns. Privacy costs act as an effective tax rate because an increase in the statutory tax rate generates a bigger percentage decline in the net-of-tax-and-privacy-costs tax rate, increasing its value. However, introducing privacy costs reduces the evasion elasticity component of the elasticity of taxable income with respect to the net-of-tax rate because the tax saving constitutes a smaller portion of the total return to evasion, given that individuals will understate income even in the absence of taxes.

We also characterize the optimal granularity of the tax base. Its optimal size depends positively on the tax rate and the elasticity of taxable income with respect to the granularity (as a more granular system makes evasion costlier), and negatively on the utility cost of granularity due to the resulting higher evasion cost and the pure privacy costs of providing more information to the government. We argue that these objects should be targeted by future empirical research. The utility costs are weighted by marginal social welfare weights, so the government should be particularly deterred from gathering information when these costs are concentrated in high-welfare-weight, that is low-income, individuals. The analysis also shows that these instruments are strategic complements: higher taxes call for a higher return to information-driven enforcement and, vice versa, making explicit the importance of considering the joint optimization of the two policy instruments instruments.

The model is stylized enough to prevent us from tackling all privacy-related concerns associated with tax policy. While being a first step towards understanding the role of tax privacy for optimal policy, several extensions are worth exploring in future research. First, we model the granularity of the tax system in terms of the volume of information collected, conditional on a tax base, without distinguishing between different sources. Taking a more general approach that incorporates the nature of each income component could provide room to explore questions such as the optimal tax base or the optimal level of tax deductions. Second, we posit that information helps policy because of enforcement benefits, but we abstract from other efficiency gains such as tagging (Akerlof, 1978; Mankiw and Weinzierl, 2010). Incorporating tagging and its conflict with other privacy concerns related to, for example, identity-based policies (Slemrod, 2025b), is also a fruitful avenue of research. Finally, our model abstracts from heterogeneity in privacy costs, which may render the progressivity of the income tax less effective than otherwise (as in Kopczuk, 2001). Assessing the relative optimality of income taxation and the granularity of the tax system in the presence of multidimensional heterogeneity is also worth exploring in future research.

Appendix: Proof of Proposition 1

The Lagrangian of the planner is given by:

$$\mathcal{L}(\tau, T, N) = s_0 G(V_0(\tau, T, N)) + s_1 G(V_1(\tau, T, N)) + \mu\,(s_1 \tau z - T)\,. \tag{A.1}$$

The FOC w.r.t. T is given by:

$$\frac{\partial \mathcal{L}}{\partial T} = s_0 G'(V_0(\tau, T, N)) + s_1 G'(V_1(\tau, T, N)) - \mu = 0. \tag{A.2}$$

Rearranging terms yields $s_0 g_0 + s_1 g_1 = 1$.

The FOC w.r.t. τ is given by:

$$\frac{\partial \mathcal{L}}{\partial \tau} = -s_1 G'(V_1(\tau, T, N))z + s_1 \mu \left(z - \tau \frac{dz}{d(1-\tau)} \right) = 0, \tag{A.3}$$

where we used the envelope theorem. Rearranging terms yields $\tau^* = (1 - g_1)/(1 - g_1 + e)$.

The FOC w.r.t. N is given by:

$$\frac{\partial \mathcal{L}}{\partial N} = -s_0 G'(V_0(\tau, T, N)\psi_N - s_1 G'(V_1(\tau, T, N))\,(h_N + \psi_N) + s_1 \mu \tau \frac{dz}{dN} = 0, \tag{A.4}$$

where we used the envelope theorem. Rearranging terms yields $N^* = \tau z \chi / [g_0 \psi_N + g_1(h_N + \psi_N)]$. $\square$

Bibliography

Akerlof, George A, "The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning," *The American Economic Review*, 1978, *68* (1), 8–19.

Battaglini, Marco, Luigi Guiso, Chiara Lacava, Douglas L Miller, and Eleonora Patacchini, "Refining public policies with machine learning: The case of tax auditing," *Journal of Econometrics*, 2024, p. 105847.

Benzarti, Youssef, "How taxing is tax filing? Using revealed preferences to estimate compliance costs," *American Economic Journal: Economic Policy*, 2020, *12* (4), 38–57.

Bø, Erlend E, Joel Slemrod, and Thor O Thoresen, "Taxes on the internet: Deterrence effects of public disclosure," *American Economic Journal: Economic Policy*, 2015, *7* (1), 36–62.

Caspi, Aviv, Jacob Goldin, Daniel Reck, and Daniel E Ho, "Optimal tax audits using predictions," *Working Paper*, 2024.

Craig, Ashley C and Joel Slemrod, "Tax knowledge and tax manipulation: A unifying model," *Journal of Political Economy Microeconomics*, 2024, *2* (2), 298–334.

Feldstein, Martin, "Tax avoidance and the deadweight loss of the income tax," *Review of Economics and Statistics*, 1999, *81* (4), 674–680.

Hasegawa, Makoto, Jeffrey L Hoopes, Ryo Ishida, and Joel Slemrod, "The effect of public disclosure on reported taxable income: Evidence from individuals and corporations in Japan," *National Tax Journal*, 2013, *66* (3), 571–607.

Hoopes, Jeffrey L, Leslie Robinson, and Joel Slemrod, "Public tax-return disclosure," *Journal of Accounting and Economics*, 2018, *66* (1), 142–162.

Keen, Michael and Joel Slemrod, "Optimal tax administration," *Journal of Public Economics*, 2017, *152*, 133–142.

Kleven, Henrik J, "Sufficient statistics revisited," *Annual Review of Economics*, 2021, *13* (1), 515–538.

Kleven, Henrik Jacobsen and Esben Anton Schultz, "Estimating taxable income responses using Danish tax reforms," *American Economic Journal: Economic Policy*, 2014, *6* (4), 271–301.

_ , Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

Kopczuk, Wojciech, "Redistribution when avoidance behavior is heterogeneous," *Journal of Public Economics*, 2001, *81* (1), 51–71.

Londoño-Velez, Juliana, "Income and wealth at the top in Colombia: An exploration of tax records 1993-2010," *Master Thesis. Paris School of Economics, Paris.*, 2012.

Mankiw, N Gregory and Matthew Weinzierl, "The optimal taxation of height: A case study of utilitarian income redistribution," *American Economic Journal: Economic Policy*, 2010, *2* (1), 155–176.

McClain, Colleen, Michelle Faverio, Monica Anderson, and Eugenie Park, "How Americans view data privacy," *Pew Research Center*, 2023, *18*, 2023.

Naritomi, Joana, "Consumers as tax auditors," *American Economic Review*, 2019, *109* (9), 3031–3072.

Paradisi, Matteo and Elia Sartori, "Using prediction models to design tax enforcement: Incentives vs targeting," *Working Paper*, 2024.

Piketty, Thomas and Emmanuel Saez, "Optimal labor income taxation," in "Handbook of Public Economics," Vol. 5, Elsevier, 2013, pp. 391–474.

Pitt, Mark M and Joel Slemrod, "The compliance cost of itemizing deductions: Evidence from individual tax returns," *American Economic Review*, 1989, *79* (5).

Pomeranz, Dina, "No taxation without information: Deterrence and self-enforcement in the value added tax," *American Economic Review*, 2015, *105* (8), 2539–2569.

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

Slemrod, Joel, "Tax compliance and enforcement," *Journal of Economic Literature*, 2019, *57* (4), 904–954.

_ , "Tax privacy," *Journal of Economic Perspectives*, 2025, *39* (1), 205–224.

_ , "Taxing identity," *Working Paper*, 2025.

_ **and Wojciech Kopczuk**, "The optimal elasticity of taxable income," *Journal of Public Economics*, 2002, *84* (1), 91–112.

_ **, Brett Collins, Jeffrey L Hoopes, Daniel Reck, and Michael Sebastiani**, "Does credit-card information reporting improve small-business tax compliance?," *Journal of Public Economics*, 2017, *149*, 1–19.

_ **, Obeid Ur Rehman, and Mazhar Waseem**, "How do taxpayers respond to public disclosure and social recognition programs? Evidence from Pakistan," *Review of Economics and Statistics*, 2022, *104* (1), 116–132.

Figure 1: Policy Reaction Functions

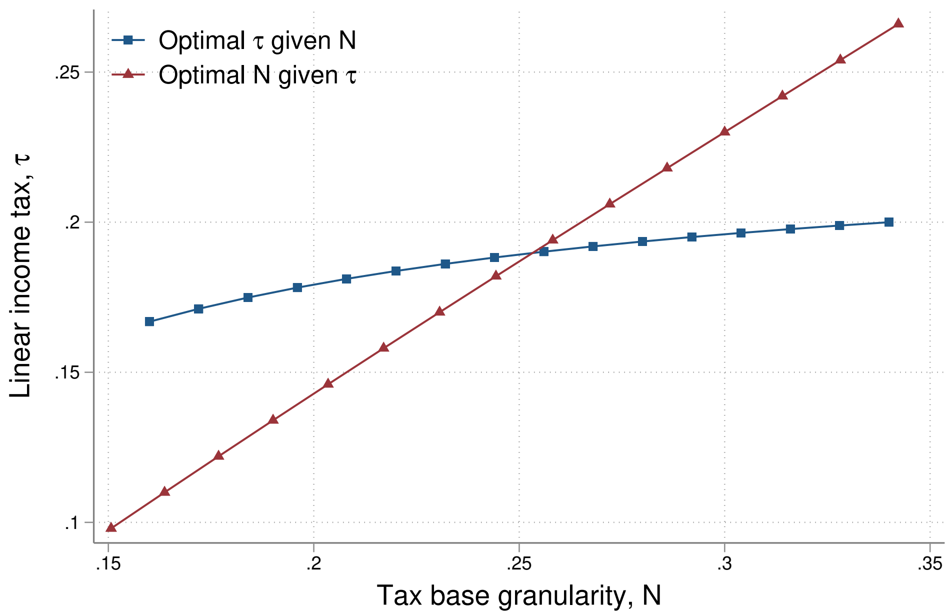

Notes: This figure presents numerical simulations of optimal τ^* given different values of N (blue curve with squares) and optimal N^* given different values of τ (red curve with triangles) computed using Proposition 1 and the parametric model of Section 3. These simulations use $(\theta, \epsilon, \nu, \kappa, \gamma, s_0) = (10, 0.5, 0.5, 0.1, 0.1, 0.1)$.

Figure 2: Comparative Statics: Privacy Costs

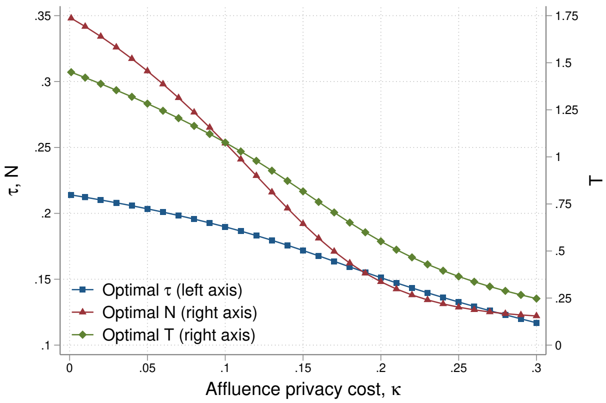

(a) Optimal policy - Varying affluence cost κ

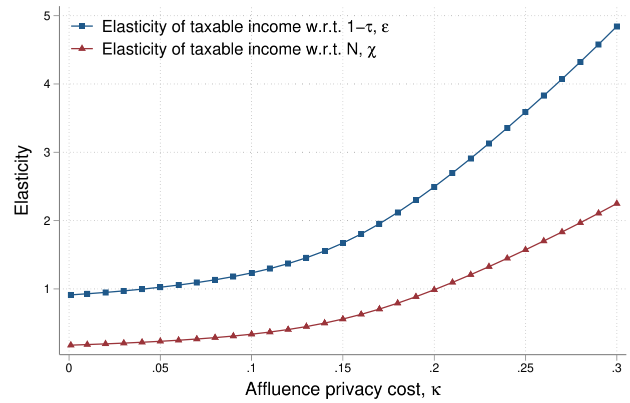

(b) Elasticities - Varying affluence cost κ

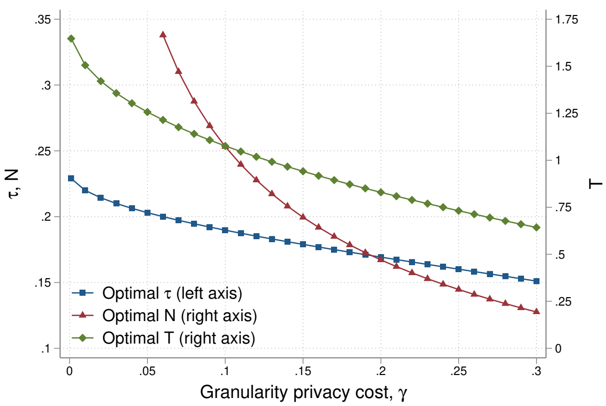

(c) Optimal Policy - Varying granularity cost γ

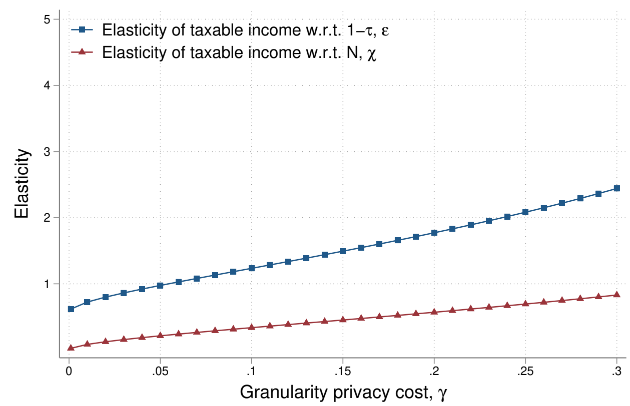

(d) Elasticities - Varying granularity cost γ

Notes: This figure presents numerical simulations of optimal policies and elasticities for different values of privacy costs κ and γ computed using Proposition 1 and the parametric model of Section 3. These simulations use $(\theta, \epsilon, \nu, \kappa, \gamma, s_0) = (10, 0.5, 0.5, 0.1, 0.1, 0.1)$. Panels (a) and (b) provide comparative statics with respect to κ, fixing γ at 0.1. Panel (a) shows optimal policy (τ^*, N^*, T^*) and Panel (b) shows the elasticities ε and χ. Panels (c) and (d) provide comparative statics with respect to γ, fixing κ at 0.1. Panel (c) shows optimal policy (τ^*, N^*, T^*) and Panel (d) shows the elasticities ε and χ.