Optimal Tax Audits Using Predictions*

Aviv Caspi Jacob Goldin

Daniel Reck Daniel E. Ho

January 24, 2024

Abstract

Tax authorities use audits to detect and deter tax evasion. In practice, they commonly rely on quantitative predictions about taxpayers' noncompliance to inform their decisions about which taxpayers to audit. We study the problem of optimal audit selection in this context. Specifically, we investigate how variation in the distribution of predicted noncompliance across taxpayers, including the uncertainty in quantitative predictions, shapes optimal audit policy. Our results highlight the contribution of these factors through a sufficient statistics characterization of the optimal audit selection rule. We leverage this characterization to quantify the social welfare benefit of varying the information available to the tax authority, for example by expanding third-party information reporting.

*Caspi: Stanford University; Goldin: University of Chicago and NBER; Reck: University of Maryland and NBER; Ho: Stanford University. We gratefully acknowledge financial support from Arnold Ventures. Francesco Armillei provided excellent research assistance. For helpful discussion and feedback, we thank Tom Hertz and Joel Slemrod.

1 Introduction

Questions about the proper allocation of tax enforcement resources are front and center in contemporary economic policy debates. In recent years, U.S. policymakers from both parties have proposed or enacted dramatic funding changes to IRS enforcement (Shaw, 2024), new requirements for banks to report information to the IRS for their high-income account holders (Rappeport and Weisman, 2021), and commitments to increase audit scrutiny on high-income individuals and complex businesses while avoiding changes to the audit rate for taxpayers with incomes below $400,000 (Yellen, 2022).

In considering and implementing such policies, governments must make decisions about which individuals to audit with incomplete information about what an audit will reveal. This fundamental difficulty has long been recognized in writing about tax evasion as an asymmetric information problem, but the difficulty becomes especially pressing for policymaking in the presence of distributional concerns. For example, policymakers might wish to increase audit scrutiny of high-income taxpayers who evade, but the tax authority cannot directly observe which taxpayers evade without an audit.

In this paper, we provide a framework for analyzing optimal audit selection that incorporates the information constraints faced by the tax authority and the capacity of taxpayers to respond behaviorally to changes in audit selection rules. We provide an empirically implementable, sufficient statistics characterization of the optimal audit selection rule. This characterization provides a principled way to assess the allocation of government resources toward audits of different groups of individuals. Specifically, our characterization admits a test of the optimality of any change in audit rates that is feasible given the information the government receives from taxpayer self-reports and third parties. We then further characterize the welfare benefits of allowing the government to observe and use additional information to improve its audit selection rule.

The primary challenge in modelling optimal audit selection is to model the information of individuals and the government realistically while maintaining tractability. Tax returns are high-dimensional by nature: individuals are making high-dimensional decisions about what to report on their returns, and the government receives high-dimensional information from individual's self-reports and from third parties. Our model accommodates these features of the real world, and we maintain tractability by focusing on the first-order sufficient statistics for optimal policy. To this end, a key feature of the model is the analysis of a rational expectations equilibrium – there is no aggregate uncertainty, and individuals and the government correctly anticipate each others' actions to the extent they can given their information – so that we can leverage classic results like the equivalence between ex ante and ex post welfare and the envelope theorem to simplify our analysis.

Apart from this information structure, the setup of our model follows prior work on tax compliance (Allingham and Sandmo, 1972) and optimal tax systems (Mayshar, 1991; Keen and Slemrod, 2017). Individuals are uncertain about whether they will be audited and choose a self-report (the contents of their tax return) to maximize their expected utility. Utility depends on final consumption; individuals remit taxes based on their self-report, and if an individual is audited, their consumption is reduced by the revenues and penalties from detected non-compliance, which are transferred to the government, and by the resource cost of being audited. The government chooses an audit selection rule based on individual self-reports and the contents of third-party information reports and potentially other exogenous information, in order to maximize a generalized utilitarian social welfare function. For simplicity, we assume that audit-recovered revenues are rebated in lump-sum fashion to the full population, and we do not consider real responses to change in audit selection rules.[1]

We analyze the first-order welfare effect of perturbations to the audit selection rule. Following the conventional envelope theorem logic, the effect of increasing audits on the margin for some observable group comprises a direct effect on private welfare and both direct and behavioral effects on the government's budget. The private, direct welfare effect is negative and includes revenues transferred to the government (additional taxes due and penalties) and the costs imposed by audits on individuals, weighted by the relevant welfare weight for the audited taxpayer. The direct fiscal effect includes the revenues received from audited individuals and the administrative cost of audits. The behavioral fiscal effect arises because changes to the audit selection rules will induce individuals to alter their report, i.e. a deterrence effect (Allingham and Sandmo, 1972).

The government evaluates expanding audits by integrating all of these effects over observable groups of individuals given its information. This allows us to characterize the optimal audit rate in terms of conditional means given the government's information, plus another term describing how audit revenues and private audit costs covary with welfare (Pareto) weights. Apart from this covariance, mean predictions are sufficient. We sign the covariance under some additional, conventional structure: holding taxpayer self-reports fixed, higher-than-expected audit revenue is a signal of higher consumption/income than expected, and, therefore, under utilitarianism with diminishing marginal utility, a lower welfare weight. This implies the covariance between audit revenues and welfare weights is negative.[2]

Building on these results, we provide an implicit characterization of optimal audit rates, following the rubric of the characterization of optimal tax rates in Saez (2001). A key step here is quantifying the fiscal revenue effect in terms of an elasticity of behavioral response, following Keen and Slemrod (2017); for us this is a weighted elasticity across taxpayers affected by a local perturbation to the audit rate. This behavioral elasticity governs the size of the deterrence effect due to additional taxes paid in response to increased audits, as well as the effect of behavioral responses on audit revenues. We then provide an implicit characterization of the optimal audit

[1] By the conventional envelope theorem logic, real responses would be of first-order importance for optimal policy to the extent that they modify tax revenues. Accommodating real responses along these lines is theoretically straightforward; an emerging empirical question would be the magnitude of the behavioral revenue effect due to real responses to changes in audit policy.

[2] Another potential reason for this covariance is if society attaches different welfare weights based more directly on compliance, i.e. the government puts a lot of weight on the costs imposed by audits on compliant taxpayers, but not much weight on the costs imposed by audits on non-compliant taxpayers.

rate for any group of individuals for whom it is feasible to increase the audit rate given the government's information. We discuss the optimality of corner solutions at a 100% or 0% audit rate for some group, and we provide an implicit characterization of the optimal audit rate for that group in terms of sufficient statistics. Apart from the judgments governing welfare weights, these sufficient statistics are empirically implementable given data on mean predicted revenue effects and costs of audits, the variance of predicted audit revenues and private audit costs, and estimates of the elasticity of behavioral response to changes in audit rates.

We build on this sufficient statistics characterization with two extensions. The first extension characterizes the social value of allowing the government to use new information for audit selection. We analyze this by first deriving the direction in which new information perturbs the audit selection rule, expressed in terms of how the new information modifies our sufficient statistics for audit selection. We then derive the first-order social welfare effect of such a perturbation in audit selection. The resulting characterization is highly intuitive: new information increases welfare at the optimum in proportion to the variance in individual-specific welfare effects of audits that is *explained* by the new information. The resulting characterization is also empirically implementable, given data on predictions and the conditional variance term described above with and without conditioning on additional information.[3]

The second extension incorporates dynamic effects of changes in audit selection. A key question here is how government and private information evolves over time as a consequence of audit policy. We illustrate one potential information structure along these lines, which is that auditing a given taxpayer provides information to the government about that taxpayer that is informative for their future tax liability. Individuals know that this information is received by the government when they are audited. In this case, we show that the effect of an audit today includes the change in future revenues remitted by individuals to the government (including taxes paid and audit revenues); from a reduced-form perspective this is essentially a new component of both private and fiscal direct effects. More broadly, we note that this effect, which has been studied by much prior literature,[4] and its importance for welfare emerges as a consequence of the specific structure we imposed on the relationship between current audits and future information, and we hope this insight will lead to broader formulations of audit information effects in the future.

Our paper contributes to the literature on tax systems (for example, Mayshar, 1991; Keen and Slemrod, 2017; Boning et al., 2023). Many applications in the optimal tax systems literature focus on efficiency criteria, with less consideration of distributional concerns. Relative to a standard optimal tax systems approach to audit policy, the innovative component of our theoretical results is the examination of how predicted audit outcomes – e.g. the product of machine learning or other predictive algorithms – optimally inform the selection of individual audits. We do so with

[3]Note that we assume new information is useful for prediction non-compliance but not for detecting it during and audit, i.e. the revenues recovered by auditing a given individual are unaffected by new information.

[4]For empirical examination of the effect of audits on future reporting behavior and revenues, see DeBacker et al. (2018); Hebous et al. (2023); Mazzolini et al. (2022); Advani et al. (2023). Boning et al. (2023) also integrate this effect into welfare analysis within an MVPF framework.

much less structure on the information of individuals and the government than that of prior theoretical work on optimal audit selection (e.g. Reinganum and Wilde, 1985; Sanchez and Sobel, 1993). One cost of our approach is that we cannot derive closed-form characterizations of the optimal audit rule in terms of primitive structural parameters; rather, we rely on reduced-form sufficient statistics. In this sense, a recent paper related to ours is Boning et al. (2023), which presents a theoretical and empirical analysis of the social value of expanding audits at different parts of the taxpayer-reported income distribution, holding the within-group audit selection rule fixed, using the MVPF approach of Hendren and Sprung-Keyser (2020). We focus more broadly on the optimal audit selection rule itself, i.e. on the social value of expanding audits conditional on observing any particular information from taxpayer reports and the government. Focusing on localized policy variation allows us to derive practical, straightforward criteria for individual audit selection based on the contents of (originally filed) tax returns, third-party reports, and normative parameters that reflect distributional concerns (i.e. policy priorities). One technical difference between approaches is that the MVPF approach requires comparing the effects of audits on money-metric private welfare and the government budget to a mean welfare weight, while our characterization explicitly allows the welfare weight to covary with audit outcomes.

Our work also contributes to the nascent literature on the use of predictive algorithms in policymaking (e.g Kleinberg et al., 2015). Prior work by researchers at the Stanford RegLab trains machine-learning algorithms to predict the effects of real-world audits (Black et al., 2022; Henderson et al., 2023; Elzayn et al., 2023). In future work, we hope to build on this work and analyze the optimal use of these algorithms in audit selection, describing how predictions and prediction uncertainty may shape optimal audit selection using empirical audit data. The theoretical approach we develop here could also be applied to other policy problems in which policymakers must target an intervention on the basis of the predicted results of the intervention.

2 Setup

2.1 Overview

Individuals are endowed with private information – their type – which encompasses all of the information that is relevant to determining their tax liability (e.g., income, credits, and deductions). Individuals file tax returns reporting this information to the government. In addition to the self-reported type from the individual, the government receives an exogenous signal of the individual's type from other sources of information. Based on all the information it observes, the government selects which individuals to audit. Finally, revenues from taxation and audits are rebated lump-sum to individuals and ex post payoffs are realized.

We will characterize, using the calculus of variations, the audit selection rule that maximizes social welfare in a rational expectations (Perfect Bayesian) equilibrium of this game.

2.2 Information

At the core of the model is asymmetric information: the government does not observe which individuals are under-reporting their taxes, nor by how much. An individual's private type is a realization of a multi-dimensional random vector, $\theta \in \mathbb{R}^N$. We define θ to include all information that is relevant to determining the individual's liability under the non-linear income tax, $T(\cdot) : \mathbb{R}^N \longrightarrow \mathbb{R}$. For example, θ includes an individual's income, filing status, number of dependents, as well as any other information relevant for determining eligibility for tax credits or deductions. Through the tax filing process, individuals self-report to the government their type, $\widehat{\theta} \in \mathbb{R}^N$, and remit their self-reported tax liability, $T(\hat{\theta})$. The government observes $\hat{\theta}$ but not θ. We model the components of θ, such as an individual's labor supply or marital status, as exogenous.

The government also observes an exogenous signal about each individual, which is a realization of a random variable $\sigma \in \mathbb{R}^M$.[5] We think of σ as containing the information available to the government for selecting audits beyond what is self-reported by the taxpayer. For example, σ may include mandatory third-party information reporting, whistleblower reports, or non-tax administrative databases. The joint distribution $f(\theta, \sigma)$ is common knowledge; however, individuals do not observe the realized value of σ when deciding what to self-report on their tax return. Because σ is an exogenous signal, the government cannot invest to gather more information prior to selecting audits. We relax this assumption in Section 3, where we quantify the value of expanding the information for audit selection.

We characterize optimal policy in a rational expectations equilibrium of an incomplete-information game. The equilibrium condition imposes that individuals and the government act optimally given their rationally formed expectation of how the other party will behave and the information that is available to them. Given this information structure, the government's beliefs about an individual's true type are determined by the population distribution of types conditional on the government's observed signal and self-reported information, i.e. $f(\theta|\hat{\theta}, \sigma)$. Where there is pooling across types in equilibrium, such as would be caused by non-compliant types attempting to resemble compliant types, the government faces uncertainty conditional on a realization of $(\hat{\theta}, \sigma)$. Meanwhile, individuals form beliefs about what the government knows about them, and therefore their audit risk, based on $f(\sigma|\theta)$.

2.3 Audit Policy

The government's audit policy consists of a function, $A(\hat{\theta}, \sigma) : \mathbb{R}^N \times \mathbb{R}^M \longrightarrow [0, 1]$, which maps the government's information to a probability of audit. When $A(\hat{\theta}, \sigma) \in (0, 1)$ for some particular $(\hat{\theta}, \sigma)$, the government, being unable or unwilling to differentiate between returns with the information available, randomly selects returns for audit at the rate $A(\hat{\theta}, \sigma)$. For tractability, we assume that the joint distribution of (θ, σ) generates self-reports $\hat{\theta}$ and government audit rates

[5] We introduce notation for realizations of random variables and we suppress notation for the random variables themselves.

$A(\hat{\theta}, \sigma)$ that, in equilibrium, are everywhere differentiable with respect to $\hat{\theta}$.[6] When we evaluate these derivatives at corners, where $A(\hat{\theta}, \sigma) = 0$ or $A(\hat{\theta}, \sigma) = 1$, we use the respective right or left derivative of A with respect to $\hat{\theta}$ to evaluate a marginal change in the audit rule, i.e. moving away form the corner on the margin.

Let $a \in \{0, 1\}$ denote whether an individual is audited. Under rational expectations, individuals of type θ form beliefs about their audit risk for a potential self-report $\hat{\theta}$ under audit policy A as follows:

$$p_\theta(\hat{\theta}, A) \equiv P(a = 1|\theta, \hat{\theta}, A) = \int_\sigma A(\sigma, \hat{\theta}) dF(\sigma|\theta) \tag{1}$$

An arbitrary perturbation of the audit policy, dA, modifies type θ's perceived audit risk according to:

$$dp_\theta(\hat{\theta}, A) = \int_\sigma dA(\sigma, \hat{\theta}) dF(\sigma|\theta) \tag{2}$$

Intuitively, if the audit rate increases for some particular σ, the individual's perceived audit probability increases in proportion to the likelihood that the government observes that value of σ for them. If the audit rate increases for some value of $\hat{\theta}$, this increases the individuals' perceived risk of audit from reporting that value of $\hat{\theta}$ at all realizations of σ.

Examples of Perturbations to the Audit Selection Rule Before introducing the rest of the model, we illustrate a few perturbations to the audit selection rule to build intuition. For these examples, we simplify the information environment to facilitate visualization.

Suppose that for a given type θ, there are two potential realizations of the government's signal $\sigma \in \{\sigma_1, \sigma_2\}$, and that the self-reported type $\hat{\theta}$ consists of a single dimension, which we label *reported income*. Suppose that σ_1 is a signal that the individual is a high-income type and σ_2 is a signal that the individual is a low-income type. We can then illustrate an audit selection rule $A(\hat{\theta}, \sigma)$ under each realization of σ, with reported income on the x axis and the audit rate on the y axis. We depict the audit rules as an inverted s-shape, based on the notion that the government might have a good idea conditional on its observation of σ that the individual has at least some amount of income, so the audit becomes much less likely when reported income exceeds some amount (Kleven et al., 2011).

Figure 1 illustrates three examples of perturbations to the audit selection rule, on the left panels, and the corresponding effect on perceived audit probabilities on the right panels. In our first example, we imagine increasing the audit rate by a marginal amount da for all individuals with income below a cutoff y_0. So in Figure 1a on the left, the audit rate increases by a constant da below y_0, for both values of σ. An individual of type θ then maps this to their perceived detection probability, $p_\theta(\hat{\theta}, A)$. Note that following equation (1), the status quo perceived detection

[6] One reason this assumption could fail, for example, is if the optimal audit rate jumps to 100% when the individual self-report contradicts third-party information (a component of σ). Kleven et al. (2011) present a model of taxpayer reporting with third-party information that captures this underlying idea but maintains differentiability, supposing that the audit probability is everywhere continuous but potentially very highly elastic around the third-party reported amount of income.

probability is a mixture of the status quo $A(\hat{\theta}, \sigma)$ over the two values of σ, where the weights in the mixture are given by $f(\sigma|\theta)$ – we imagine roughly equal weights for type θ in this example, i.e. $f(\sigma_1|\theta) \approx 0.5$. Because in this first example the audit rule increases by da below y_0 regardless of σ, the perceived audit rule also increases by da below y_0.

In our second example, we increase audit rates marginally at a particular realization of σ, σ_2. This is represented by a level shift of da in $A(\hat{\theta}, \sigma_2)$ at all reported incomes, while $A(\hat{\theta}, \sigma_1)$ remains unchanged. In this case, the perceived probability for type θ increases by $f(\sigma_2|\theta)da$ regardless of reported income. In words, the individual is only threatened by this increase in audit rates to the extent that they believe the government receives the signal σ_2.

We focus below on perturbations similar to Example 3 in Figure 1, i.e. perturbations based on arbitrarily fine partitions of the government's information. In Example 3, we increase the audit rate by da for values of reported income in some open interval (y_0, y_1) and $\sigma = \sigma_2$. The individual's detection probability therefore increases in the range (y_0, y_1), and the increase is weighted by $f(\sigma_2|\theta)$. So we have, altogether $dp_\theta = 1\{\hat{\theta} \in (y_0, y_1)\} f(\sigma_2|\theta) da$. More generally, we characterize the welfare effects of increases in audits of da in open intervals of both θ and σ. We focus on this type of perturbation because it facilitates aggregation to arbitrary perturbations, and it allows us to study the optimal use of prediction algorithms in audit selection.

2.4 Individual Behavior and Welfare

Individuals' well-being depends on their type and their consumption, $u_\theta(c)$. In turn, consumption depends on the type that an individual reports ($\hat{\theta}$) and whether the individual is audited (a) according to:

$$c_\theta^a = z_\theta - T(\hat{\theta}) + G - a[R_\theta(\hat{\theta}) + H_\theta],$$

where z_θ denotes the individual's (exogenous) pre-tax income and G is a lump-sum transfer from rebated tax and audit revenues. In the event of an audit ($a = 1$), the individual's consumption is reduced by $R_\theta(\hat{\theta})$, which encompasses the total monetary payment imposed from the audit, inclusive of underpaid taxes, interest, and penalties. When an audit occurs, the taxpayer's consumption is also reduced by H_θ, which encompasses the compliance burden of the audit to the individual in units of consumption. To illustrate, an audited taxpayer who was determined by the tax authority to have accurately reported their type would be worse off to the extent $H_\theta > 0$.[7]

Individuals choose which type to report to maximize expected utility based on the audit policy:

[7] We note that H_θ is assumed to be a deterministic function of θ. Thus although H_θ may co-vary with other type-specific outcomes such as under-reporting decisions, audit costs do not depend directly on reporting choices conditional on type.

Figure 1: Perturbations of Audit Selection Rules and Perceived Audit Probabilities

(a) Ex 1: Increase Audits Below Some Cutoff Report

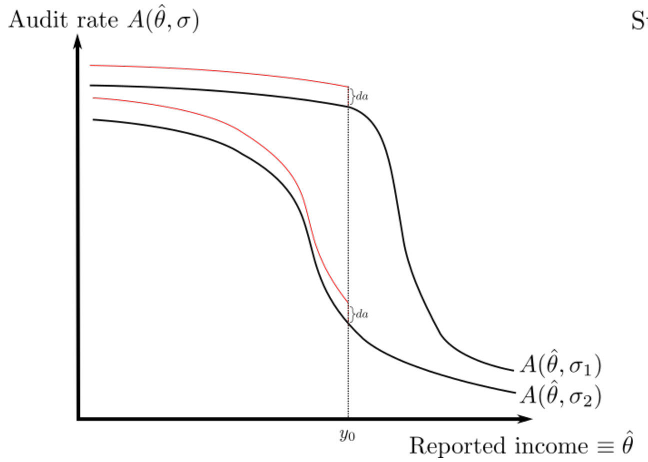

(b) Ex 1: Effect on Perceived Audit Probabilities

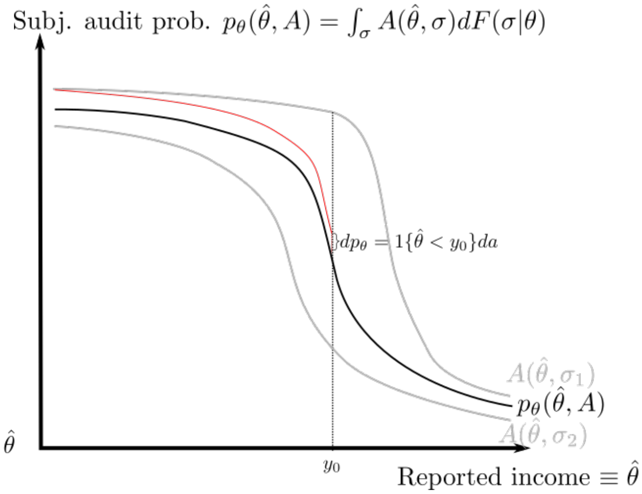

(c) Ex 2: Increase Audits for Some σ

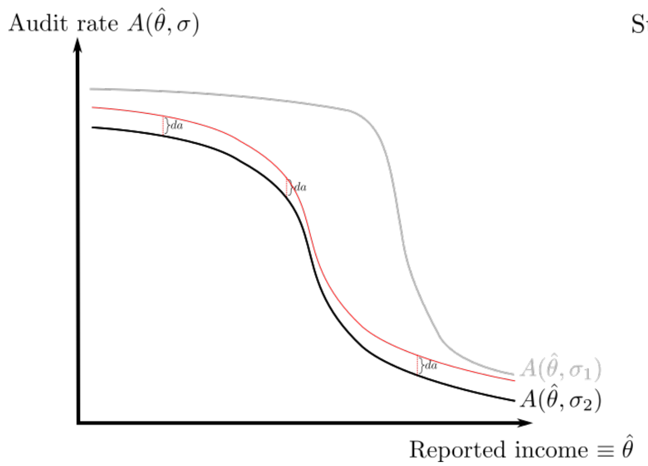

(d) Ex 2: Effect on Perceived Audit Probabilities

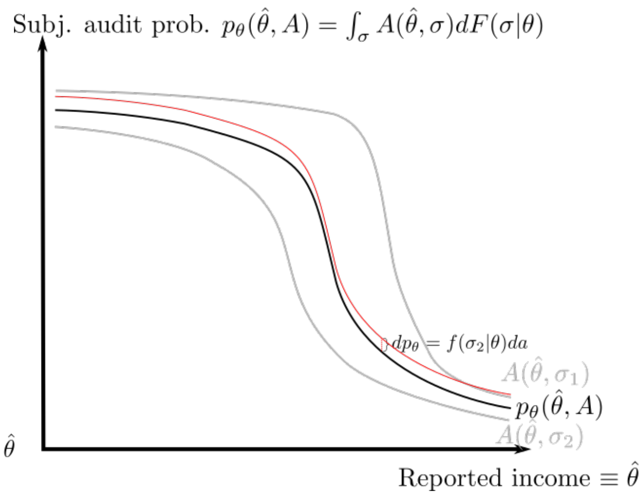

(e) Ex 3: Increase Audit Rate for particular σ in an interval for reported income

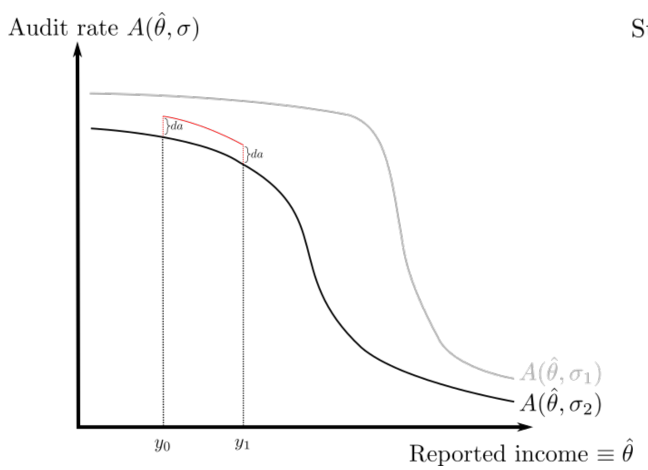

(f) Ex 3: Effect on Perceived Audit Probabilities

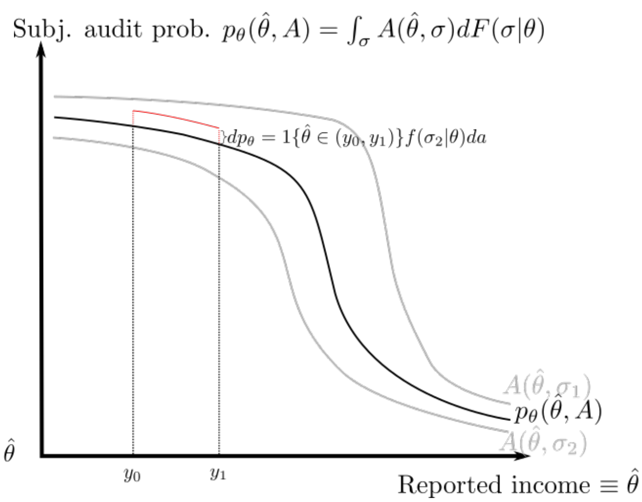

$$\max_{\hat{\theta}} \quad p_\theta(\hat{\theta}, A)u_\theta(c^1_\theta(\hat{\theta})) + (1 - p_\theta(\hat{\theta}, A))u_\theta(c^0_\theta(\hat{\theta})) \tag{3}$$
$$\text{where} \qquad c^a_\theta(\hat{\theta}) = z_\theta - T(\hat{\theta}) + G - a[R_\theta(\hat{\theta}) + H_\theta].$$

We denote optimal reporting for type θ by $\hat{\theta}(\theta, A)$, or when the context is clear, $\hat{\theta}_\theta$. Optimal consumption in audit state a is denoted $c^a(\theta, A)$.

In order to simplify behavioral responses and build an understanding of welfare weights, we make the simplifying assumption that the penalty R_θ is linear in the amount of underreported taxes, $R_\theta(\hat{\theta}) = (1+\rho)[T(\theta) - T(\hat{\theta})]$ for a known penalty parameter $\rho > 0$. then $dR_\theta = -(1 + \rho)dT_\theta$. This assumption simplifies our expressions but it is straightforward to relax.[8]

We next consider how individuals' expected welfare at the optimum, denoted $v(\theta, A, G)$ depends on the audit policy $A(\hat{\theta}, \sigma)$. A perturbation to the audit rule dA (holding G fixed) will affect optimal behavior, but as a consequence of the envelope theorem, the sole first-order effect of a marginal change in the audit rule dA on private welfare for type θ is the direct effect:

$$dv_\theta = [u_\theta(c^1(\theta, A)) - u_\theta(c^0(\theta, A))]dp_\theta. \tag{4}$$

where dp_θ denotes the change in the audit rate for type θ associated with the audit policy change dA, as specified by equation (2). Using a first-order Taylor approximation we can write:

$$dv_\theta \approx -\overline{u'}_\theta[R_\theta(\hat{\theta}) + H_\theta]dp_\theta. \tag{5}$$

where $\overline{u'}_\theta$ denotes the expected marginal utility of consumption for an individual of type θ:[9]

$$\begin{aligned}\overline{u'}_\theta &= E\left[u'_\theta(c(\theta, A))\right] \\ &= p_\theta(\hat{\theta}_\theta, A)\, u'_\theta(c^1_\theta) + (1 - p_\theta(\hat{\theta}_\theta, A))\, u'_\theta(c^0_\theta)\end{aligned}$$

2.5 Government Problem

Given its information σ and taxpayer self-reports $\hat{\theta}$, the government sets the audit selection rule $A(\hat{\theta}, \sigma)$ in order to maximize a generalized utilitarian social welfare function W. The government places welfare weights on individual utilities denoted ψ_θ.

$$W(A) \equiv \int_\theta \int_\sigma \psi_\theta \left[A(\sigma, \hat{\theta})u_\theta(c^1_\theta(\hat{\theta})) + (1 - A(\sigma, \hat{\theta}))u_\theta(c^0_\theta(\hat{\theta}))\right] dF(\sigma|\theta)dF(\theta). \tag{6}$$

[8] Our characterization of behavioral responses in Section 2.8 relies only on the linearity of R in $T(\hat{\theta})$, i.e. that $dR_\theta = -(1+\rho)dT_\theta$, not the assumption that audits detect all non-compliance. One could obviously allow the marginal penalty ρ to be type-specific. Our examination of the relationship between welfare weights and audit revenues in Section 2.7 does require the assumption that audits detect all non-compliance. In that context, this assumption rules out, for instance, the possibility that a planner might apply a low welfare weight to a group of individuals whom the planner believes to be very wealthy and non-compliant, but whose noncompliance cannot be detected by audits.

[9] This expression follows from a first-order Taylor approximation of $u_\theta(c)$ around a certainty-equivalent level of consumption, $\tilde{c}(\theta, A)$, such that $u'_\theta(\tilde{c}(\theta, A)) = E\left[u'_\theta(c^a(\theta, A))|\theta\right]$.

Auditing an individual of type θ carries an administrative cost K_θ. We write the government's problem is as follows:

$$\begin{aligned} \max_{A(\hat{\theta},\sigma)} \quad & W(A(\hat{\theta},\sigma)) && (7) \\ \text{subject to} \quad & \text{for every } \theta,\, \hat{\theta} = \hat{\theta}(\theta, A) && \text{(IC)} \\ & G = \int_\theta \int_\sigma \left\{ T(\hat{\theta}) + A(\hat{\theta},\sigma)[R_\theta(\hat{\theta}) - K_\theta] \right\} dF(\sigma|\theta)\, dF(\theta) && \text{(GBC)} \end{aligned}$$

Optimality of the audit selection rule $A()$ in equilibrium requires that no feasible perturbation of A increases social welfare. Characterizing when this condition holds yields a reduced-form sufficient statistics characterization of the optimal audit selection rule.

2.6 Marginal Reforms to the Audit Rule

We next characterize the social welfare effect of marginal perturbations to the audit rule. We begin by noting that in a rational expectations equilibrium, perceived probabilities $p_\theta(\hat{\theta}, A)$ are correct given individuals' information. We apply equation (1) to express social welfare as

$$W(A) = \int_\theta \psi_\theta \left[p_\theta(\hat{\theta}, A) u_\theta(c^1_\theta(\hat{\theta})) + (1 - p_\theta(\hat{\theta}, A)) u_\theta(c^0_\theta(\hat{\theta})) \right] dF(\theta). \tag{8}$$

The welfare effect of any perturbation of the optimal audit rule dA is

$$dW = \int_\theta \psi_\theta dv_\theta dF(\theta) + \lambda dG, \tag{9}$$

where dv_θ is the change in expected private welfare associated with the reform from equation (4), λ is the Lagrange multiplier on the GBC, and dG captures the change in government revenues (including both tax payments and audits) caused by the perturbation dA. Because revenues are rebated lump sum to individuals through the transfer G, at the optimum we have $\lambda = \int_\theta \psi_\theta \overline{u}'_\theta dF(\theta)$, or the expected social value of one dollar distributed equally to the full population.

We express the change in government revenues as the sum of a mechanical and behavioral effect, following Saez (2001):

$$dG = \underbrace{\int_\theta dp_\theta (R_\theta - K_\theta) dF(\theta)}_{dG^M} + \underbrace{\int_\theta [dT_\theta + p_\theta dR_\theta] dF(\theta)}_{dG^B}, \tag{10}$$

where dT_θ and dR_θ capture the response of tax liabilities and audit revenue that result from behavioral responses of type θ, i.e. the change in $\hat{\theta}(\theta, A)$. For now we will group both of these into a generic term representing the overall behavioral effect dG^B, and turn our attention to the direct/mechanical effects.

Following Saez (2001) again, we let $g_\theta = \frac{\psi_\theta EMU_\theta}{\lambda}$ denote the social value of a dollar to a person of type θ relative to the mean social value of the dollar λ. Dividing both sides of equation (9) by

λ (normalizing λ to 1) and using the definitions in equation (10), we obtain

$$\frac{dW}{\lambda} \approx \int_{\theta} dp_{\theta} \left[R_{\theta} - K_{\theta} - g_{\theta}(R_{\theta} + H_{\theta}) \right] dF(\theta) + dG^{B}. \tag{11}$$

Note that in the second term inside the curly brackets, we have the welfare weight g_θ and the private costs imposed by the audit on the taxpayer, $R_\theta + H_\theta$. Because both of these depend on θ, when we evaluate a reform that affects individuals with heterogeneous types – e.g. because the government cannot distinguish between some types with its information – how welfare weights and audit outcomes covary will matter.

Toward our sufficient statistics characterization, we now focus on a specific class of perturbations of the audit rule. Specifically, we consider a perturbation that increases the audit rate uniformly by da for values of $(\hat{\theta}, \sigma)$ in an open interval $\hat{\Theta}_0 \times \Sigma_0$.

$$dA_0 = 1\{\hat{\theta} \in \hat{\Theta}_0\} 1\{\sigma \in \Sigma_0\} da. \tag{12}$$

$$dp_{0,\theta} = 1\{\hat{\theta}_\theta \in \hat{\Theta}_0\} Pr(\sigma \in \Sigma_0 | \theta) da. \tag{13}$$

We think of this perturbation intuitively as an increase in audits of a set of individuals with similar observable information from the government's point of view. The number of additional audits is $N_0\, da$, where N_0 denotes the size of the population for which the audit rate is increased, $N_0 = Pr(\hat{\theta}_\theta \in \hat{\Theta}_0, \sigma \in \Sigma_0)$. We assume that the welfare effect for individuals who exit the set $\hat{\Theta}_0$ by changing their reporting behavior in response to a marginal change in audits dA is second-order.[10] This assumption rules out discontinuities in behavior due to non-convexities in the private objective in (3), when the objective is evaluated at the optimal audit rule A.

When we consider changes in audit rules taking the form in (12), many of the terms in equation (11) can be expressed as expectations conditional on the government's observing $\sigma \in \Sigma_0$ and $\hat{\theta}_\theta \in \hat{\Theta}_0$. Denote these conditional expectations for variable X using the notation

$$\begin{aligned} \overline{X} &= E[X_\theta | \hat{\theta}_\theta \in \hat{\Theta}_0, \sigma \in \Sigma_0] \\ &= \frac{1}{N_0} \int_{\theta \in \Theta_0} \int_{\sigma \in \Sigma_0} X_\theta \, dF(\sigma|\theta) dF(\theta) \end{aligned} \tag{14}$$

where $\Theta_0 = \{\theta \,|\, \hat{\theta}(\theta, A) \in \hat{\Theta}_0\}$. In words, $\overline{X}$ denotes the mean value of X among the marginally audited individuals: those whose self-report $\hat{\theta}$ and signal σ make them a target of our reform to the audit selection rule.

With this notation, we can apply equation (13) to equation (11) to obtain the welfare effect of

[10] A similar convexity assumption is commonly imposed when considering a local perturbation to the income tax schedule, e.g., Saez 2001 p. 218, Mirrlees 1971 pp. 182-183.

the perturbation:

$$\begin{aligned}\frac{dW}{\lambda} &\approx \int_{\theta\in\Theta_0}\int_{\sigma\in\Sigma_0} da\,[R_\theta - K_\theta - g_\theta(R_\theta + H_\theta)]\; dF(\sigma|\theta)\, dF(\theta) + dG^B \\ &= N_0 da\left\{\overline{R} - \overline{K} - \overline{g\,R} - \overline{g\,H}\right\} + dG^B \\ &= N_0 da\left\{(1-\overline{g})\,\overline{R} - \overline{K} - \overline{g}\,\overline{H} - Cov[g_\theta, R_\theta|\hat{\Theta}_0, \Sigma_0] - Cov[g_\theta, H_\theta|\hat{\Theta}_0, \Sigma_0]\right\} + dG^B\end{aligned} \tag{15}$$

where the second expression follows from (14) and the last from the definition of covariance.

Finally, let $\beta_{g,R}$ denote the linear projection of g_θ on R_θ among those affected by the audit perturbation:

$$\beta_{g,R} = \frac{\mathrm{Cov}(g_\theta\,,\,R_\theta\,|\,\hat{\Theta}_0, \Sigma_0)}{\mathrm{Var}(R_\theta\,|\,\hat{\Theta}_0, \Sigma_0)}$$

and similarly for $\beta_{g,H}$. Substituting this definition into equation (15) yields

$$\frac{dW}{\lambda} \approx N_0 da\left\{(1-\overline{g})\,\overline{R} - \overline{K} - \overline{g}\,\overline{H} - \sigma_R^2\,\beta_{g,R} - \sigma_H^2\,\beta_{g,H}\right\} + dG^B \tag{16}$$

where $\sigma_X^2 = \mathrm{Var}(X\,|\,\hat{\Theta}_0, \Sigma_0)$.

Intuitively, $\beta_{g,R}$ captures the relationship between social welfare weights and audit adjustments (including additional revenue collected from the audit), whereas σ_R^2 captures residual variation in audit adjustments within some observationally equivalent category of taxpayers. The greater the ability of the government to distinguish taxpayers based on the revenue that would be collected upon audit, the lower σ_R^2 will be. Whereas $\sigma_R^2 > 0$ by definition, the typical structure of welfare weights implies $\beta_{g,R} < 0$, as can be seen in the following subsection.

Going forward we assume that $\sigma_H^2\beta_{g,H} = 0$, because, in contrast to audit outcomes, most of the determinants of audit costs (e.g., complexity of the tax return) are observable to the tax authority when it makes its auditing decisions. We can see from equation (16) that this assumption is straightforward to relax.

2.7 Predicted Audit Revenue and Social Welfare Weights

We next investigate the covariance between welfare weights g_θ and recovered revenue R_θ. We show that under common assumptions about welfare weights and audit revenues, $\mathrm{Cov}(g_\theta, R_\theta|\hat{\Theta}) < 0$. Intuitively, welfare weights are decreasing in R because a) when the government recovers more revenue than expected, it also discovers that the individual has higher consumption than expected, and b) marginal utility is diminishing in consumption.

More formally, with our generalized utilitarian social welfare function we have welfare weights $g_\theta = \frac{\psi_\theta u'_\theta(\tilde{c}_\theta)}{\lambda}$, where $\tilde{c}_\theta$ is the (certainty-equivalent) consumption of type θ – so that $u'_\theta(\tilde{c}) = \overline{u'_\theta}$.

Recall that we assumed, in Section 2.4, that R_θ is linear and increasing in the extent of underreporting. The following additional assumptions are sufficient to obtain a negative relationship between audit revenues and welfare weights, i.e. to ensure $\beta_{g,R} < 0$.

1. Suppose that the taxpayers full income is taxable and that taxable income is sufficient for tax liability. Formally, we assume can express the tax schedule as $T(z(\theta))$ with $T' > 0$ everywhere.

2. There is no relevant heterogeneity in preferences over consumption conditional on $\hat{\theta}$.[11] Formally, we assume that for any $\hat{\theta}$, there is a utility function $\hat{u}_{\hat{\theta}}(c)$ such that for any θ, $\hat{\theta}_\theta = \theta$ implies $u_\theta(c) = \hat{u}_{\hat{\theta}}(c)$.

3. Utility exhibits diminishing marginal utility over consumption: $\hat{u}''_{\hat{\theta}}(c) < 0$.

4. Non-utilitarian welfarist motives ψ_θ covary (weakly) negatively with R_θ: $Cov(\psi_\theta, R_\theta) \leq 0$.

Assumption 1 controls the relationship between the taxpayers' full income z_θ and the extent of under-reporting. It implies that holding the taxpayer's self-reported type $\hat{\theta}$ fixed, when under-reporting is higher, the taxpayer's full income z_θ is also higher. Then we know that the larger is R_θ, the larger is certainty-equivalent consumption $\tilde{c}_\theta$. Thus holding the government's information fixed, more non-compliant types will tend to have higher consumption.

Assumption 2 ensures that consumption is sufficient for marginal utility over consumption. Assumption 3 implies that higher consumption maps to a lower marginal utility of consumption. Incorporating these assumptions, we now have that more non-compliant types will tend to have lower marginal utility of consumption, holding govenrment information fixed.

This logic establishes that under Assumptions 1-3, $\beta_{R,g} < 0$ for a strictly utilitarian welfare function, i.e. where $\psi_\theta = 1$ everywhere. All that remains is to ensure that the additional component of the welfare weight ψ_θ does not reverse the sign of the welfare weights, which we impose directly in Assumption 4. One intuitive departure from strict utilitarianism, consistent with Assumption 4, would be to impose directly that discovering that a taxpayer is noncompliant causes the planner to wish to apply a lower welfare weight, holding true income fixed. [12]

We find that under standard conditions, the conditional covariance between R_θ and g_θ from equation (15) is negative. Holding fixed mean predicted revenue from some observationally equivalent category of taxpayers, higher variance of underreporting therefore tends to *increase* the marginal value of an audit and thus the optimal audit rate of taxpayers in that category. This statement should not be misconstrued as implying that wider variance, i.e. more uncertainty about audit outcomes, improves social welfare. Holding mean predictions fixed, wider variance increases the optimal audit rate but worsens social welfare. Making the government less uncertain by including additional information in σ would also modify both mean predictions and the covariance term in equation (15). We formalize the welfare effect of changes in uncertainty below and confirm that more information always improves social welfare.

[11] Note that we have already imposed that consumption utility does not depend directly on σ.

[12] Welfare weights that depend directly on non-compliance are non-standard, but we frequently hear arguments suggestive of compliance-based welfare weights from tax authorities and others. For example, there is much public discussion of minimizing the "no-change rate" (see e.g. Howard et al., 2021), which would appear to come from a particular motive not to subject compliant taxpayers to the burden of an audit.

2.8 Behavioral Response Elasticites

We next turn our attention back to the behavioral revenue effect dG^B and characterize it in terms of elasticities of behavioral response (as in Keen and Slemrod, 2017).

We first note that for the class of reform described by (12), we have

$$\begin{aligned} dG^B &= \int_{\theta \in \Theta_0} \{dT_\theta + p_\theta dR_\theta\} dF(\theta) \\ &= \int_{\theta \in \Theta_0} \{[1 - p_\theta(1+\rho)] dT_\theta\} dF(\theta) \end{aligned} \tag{17}$$

where the second equality uses the assumption that the penalty is linear in the amount of under-reported taxes.

We next define the type-specific elasticity of reported taxes with respect to the perceived audit probability:

$$\varepsilon_\theta = \frac{dT_\theta}{dp} \frac{p_\theta}{T_\theta}.$$

Applying this definition to Equation (17), we have:

$$\begin{aligned} dG^B &= \int_{\theta \in \Theta_0} \{[1 - p_\theta(1+\rho)] \frac{dT_\theta}{dp} dp_\theta\} dF(\theta) \\ &= \int_{\theta \in \Theta_0} \{[1 - p_\theta(1+\rho)] \varepsilon_\theta \frac{T_\theta}{p_\theta} dp_\theta\} dF(\theta) \\ &= (N_0\, da)\, E\left[(1 - p_\theta(1+\rho))\, \varepsilon_\theta \frac{T_\theta}{p_\theta} dp_\theta | \Theta_0, \Sigma_0\right] \\ &= (N_0\, da)\, [1 - \overline{p}(1+\rho)] \frac{\overline{T}}{\overline{p}} \bar{\varepsilon}, \end{aligned} \tag{18}$$

where $\bar{\varepsilon}$ is a weighted average of individual-specific elasticities ε_θ, with the weights accounting for the importance of deterrence effects for each type θ in the marginally audited group:

$$\bar{\varepsilon} \equiv E\left[\varepsilon_\theta \frac{T_\theta}{\overline{T}} \frac{dp_\theta/p_\theta}{\overline{dp}/\overline{p}} \middle| \Theta_0, \Sigma_0\right].$$

Specifically, we weight the elasticity by 1) the incidence of the perturbation dA on type θ's audit risk, $\frac{dp_\theta/p_\theta}{\overline{dp}/\overline{p}}$, and 2) the revenue importance of type θ within the audited group, $\frac{T_\theta}{\overline{T}}$.[13]

Returning to equation (16) with this characterization of the deterrence term (assuming $\sigma_H^2 \beta_{g,H}$=0

[13] Here we are using the notation introduced earlier, $\overline{p} = E[p_\theta|\Theta_0, \Sigma_0]$, and $\overline{T} = E[T_\theta|\Theta_0, \Sigma_0]$. Both of these are evaluated at the status quo audit rule and type θ's self report $\hat{\theta}_\theta(A)$.

as discussed above), we have

$$\frac{dW}{\lambda N_0 da} \approx \underbrace{\overline{R}-\overline{K}}_{\text{Expected Net Revenue per audit}} - \underbrace{\overline{g}[\overline{R}+\overline{H}]}_{\text{Expected private cost per audit}} - \underbrace{\beta_{g,R}\, Var(R_\theta|\hat{\Theta}_0,\Sigma_0)}_{\text{Prediction uncertainty}} + \underbrace{\overline{T}\overline{\varepsilon}\left[\frac{1-\overline{p}(1+\rho)}{\overline{p}}\right]}_{\text{Expected deterrence effect per audit}} \tag{19}$$

2.9 Characterizing the Optimal Policy

At the optimum, we must have $dW \leq 0$ for any feasible perturbation dA. Equation (19) allows us to characterize when optimality holds using sufficient statistics. In further describing optimality, it is convenient notation to group all of the terms of equation (19) except the expected deterrence effect in the last term into a single term capturing all direct effects, which we denote $\overline{D}$. We first consider corner solutions, which imply audit rates at 0 or 1 for a specific group, then address interior solutions.

Corners First, we find some conditions under which the welfare effect of a marginal audit is positive regardless of the audit rate $\overline{p}$, which can only be consistent with optimality if the optimal audit rate is $\overline{p}^* = 1$. From inspecting equation (19), we find

$$\overline{D} > \overline{T}\rho\overline{\varepsilon} \iff \overline{p}^* = 1. \tag{20}$$

Note that provided deterrence elasticities are non-negative, we can only find $\overline{p}^* = 1$ where direct effects are positive ($\overline{D} > 0$). A positive direct effect would be sufficient to imply a 100% optimal audit rate if the behavioral revenue effect were everywhere positive (if $dG^B \geq 0$ everywhere). This is not the case, however, because by deterring evasion, increasing audit rates can reduce penalties that are collected upon audit. However, direct effects can be sufficiently large at a 100% audit rate that reduced penalty collections are not pivotal, and equation (20) characterizes when such a case obtains. Note that this result implies that the optimal audit rate must not be 1 for some value of $\hat{\theta}$ and all values of σ; if so, a taxpayer would know with 100% certainty that they would be audited if they were to report that value of $\hat{\theta}$, and hence would only do so if it coincided with their true θ. In that event, however, $R_\theta = 0$, implying the direct effect of the audit is negative, and hence that $\overline{p}^* \neq 1$.

The deterrence elasticity and the last term in equation (19) are not well-defined at a zero audit rate, so we characterize the corner solution at zero in terms of the direct effect and the semi-elasticity of voluntary tax payments with respect to p, $\eta_\theta \equiv \frac{dT_\theta}{dp}\frac{1}{T_\theta}$. Following a similar set of steps that we used with the elasticity above, we find that $\overline{dT} = \overline{T}\overline{\eta}$, where the type-specific semi-elasticity η_θ is now weighted by $\frac{T_\theta}{\overline{T}}\frac{dp_\theta}{N_0 da}$. We find similarly $\overline{pdR} = -\overline{p}(1+\rho)\overline{T}\overline{\eta}$.

Turning to the optimality of a zero audit rate, first note that the audit revenue effect $\overline{pdR} = 0$ mechanically at a zero audit rate (because $\overline{p} = 0$). Optimality requires that the welfare effect of increasing the audit rate would be negative at $\overline{p} = 0$, so we find[14]

$$\overline{D} < -\overline{T}\overline{\eta} \iff \overline{p}^* = 0, \tag{21}$$

This expression is intuitive under the condition that audits deter non-compliance rather than encouraging it i.e. $\overline{\eta} \geq 0$.[15] With this condition, a zero audit rate can only be optimal when direct effects are negative, and in fact direct effects must be sufficiently negative that the positive deterrence motives do not dominate at $\overline{p} = 0$.

Equation (21) is suggestive of some intuitive cases in which we could meet the condition for a zero audit rate to be optimal. First, the population could be highly or even perfectly compliant, so that the revenue component of the direct effect would be small and the direct effect mainly consists of costs, and the deterrence effect could be small, because there is little to no non-compliance to deter. In this case, audits are an inefficient use of resources. Note that these are optimal audit rates conditioned on individual behavior, so this case could be true, for example, with wage earners whose income is completely third-party reported and who are highly compliant in equilibrium. Second, a population might be highly non-compliant but audits might be a poor tool to address their non-compliance. In this case, the *audit-recoverable* taxes in the direct effect would again be small (relative to audit costs) and audits would have weak deterrence effects, so we could have a zero audit rate at the optimum. Third, a zero audit rate could essentially be driven by high welfare weights on a non-compliant group. Recall that if $\overline{g} > 1$, the direct effect is *decreasing* in audit-recovered revenues, so a marginal audit that recovers more revenue is worse for welfare, holding all else fixed. In this case, we could have that direct effects are so large and negative as to dominate even a potentially significant deterrence effect, implying that a zero audit rate is optimal. We note that the case in which a zero audit rate emerges as a consequence of high welfare weights on some group of taxpayers, is more likely to arise when we use utilitarian welfare weights and consider audits for the bottom of the income distribution, while it is less likely to arise with welfare weights that depend more directly on compliance.

Interior Optimum At an interior optimum i.e. for $\overline{p} \in (0, 1)$, we must have that the expression in (19) equals 0. Exploring this condition, we obtain an implicit characterization of the optimal audit rate $\overline{p}^*$. At an interior optimum, the optimal audit rate should satisfy the following:

$$\overline{p}^* = \frac{\overline{T}\overline{\varepsilon}}{\overline{T}\overline{\varepsilon}(1+\rho) - \overline{D}}. \tag{22}$$

When neither of the conditions for corner solutions in equations (20) and (21) is satisfied, equation (22) implies $\overline{p}^* \in (0, 1)$.

[14] Recall that $\tilde{D}$ and $\overline{\eta}$ in this expression are evaluated where $\overline{p} = 0$.

[15] Because there can be corner solutions in individual behavior, we should allow for $\overline{dT} = 0$ – e.g. some set of types Θ_0 might be 100% compliant or under-report 100% of their income.

Comparative Statics The optimal audit rate is everywhere increasing in the overall direct effect $\overline{D}$, provided $\overline{T}\overline{\varepsilon} > 0$. Intuitively, this finding implies that the audit rate should, all else equal, be higher for better off groups (who have lower values of g) and for groups where there is higher predicted noncompliance. Meanwhile, the effect on the optimal audit rate of the tax reporting response to audit risk, $\overline{T}\,\overline{\varepsilon}$, depends on the sign of $\overline{D}$:

$$\frac{d\,\overline{p}^*}{d\,\overline{T}\overline{\varepsilon}} = \frac{-\overline{D}}{(\overline{T}\overline{\varepsilon}(1+\rho) - \overline{D})^2}$$

From Equation (19), it is apparent that at an interior optimum, $\overline{D} < 0$ if and only if the behavioral effect of the audit is positive. In turn, the behavioral effect of the audit is positive if and only if $\overline{p} < \frac{1}{1+\rho}$. While not guaranteed for all groups, audit rates above this threshold are likely for all but the most heavily audited taxpayers. For taxpayers facing audit rates above this threshold, the behavioral effect of an increase in the audit rate is *lower* revenue due to a reduction in the yield of previously occurring audits.

Altogether, the above results admit a sufficient statistics test for the optimality of audit selection for *any group of individuals for whom it is feasible to increase audit rates given the government's information*. As with any sufficient statistics characterization, all of the above is an implicit characterization of a local optimum. In the event that the implied optimal audit rate is different from the actual audit rate, we would need to account for how the various terms in equations (19) and (22) change as we move toward the optimum, but under some regularity conditions one could iterate to a fixed point building on the above characterization, as with the implicit characterization of optimal marginal tax rates in Saez (2001). This would give convergence to a local optimum; local optima may not coincide with the global optimum, especially given the high-dimensional decision-making environment. Such are the limitations of the first-order sufficient statistics approach to welfare analysis.

3 The Social Value of New Information for Audit Selection

In this section, we apply the framework developed in the previous section to quantify the welfare effect of changes in the information that the government can observe. For example, the government might gain access to information from new third-party reporting and use this information to improve its prediction technology. We derive sufficient statistics characterizations of the first-order welfare benefits of allowing the government to use such additional information for audit selection. Naturally, the welfare benefits of new information should be traded off against the costs of collecting and incorporating the new information, which we do not directly model here.

Setup Our goal is to quantify the social welfare effect of incorporating some new component σ_n into σ. We denote all components of σ besides σ_n by σ_{-n}. Without new information, we have an equilibrium where the government observes the signal σ_{-n}; with new information, the government observes (σ_{-n}, σ_n).

We derive the social value of the government's observing σ_n in two steps. In both steps, we do our analysis by perturbing an initial equilibrium where the government *does not* select audits using the information in σ_n, i.e. the optimum characterized above in the case where σ_n is not a component of σ. First, we derive the marginal effect on the optimal audit selection rule of allowing the government to observe and to select audits based on the information in σ_n. Second, we derive the first-order welfare effects of a marginal perturbation to the audit selection rule toward the new optimum where σ_n is used optimally in audit selection.

The Marginal Effect of Information on Audit Selection For simplicity, we will focus on cases where $A(\hat{\theta}, \sigma) \in (0, 1)$ for all relevant values of $(\hat{\theta}, \sigma)$, so that equation (22) yields an expression for the optimal audit rate at any $(\hat{\theta}, \sigma)$:

$$A(\hat{\theta}, \sigma) = \frac{T(\hat{\theta})\overline{\varepsilon}(\hat{\theta}, \sigma)}{T(\hat{\theta})\overline{\varepsilon}(\hat{\theta}, \sigma)(1+\rho) - \overline{D}(\hat{\theta}, \sigma)}. \tag{23}$$

Here we are explicitly thinking of the conditional mean of the direct effect and of the behavioral elasticity as functions of the government's information.

Now we express the perturbation of the audit selection rule caused by observation of the information in σ_n as:

$$\begin{aligned} dA(\hat{\theta}, \sigma_{-n}, \sigma_n) &= \frac{\partial A}{\partial \overline{D}} d\overline{D}(\hat{\theta}, \sigma_{-n}, \sigma_n) + \frac{\partial A}{\partial \overline{\varepsilon}} d\overline{\varepsilon}(\hat{\theta}, \sigma_{-n}, \sigma_n) \\ &= \frac{A(\hat{\theta}, \sigma_{-n})}{T(\hat{\theta})\overline{\varepsilon}(\hat{\theta}, \sigma_{-n})(1+\rho) - \overline{D}(\hat{\theta}, \sigma_{-n})} d\overline{D} - \underbrace{\frac{\overline{D}(\hat{\theta}, \sigma_{-n})}{(T(\hat{\theta})\overline{\varepsilon}(\hat{\theta}, \sigma_{-n})(1+\rho) - \overline{D}(\hat{\theta}, \sigma_{-n}))^2}}_{\text{Assumed negligible}} \end{aligned} \tag{24}$$

Here, $d\overline{D}$ describes how the expected direct effect of an audit changes when the government observes σ_n, while $d\overline{\varepsilon}$ describes how behavioral elasticities change when the government observes σ_n.

The change in the direct effect that results from new information is

$$d\overline{D} = E[D_\theta | \hat{\theta}, \sigma_{-n}, \sigma_n] - E[D_\theta | \hat{\theta}, \sigma_{-n}] = \overline{D}(\hat{\theta}, \sigma_{-n}, \sigma_n) - \overline{D}(\hat{\theta}, \sigma_{-n}). \tag{25}$$

Note that by the law of iterated expectations, $E[E[D_\theta | \hat{\theta}, \sigma_{-n}, \sigma_n] | \hat{\theta}, \sigma_{-n}] = E[D_\theta | \hat{\theta}, \sigma_{-n}]$; therefore, $d\overline{D}$ and dA will be mean zero conditional on $(\hat{\theta}, \sigma_{-n})$. This can be viewed as a consequence of optimization under rational expectations: if the government knew that observing σ_n would lead it to increase the average audit rate for some realization of $(\hat{\theta}, \sigma_{-n})$, then it could achieve welfare improvements by increasing overall audit rates at those realizations even when σ_n is unobserved. Using similar notation for the components of $\overline{D}$, we can decompose the change in the direct effect as follows:

$$d\overline{D} = \quad d\overline{R}(1-\overline{g}(\hat{\theta},\sigma_{-n})) - d\overline{H}\overline{g}(\hat{\theta},\sigma_{-n}) - d\overline{g}(\overline{R}(\hat{\theta},\sigma_{-n}) + \overline{H}(\hat{\theta},\sigma_{-n})) - d\overline{K} \\ - \left[\sigma^2_{R|\hat{\theta},\sigma_{-n},\sigma_n}\beta_{g,R|\hat{\theta},\sigma_{-n},\sigma_n} - \sigma^2_{R|\hat{\theta},\sigma_{-n},\sigma_n}\beta_{g,R|\hat{\theta},\sigma_{-n})}\right] \tag{26}$$

We do not further characterize $d\bar{\varepsilon}$; doing so tractably likely requires more structure on deterrence effects. Going forward we assume that $\sigma_n \perp \varepsilon_\theta|(\hat{\theta},\sigma_{-n}) \implies d\bar{\varepsilon} = 0$, and we focus our attention on situations where σ_n primarily conveys information about direct audit outcomes rather than behavioral responses, i.e how predicted revenues, costs, and welfare weights depend on σ_n.

First-Order Welfare Effects Now we find the welfare effects for a perturbation of the audit selection rules in equation (24), i.e. the first-order welfare effects of discriminating on some piece of information σ_n in audit selection.

Using that dA is mean zero conditional on $(\hat{\theta},\sigma_{-n})$, we can express the change in welfare from a marginal perturbation dA as follows:

$$\begin{aligned} dW &= \int_{(\hat{\theta},\sigma_{-n})}\int_{\sigma_n} dA(\hat{\theta},\sigma_{-n},\sigma_n)\left[\int_\theta D_\theta + T(\hat{\theta})\varepsilon_\theta\frac{1-A(\hat{\theta},\sigma_{-n})(1+\rho)}{A(\hat{\theta},\sigma_{-n})}F(\theta|\hat{\theta},\sigma_{-n},\sigma_n)\right] dF(\sigma_n|\hat{\theta},\sigma_{-n})dF(\hat{\theta},\sigma) \\ &= \int_{(\hat{\theta},\sigma_{-n})} Cov\left[dA(\hat{\theta},\sigma_{-n},\sigma_n), \overline{D}(\hat{\theta},\sigma_{-n},\sigma_n) + T(\hat{\theta})\bar{\varepsilon}(\hat{\theta},\sigma_{-n},\sigma_n)\frac{1-A(\hat{\theta},\sigma_{-n})(1+\rho)}{A(\hat{\theta},\sigma_{-n})}\middle|\hat{\theta},\sigma_{-n}\right] dF(\hat{\theta},\sigma_{-n}) \end{aligned} \tag{27}$$

As we assume σ_n is orthogonal to deterrence effects, it must also be the case that $Cov(d\overline{D},\bar{\varepsilon}|\hat{\theta},\sigma_{-n}) = 0$ – if these two do covary, the information in σ_n would also be relevant for the expected deterrence response. Using the law of iterated expectations again, we find that

$$\begin{aligned} Cov[d\overline{D}(\theta,\sigma_{-n},\sigma_n),\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)|\hat{\theta},\sigma_{-n}] &= \int_{\sigma_n}\left(\left[\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n) - \overline{D}(\hat{\theta},\sigma_{-n})\right]\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)\right)dF(\sigma_n|\hat{\theta},\sigma_{-n}), \\ &= \int_{\sigma_n}\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)^2 dF(\sigma_n|\hat{\theta},\sigma_{-n}) - \left(\int_{\sigma_n}\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)dF(\sigma_n|\hat{\theta},\sigma_{-n})\right)^2 \\ &= Var[\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)|\hat{\theta},\sigma_{-n}] \end{aligned}$$

Substituting our expression for dA from equation (24) into the above expression for dW and using this characterization of $Cov[d\overline{D}(\theta,\sigma_{-n},\sigma_n),D(\hat{\theta},\sigma_{-n},\sigma_n)|\hat{\theta},\sigma_{-n}]$, we find

$$dW = \int_{\hat{\theta},\sigma_{-n}}\frac{A(\hat{\theta},\sigma_{-n})}{(1+\rho)\bar{\varepsilon}(\hat{\theta},\sigma_{-n}) - \overline{D}(\hat{\theta},\sigma_{-n})}Var[\overline{D}(\hat{\theta},\sigma_{-n},\sigma_n)|\hat{\theta},\sigma_{-n}]dF(\hat{\theta},\sigma_{-n}) \tag{28}$$

Discussion Equation (28) characterizes the first-order effect of new information on audit selection when that new information is primarily informative about direct welfare effects of audits. Intuitively, we find that the welfare effect is proportional to the amount of uncertainty in the direct effects that is resolved by observing σ_n. Note that we can decompose the variance of the type-specific welfare effects D_θ conditional on $(\hat{\theta}, \sigma_{-n})$ using the law of total variance as follows

$$Var[D_\theta|\hat{\theta}, \sigma_{-n}] = Var[\overline{D}(\hat{\theta}, \sigma_{-n}, \sigma_n)|\hat{\theta}, \sigma_{-n}] + E_{\sigma_n}[Var[D_\theta|\hat{\theta}, \sigma_{-n}, \sigma_n]|\hat{\theta}, \sigma_{-n}] \tag{29}$$

The first term in the above expression is the (conditional) variation in the direct effect that is explained by σ_n. Equation (28) reveals that for a given realization of $(\hat{\theta}, \sigma_{-n})$, the welfare effect of new information is proportional to this explained variance.

4 A Dynamic Extension

In this section, we present a dynamic extension of our theoretical framework. We use a dynamic model to incorporate the dynamic effects of audit selection on the information of individuals and the government. We present a dynamic model in which, as in Boning et al. (2023, Appendix B), audits reveal common information to audited individuals and the government (e.g. the government knows after auditing an individual that they had significant income from self-employment in the audited period, and the audited individual knows the government knows this). Building on the intuition from our reduced-form characterization of welfare with this information environment, we discuss some other possibilities for dynamic effects of audits with alternative information structures, i.e. those in which the informational effects of a marginal audit is not specific to the audited individual.

4.1 Setup

Information Each individual has an audit history as of period t, $\alpha_t = \{a_0, a_1, ...a_{t-1}\}$ that represents whether they were audited in previous periods. The audit selection rule in period t now depends on the audit history and is denoted $A_t(\hat{\theta}_t, \sigma_t, \alpha_t)$. One can think of a taxpayer's prior reports, e.g. $\hat{\theta}_{t-1}$, as a common component of the government's information σ_t and the individual's information θ_t.[16]

Individual type and government information evolve over time according to a generic process that is common knowledge.[17] Similarly to before, individuals of type θ_t form beliefs about their audit probability in time t according to $p_{\theta_t}(\hat{\theta}_t, A_t) = \int_{\sigma_t} A_t(\hat{\theta}_t, \sigma_t, \alpha_t) dF(\sigma_t|\theta_t, \alpha_t)$.

[16] Implicitly, when we condition beliefs on α_t, we are also conditioning on the *findings* of the audit. Setting things up this way allows us to more derive the forward-looking welfare effects of audits in simple terms. A more structural characterization of the importance of one's audit history for behavior should probably impose some structure how future detection probabilities and penalties depend on the findings of audits, not just whether an audit occurred. One can also think of the audit history as a common component of θ and σ but we introduce it as a separate variable for analytical convenience.

[17] Formally, we are assuming that at any t, the distribution $F(\theta_t, \sigma_t|(\theta_0, ..., \theta_{t-1}), (\sigma_0, ..., \sigma_{t-1}))$ is common knowledge.

Individuals At the start of period t individuals choose a report $\hat{\theta}_t$ in order to maximize

$$U_{\theta_t,t}(\hat{\theta}_t, \alpha_t, A_t) = p_{\theta_t}(\hat{\theta}, \alpha_t, A_t)[u_{\theta_t}(c^1_{\theta_t,t}(\hat{\theta}, \alpha_t)) + V_{\theta_t}(1, \hat{\theta}_t)] + (1 - p_{\theta_t}(\hat{\theta}, A_t))[u_{\theta_t}(c^0_{\theta_t,t}(\hat{\theta}, \alpha_t)) + V_{\theta_t}(0, \hat{\theta}_t)]. \tag{30}$$

Most of the terms in this expression are period-t analogues to the terms in equation (3). The sole new term is $V_{\theta_t}(a_t, \hat{\theta}_t)$, which is a continuation value:

$$V_{\theta_t}(a_t, \hat{\theta}) = \max_{\{\hat{\theta}_s\}_{s=t+1}^{\infty}} \sum_{s=t+1}^{\infty} \beta^s \int_{\theta_s} p_{\theta_s}(\hat{\theta}_s, A_s) u_{\theta_s}(c^1_{\theta_s,s}(\hat{\theta}, \alpha_s)) + (1 - p_{\theta_s}(\hat{\theta}_s, A_s)) u_{\theta_s}(c^0_{\theta_s,s}(\hat{\theta}, \alpha_s) dF(\theta_s | \theta_t, \alpha_s(a_t)), \tag{31}$$

where $\alpha_s(a_t) = \{a_0, ..., a_t, ..., a_{s-1}\}$ is the audit history in period s; our notation makes it explicit that α_s, and thus $V_{\theta_t}(.)$, depends on whether an audit occurred in $t < s$, i.e. on a_t.

Government The government selects an audit selection rule $A_t(\hat{\theta}_t, \sigma_t, \alpha_t)$ in each period in order to maximize social welfare, subject to incentive compatibility and a dynamic government budget constraint. As in equation (8), we leverage rational expectations to express the social welfare function as[18]

$$W(\{A_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \beta^t \int_{\theta_t} p_{\theta_t}(\hat{\theta}_{\theta_t}, A_t) u_{\theta_t}(c^1_{\theta_t,t}(\hat{\theta}_{\theta_t}, \alpha_t)) + (1 - p_{\theta_t}(\hat{\theta}_{\theta_t}, A_t)) u_{\theta_t}(c^0_{\theta_t,t}(\hat{\theta}_{\theta_t}, \alpha_t)) dF(\theta_t | \hat{\theta}_t, \sigma_t) \tag{32}$$

The dynamic government budget constraint equates the Net Present Value (NPV) of rebates G_t, denoted $G^N PV$, with the NPV of revenue raised from taxes and penalties.[19]

$$G^{NPV} \equiv \sum_{t=0}^{\infty} \frac{G_t}{(1+r)^t} = \sum_{t=0}^{\infty} \int_{\theta_t} \frac{T(\hat{\theta}_{\theta_t}) + A(\hat{\theta}_{\theta_t}, \sigma_t) R_{\theta_t}(\hat{\theta}_{\theta_t})}{(1+r)^t} dF(\theta_t). \tag{33}$$

4.2 Results

Similarly to before, we characterize the effects of a perturbation to the audit selection rule in a given period, denoted dA_t. Here our characterization requires two distinct applications of the envelope theorem: one for individual re-optimization of their behavior in current and future periods in response to the perturbation ($d\hat{\theta}_{\theta_s}$ for $s \geq t$), and one for the government's re-optimization of audit selection rules in future periods (dA_s for $s > t$). Given both of these effects are second order, we find

[18]Our intial expressions here are slightly different from the static setup above because we incorporate incentive compatibility implicitly from the beginning, to economize on space and notation. Incentive compatibility essentially requires that we use as inputs to these expressions the optimal report of type θ_t in time t, $\hat{\theta}_{\theta_t}$.

[19]We suppose revenues are rebated at the end of the period in which audits occur. However, this is mainly for analytical convenience. We assume rational expectations and do not incorporate inter-temporal consumption-smoothing frictions or heterogeneous discounting, so it follows from the consumption Euler equation that the timing of individuals' receipt of rebated revenues G_t is immaterial for welfare.

$$dW = \int_{\theta_t} \underbrace{\{u_\theta(c^1_{\theta_t,t}(\hat{\theta}_{\theta_t}, \alpha_t)) + V(1, \hat{\theta}_{\theta_t}) - [u_\theta(c^0_{\theta_t,t}(\hat{\theta}_{\theta_t}, \alpha_t)) + V(0, \hat{\theta}_{\theta_t})]\} dp_{\theta_t}}_{\equiv dv^{NPV}_\theta} dF(\theta_t|\hat{\theta}_t, \sigma_t) + \lambda dG^{NPV} \tag{34}$$

We note that equation (34) has a similar structure to equation (9), with the main difference being the presence of the difference in continuation values in the private welfare effect and the use of dG^{NPV} for the fiscal effect.

Leveraging the consumption Euler equation, which intuitively relates the marginal value of a dollar in period t to the expected marginal value of a dollar in future periods, we find an approximation that resembles the one we used in the static model (equation (5)):

$$dv^{NPV}_{\theta_t} \approx EMU_{\theta_t} \Delta C^{NPV}_{\theta_t} dp_\theta, \tag{35}$$

where EMU_{θ_t} is the expected value of the marginal utility of consumption in period t and $\Delta C^{NPV}_{\theta_t}$ is the difference in the expected present value of consumption from time t onwards, for an individual of type θ_t.

Note that obtaining equation (35) requires relatively little structure on consumption in each period, $c^a_{\theta_s}(\hat{\theta}_{\theta_s}, \alpha_s)$. We write the expression this way to emphasize that any effect of audits on the expected PDV of consumption will matter for this private welfare effect, including some potential factors that we disregard here.[20] Supposing that $c^a_{\theta_s}(\hat{\theta}_{\theta_s}, \alpha_s)$ depends on tax payments and the occurrence of audits in the same fashion as above (see equation (3)),[21] we find

$$\Delta C^{NPV}_{\theta_t} = -[R_{\theta_t}(\hat{\theta}_t, \alpha_t) + H_{\theta_t}(\alpha_t)] + \sum_{s=t+1}^{\infty} \int_{\theta_s} \frac{-\Delta_{a_t} T_{\theta_s}(\hat{\theta}_{\theta_s}) - \Delta_{a_t}[p_{\theta_s}(\hat{\theta}_{\theta_s}, \alpha_s) * (R_{\theta_s}(\hat{\theta}_{\theta_s}, \alpha_s) + H_{\theta_s}(\alpha_t))]}{(1+r)^s} dF(\theta_s|\theta_t, \alpha_t), \tag{36}$$

where the difference operator Δ_{a_t} calculates difference between future outcomes if $a_t = 1$ compared to if $a_t = 0$. The first part of this expression resembles the direct effect in the static model. The second part is a forward-looking term arising because the individual's future tax situation can be affected by their audit history, through either future self-reporting decisions or future audit risk. The term includes the expected NPV of $\Delta_{a_t} T_{\theta_s}(\hat{\theta}_{\theta_s})$, which reflects the causal effect of an audit in period t on the audited individual's report of their tax liability in period $s > t$.[22] The

[20] For example, one might include private resource costs the individual incurs to facilitate non-compliance in $\Delta C^{NPV}_{\theta_t}$. If the individual becomes more compliant after being audited, such costs would be reduced in periods $s > t$.

[21] Consumption in t will also depend on saving decisions in the dynamic model, according to a conventional intertemporal private budget constraint. We do not directly model saving, but if saving behavior responds to changes in audit policy this could matter for dG^NPV. Due to the envelope theorem, saving responses will not matter for private welfare.

[22] Our characterization of this effect into welfare is a generalization of the result in Boning et al. (2023, Appendix B), which presents a similar derivation for the case where 1) only audit penalties depend on the audit history, and 2) the probability of detection is exogenous and does not depend on taxpayer reports. The future audit revenues effect is disregarded by Boning et al. (2023); this is understandable because in the environment in which they seek to estimate an MVPF of expanding audits, audit rates are sufficiently low that the effect is likely negligible.

other part of the forward-looking term is an analogous causal effect for future revenues collected via audits, i.e. expected NPV of $\Delta_{a_t}[p_{\theta_s} * (R_{\theta_s} + H_{\theta_s})]$. Intuitively, being audited today shifts the taxpayer into a previously-audited state. The taxpayer's willingness to pay to avoid this state will be the difference in the amount of tax they expect to pay and the difference in their expected audit costs, discounted and aggregated over all future periods.

From here, it is straightforward to follow the same steps as above in order to characterize optimal audits in this dynamic environment.

Discussion From a theoretical perspective, the implications of our dynamic results are straightforward: we should account for forward-looking effects of audits today on individuals' tax payments and revenues in the future, which appear in both the private and fiscal effect because they reflect transfers between individuals and the government, as well as on future audit costs, which appears only in the fiscal effect. It should be possible to find a more general characterization along these lines for more general information structures than we have so far considered. These terms and their dependence on information structure are closely related to the analysis of the causal effects of audits. A number of papers examine the follow-on effects of audits on taxpayer reporting behavior (e.g. DeBacker et al., 2018; Mazzolini et al., 2022; Hebous et al., 2023; Advani et al., 2023). Boning et al. (2023) incorporate forward-looking tax revenue effects of this kind into their MVPF calculations – as this effect appears for both private and fiscal welfare effects, it appears in both the numerator and denominator of their estimated MVPF. Our work here generalizes the model they use to quantify the MVPF along these lines (e.g. we allow for a perceived audit probability that is endogenous to the taxpayer's report). Following similar logic to the above, one would also want to account for covariance between welfare weights and the magnitude of the private forward-looking effects. The effect of an audit today on future audit costs is less commonly examined; it is also negligible when expected future audit rates are sufficiently small.

However, the use of the most commonly studied form of looking effect, the effect of an audit today on the future audited taxpayer's reporting behavior, to evaluate the dynamic welfare effects of audits is closely tied to the structure we imposed on how future information depends on current audits. If no new information arrives between periods (i.e. if everything above that depends on α_t is constant over α_t), then we essentially replay the static model in every period. If, in contrast, audits today provide information about a wider group of taxpayers in future periods – e.g. by revealing information that is useful for training prediction algorithms or assessing compliance across networks of individuals (e.g. K-1 networks, paid preparer networks) – then auditing one taxpayer has broader welfare effects than the ones identified by the post-audit reporting response of the audited taxpayer specifically.

Finally, we note that so far we have ruled out aggregate uncertainty. A different type of dynamic information effect would arise if we relax this assumption: taxpayers or the government may update their beliefs about the distribution $F(\sigma, \theta)$ as a consequence of information revealed by audits. For example, individuals could learn that the government's audit selection criteria are better than they thought at detecting potential evasion, or the government may discover that a group of individuals they thought were compliant is actually engaged in some type of sophisticated tax evasion. For the government, the benefits of learning about the distribution of types from audits resembles the explore side of the oft-discussed exploit-explore tradeoff in predictive modelling. Maintaining tractability and characterizing equilibrium welfare under aggregate uncertainty is difficult, but it is a promising avenue for future research, because it could allow us to describe an exploit-explore tradeoff for social welfare.

5 Conclusion

This paper develops an empirically implementable characterization of optimal audit selection in terms of the mean predictions for audit revenues, private and administrative costs of audits, the elasticity of voluntary tax payments to changes in the probability of audit, and welfare weights, as well as the covariance between welfare weights and private audit effects; all of these are evaluated conditional on the information observable to the government. With empirical data on the aforementioned quantities, it is feasible to assess the optimality, in our model, of expanding audit selection for any group of individuals for whom it is feasible to increase audit rates given the government's information. We hope to implement this chracterization with empirical data in future work. We also characterize the value of new information for audit selection. Taken to data, this prediction would allow researchers to evaluate the welfare effects of targeting audits based on the predictions of machine learning tools, which make use of much richer information than older predictive modelling tools. We also hope to take this theoretical idea to real-world audit data in future work.

Our work here raises a number of questions that could be addressed with future work. First, our characterization of deterrence effects is reduced-form and the elasticity concept we derive as a sufficient statistic is difficult to estimate. One promising avenue for future research here would be to model the deterrence process more structurally and then to use this to pin down deterrence effects in terms of more primitive parameters. Second, as we acknowledged above, our dynamic extension only scratches the surface of possibilities for dynamic models: audits today could provide many types of information to the government that is useful for future audits, not just information about the audited individual. Third, we consider optimal audits holding fixed all other dimensions of the tax system, which includes not only policies like tax rates and the scope for tax avoidance but also audit-specific policies like the auditor's information environment and the type of auditor assigned to a case. Recent research suggests that the type of audit and audit procedures matter, especially for sophisticated high-income individuals (Guyton et al., 2021). In the same spirit as Slemrod and Kopczuk (2002), it would be instructive to incorporate these and other tax policy instruments into a model of optimal audit selection.

References

Advani, Arun, William Elming, and Jonathan Shaw, "The Dynamic Effects of Tax Audits," *The Review of Economics and Statistics*, 2023, *105* (3), 545–561.

Allingham, Michael G. and Agnar Sandmo, "Income tax evasion: a theoretical analysis," *Journal of Public Economics*, 1972, *1* (3), 323–338.

Black, Emily, Hadi Elzayn, Alexandra Chouldechova, Jacob Goldin, and Daniel Ho, "Algorithmic Fairness and Vertical Equity: Income Fairness with IRS Tax Audit Models," in "2022 ACM Conference on Fairness, Accountability, and Transparency" FAccT '22 Association for Computing Machinery New York, NY, USA 2022, p. 1479–1503.

Boning, William, Nathaniel Hendren, Ben Sprung-Keyser, and Ellen M. Stuart, "A Welfare Analysis of Tax Audits Across the Income Distribution," *Working Paper*, 2023.

DeBacker, Jason, Bradley T Heim, Anh Tran, and Alexander Yuskavage, "Once Bitten, Twice Shy? The Lasting Impact of Enforcement on Tax Compliance," *Journal of Law and Economics*, 2018, *61* (1), 1–35.

Elzayn, Hadi, Evelyn Smith, Thomas Hertz, Arun Ramesh, Robin Fisher, and Jacob Goldin, "Measuring and Mitigating Racial Disparities in Tax Audits," *Working Paper*, 2023.

Guyton, John, Patrick Langetieg, Daniel Reck, Max Risch, and Gabriel Zucman, "Tax Evasion at the Top of the Income Distribution: Theory and Evidence," NBER Working Papers 2021.

Hebous, Shafik, Zhiyang Jia, Knut Løyland, Thor O Thoresen, and Arnstein Øvrum, "Do audits improve future tax compliance in the absence of penalties? Evidence from random Audits in Norway," *Journal of Economic Behavior & Organization*, 2023, *207*, 305–326.

Henderson, Peter, Ben Chugg, Brandon Anderson, Kristen Altenburger, Alex Turk, John Guyton, Jacob Goldin, and Daniel E Ho, "Integrating Reward Maximization and Population Estimation: Sequential Decision-Making for Internal Revenue Service Audit Selection," *ACM AAAI*, 2023.

Hendren, Nathaniel and Ben Sprung-Keyser, "A unified welfare analysis of government policies," *Quarterly Journal of Economics*, 2020, *135* (3), 1209–1318.

Howard, Ben, Lucia Lykke, David Pinski, and Alan Plumley, "Can Machine Learning Improve Correspondence Audit Case Selection? Considerations for Algorithm Selection, Validation, and Experimentation," The IRS Research Bulletin, Publication 1500 (Rev. 5-2021) 2021.

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

Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer, "Prediction Policy Problems," *American Economic Review*, May 2015, *105* (5), 491–95.

Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus 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.

Mayshar, Joram, "On Measuring the Marginal Cost of Funds Analytically," *American Economic Review*, 1991, *81* (5), 1329–35.

Mazzolini, Gabriele, Laura Pagani, and Alessandro Santoro, "The Deterrence Effect of Real-World Operational Tax Audits on Self-Employed Taxpayers: Evidence from Italy," *International Tax and Public Finance*, 2022, pp. 1–33.

Rappeport, Alan and Jonathan Weisman, "Democrats, facing a Republican barrage, scale back plans for a crackdown on tax cheating.," *The New York Times*, October 2021.

Reinganum, Jennifer F and Louis L Wilde, "Income tax compliance in a principal-agent framework," *Journal of Public Economics*, 1985, *26* (1), 1–18.

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

Sanchez, Isabel and Joel Sobel, "Hierarchical design and enforcement of income tax policies," *Journal of Public Economics*, 1993, *50* (3), 345–369.

Shaw, Tim, "IRS Funding Cuts Accelerated in FY 2024 $1.6T Spending Deal," *Thomson Reuters Tax & Accounting News*, January 2024.

Slemrod, Joel and Wojciech Kopczuk, "The Optimal Elasticity of Taxable Income," *Journal of Public Economics*, 2002, *84* (1), 91–112.

Yellen, Janet L, "Letter from Secretary of the Treasury Janey L. Yellen to Commissioner of the IRS Charles P. Rettig," August 2022.