Endogenous audits, uncertainty, and taxpayer assistance services: Theory and experiments[1]

Christian A. Vossler

Department of Economics and Howard H. Baker Jr. Center for Public Policy, University of Tennessee, Knoxville, TN 37996, USA

Scott M. Gilpatric

Department of Economics, University of Tennessee, Knoxville, TN 37996, USA

The published version of this study is in the *Journal of Public Economics*. The final publication is available at https://doi.org/10.1016/j.jpubeco.2018.07.009.

[1] Correspondence should be directed to Christian Vossler, Department of Economics, 523 Stokely Management Center, University of Tennessee, Knoxville, TN 37996. E-mail: cvossler@utk.edu. Telephone: 865-974-1699. Fax: 865-974-4601. This research was undertaken in partial fulfillment of IRS contract TIRNO-09-Z-00019. The views expressed are those of the authors and do not reflect the opinions of the IRS or of any researchers working within the IRS. We have, however, benefited from several discussions with Kim Bloomquist, Alan Plumley and Mary Helen Riser at the IRS. We thank Mike McKee for his valuable input and guidance on the experimental design, Jens Schubert for graciously administering a set of experiments at VCU, and Nathan Murray for his research assistance. We further thank Co-Editor Tim Cason and two anonymous reviewers for their detailed and constructive feedback.

Endogenous audits, uncertainty, and taxpayer assistance services: Theory and experiments

Abstract: In recent years there has been a sharp rise in the information available to individual income taxpayers, such as through tax preparation software provided by third parties and support available by tax agencies, but the effects of this information on tax reporting are not well understood. Within a setting characterized by an endogenous audit process and taxpayer uncertainty, this study uses theory and laboratory experiments to investigate the effects of taxpayer assistance services that better inform taxpayers about their tax liability and the audit process. The endogenous audit rule we study is simple, yet relative to existing work is more likely to characterize the actual incentives facing taxpayers. Among our findings, and in contrast to the case of purely random audits, in theory the effect of providing more accurate information on tax liability is ambiguous, and we find support empirically for increased tax underreporting even in a setting where theory predicts the opposite. This unanticipated result is mitigated when services provide better information on both liability and the audit process, suggesting that audit information may be more salient to participants.

JEL Classifications: C91; D8; H24; H26; H83

Keywords: individual income tax; taxpayer assistance services; endogenous audits; tax reporting and enforcement; experimental methods

1. Introduction

In a voluntary reporting tax system, such as the US individual income tax, the opportunity for underreporting liability exists for income sources and deductions that are neither subject to withholding nor to third party reporting, leading to a significant tax "gap".[2,3] Taxpayers often have considerable uncertainty over tax liability for these sources, given complexities of the tax code, imperfect bookkeeping and other factors. Moreover, many taxpayers are poorly informed of audit rules.[4] In response, there has been a sharp rise in available information related to individual income tax reporting, although it remains unclear how such information signals impact compliance. This includes the effects of services provided by tax agencies, such as telephone help lines and internet information documents, as well as information provided by third parties, including tax preparation software (e.g. TurboTax and TaxAct), professional tax preparers, and publications with insight on tax agency operations. For instance, some software is likely to provide signals on how the audit process works through, for example, suggestions of audit flags and information on what like-taxpayers report in charitable contributions. Within a setting characterized by an endogenous audit process and taxpayer uncertainty, this study uses theory and experiments to investigate the effects of taxpayer assistance services that better inform taxpayers about their tax liability and the audit process.

Limited by resource constraints, it has long been recognized that many tax agencies use endogenous rules that base audit chances on taxpayer characteristics and tax reporting behavior. Such rules can be crudely characterized as ones that, for a given peer group of taxpayers (i.e. an

---

[2] Bloomquist et al. (2013) estimates a tax gap of $300 billion for the 2006 tax year ($235 billion in individual income tax and $57 billion in the self-employment tax).
[3] Common examples include tip and self-employment income, capital gains and rental income, and charitable contributions.
[4] The most recent survey evidence we can find for US taxpayers supports this claim, although the evidence is admittedly dated (Louis Harris and Associates, Inc., 1988).

audit class), lead to audit chances that are increasing with the level of noncompliance (Phillips, 2012). Indeed, as audits are costly, there is an efficiency argument for rules that do well in targeting those taxpayers with the highest expected taxes evaded. In turn, when facing an endogenous audit rule, individuals contemplating how much tax to evade may take into account available information signals on how the audit process is perceived to work.

The particular audit rule we develop and examine determines the audit probability, separately for each taxpayer, as a strictly increasing function of one's expected level of tax evasion.[5] Although this rule represents a simplification of actual audit processes, it captures the key incentive present in most endogenous processes in the field and, in expectation, leads to audits that target the worst offenders. Alternatively, some have modelled the audit process as imposing a threshold, where reporting taxes above the threshold results in a zero audit chance. Although this may explain some audit rules in practice, field evidence is largely inconsistent with theoretical predictions from such threshold models (Andreoni, Erard and Feinstein, 1998).

A handful of prior theoretical and experimental studies have examined endogenous audit mechanisms that establish an agent's audit probability based on the observed, relative behavior within a regulated group (Alm and McKee, 2004; Cason, Friesen and Gangadharan, 2016; Gilpatric, Vossler and McKee, 2011) or based on compliance history (Alm, Cronshaw, and McKee, 1993; Cason and Gangadharan, 2006; Clark, Friesen and Muller, 2004). Models dependent on peer evaluations create interactions between taxpayers – e.g. a competition effect to avoid being selected for audit or an incentive to coordinate reports to lower the equilibrium audit probability – and are most relevant when the regulated group is small, which does not

[5] Our discussion of the literature focuses on related work examining endogenous audit rules or the effects of information services on income tax compliance. Of course, the literature on tax compliance (evasion) is vast, and we point the interested reader to Alm (2012), and references therein, for a discussion of the broader literature.

characterize individual tax compliance.[6]

The effects of liability information services have been examined (Alm et al., 2010; Vossler and McKee, 2017), albeit under random audits, with the basic findings that providing more precise information on tax liability leads to less tax underreporting and higher filing rates, both theoretically and experimentally.[7] This work builds on earlier models of compliance behavior given uncertain liability (Beck and Jung, 1989; Snow and Warren, 2005; Evans, Gilpatric, and Liu, 2009). With the endogenous audit rule, and in contrast to the case of purely random audits, providing more accurate tax liability information in theory has an ambiguous effect on tax underreporting. In particular, liability information can alter expected liability, which in turn affects beliefs regarding the audit probability associated with any level of underreporting. For example, when a service alerts the taxpayer that her liability is lower than expected, this increases the expected audit probability tied to a given level of underreporting and thus increases the marginal incentive to comply. Furthermore, the effect of improved information on reporting is ambiguous even if it does not change expected liability but reduces the uncertainty regarding it. When uncertainty is reduced this decreases the expected penalty the taxpayer faces for any given level of underreporting of taxable income, which in turn reduces the value of avoiding an audit. A lower expected penalty conditional on being audited reduces the benefit of lowering the probability of audit by reporting more taxable income, and may lead to reduced compliance.

The effects of providing more accurate information regarding audit probabilities has not

[6] A taxpayer is unlikely to know the size and membership of their audit class, and further – given the large number of taxpayers with similar observable characteristics – it is improbable that any individual taxpayer believes she can actually influence the audit chances of others through her reporting behavior. Models based on compliance history, while they may capture an important tax consideration, introduce a complicated dynamic game between the taxpayer and the regulator. Importantly, aside from using history to apply differential audit efforts to previously compliant and noncompliant taxpayers, audits are randomly determined.

[7] Alm, Bloomquist and McKee (2015) provide a discussion, along with supportive evidence, of the external validity of laboratory experiments on individual income tax compliance.

been previously examined, either in a random or endogenous audit setting. Whereas more accurate tax liability information has a complex and generally ambiguous affect, more accurate audit information has a straightforward impact on predicted behavior. If the information reveals the audit probability to be higher than expected this increases predicted compliance, and if it reveals the audit probability to be lower than expected this reduces predicted compliance. More accurate information that does not change the expected audit probability has no impact on predicted behavior.

We further examine two additional dimensions of taxpayer assistance services. First, we consider the possibility of partial audits, which characterizes a situation where particular tax form line items (e.g., itemized deductions) may be subject to differential scrutiny. If information services reveal which of the two line items, income or deductions, is more likely to be audited, this is predicted to lead to increased compliance in the line with a higher probability of audit, and reduced compliance on the other line. Second, many tax service providers (e.g. H&R Block, TaxSlayer) back their service through guarantees of minimum tax payments (maximum refunds), filing accuracy, and reimbursement of penalties and interest charges. These guarantees thus decrease the expected cost of a tax audit. To characterize this in a stylized way, we consider a guarantee that insulates the taxpayer from any penalties accruing from an audit, conditional on the taxpayer meeting reporting requirements. The effect of a guarantee on reported taxable income is ambiguous theoretically, depending on the relationship between the level of reporting that invokes the guarantee and the report that would minimize the taxpayer's expected costs absent the guarantee. The guarantee may increase reported tax liability if a sufficiently small increase in reporting is required to capture the guarantee. In such a case, the taxpayer benefits from reporting a higher tax liability because, although she incurs a higher tax payment, this is offset by the reduction in possible penalties.

In our experiment, in a setting with deliberate income tax framing, participants face uncertainty over tax liability and the endogenous audit process, and report income and deductions through a simplified tax form. Our experiment is designed to test whether subjects respond as predicted to the endogenous-audit enforcement mechanism, and the various aspects of taxpayer information services as described above within that mechanism. Thus we test whether subjects respond with an increase in reported taxable income to either an increased audit probability *ceteris paribus* or an increased sensitivity of the audit probability to their report. From a baseline of predicted underreporting, we test the effect of a liability information service that reduces the variance of possible tax liability while holding constant the expected value.[8] The model predicts that subjects will increase reported taxable income in this case. We test whether subjects respond as predicted in a partial audit environment to information regarding which line item is more likely to be audited. Finally, we test whether subjects respond to a service guarantee by increasing reported taxable income as the theory predicts in the setting we employ.

The experimental results serve to both confirm and challenge the theoretical model. The most striking finding is that providing more accurate information on tax liability has either a null effect or increases tax underreporting. This contrasts with prior work, which suggests significant reductions in underreporting (Alm et al., 2010; Vossler and McKee, 2017). In fact, one prominent theme of this study is that we have revealed settings theoretically and/or empirically where information services can have undesirable consequences on tax reporting. In the full audit setting, when services provide better information on both liability and the audit process, audit

[8] Ayers, Jackson and Hite (1989) survey third-party tax preparers and find they tend to have pro-taxpayer interpretations of tax regulations. While this may be generally true for third-party information providers, as we are also interested in the effects of (presumably unbiased) information provided by tax agencies, in our design we focus on the provision of unbiased information.

information appears to be more salient to participants, negating the strong effects observed when only liability information is provided. The effects of the audit services are largely in line with theory, in the sense that revealing a higher (lower) audit chance increases reported taxes, and – in the partial audit case – revealing which line item is likely to be audited decreases reporting associated with the untargeted line item. Additional details of our results are described later.

2. Theoretical framework

The basic theoretical framework for tax evasion was derived by Allingham and Sandmo (1972), and follows from the Becker model of crime. This approach views tax evasion as a gamble. In this section, we build upon the theoretical literatures that assume the probability of audit is a decreasing function of reported income (see Andreoni, Erard, and Feinstein, 1998), and considers taxpayer uncertainty (as in Beck and Jung, 1989; Snow and Warren, 2005). The focus of this study is on tax reporting in settings where evasion is most likely to occur, that is, on reporting tied to sources of income and deductions that are subject to little or no third-party reporting requirements. In such settings, taxpayers will naturally have uncertainty over their true liability, given bookkeeping and other recording errors. Uncertainty can also arise given that reporting requirements for some sources, such as the value of donated goods, are tied to estimation methods. Further, uncertainty over liability may arise due to the complexities of the tax code.

Although taxpayers believe their choices of reported income and deductions impact the probability they will be audited, they may nevertheless remain uncertain regarding the audit probability for any given report. This framework will allow us to make predictions about how information services that reduce taxpayer uncertainty about true tax liability or the likelihood of being audited will affect behavior at the individual and aggregate levels.

Consider a taxpayer whose taxable income x is drawn from distribution $F(x)$ with positive density $f(x)$ on the interval $[a, b]$. Importantly, at the time of filing the taxpayer knows the distribution, F, but does not know her realized value, reflecting the uncertainty discussed above regarding true tax liability. Similarly, let the taxpayer's deductions, denoted δ, be drawn from distribution $H(\delta)$ with positive density $h(\delta)$ on the interval $[j, k]$. Again, the taxpayer is uncertain regarding her true deductions and knows only the distribution H from which δ will be drawn. The taxpayer will choose a level of reported income R and deductions D in order to minimize her expected compliance cost, which is the sum of her tax payments and expected penalties. Assuming the taxpayer minimizes expected compliance costs implies risk-neutral behavior. In our framework risk-aversion is not required to reach an interior solution on tax reporting, unlike many tax compliance models in which non-compliance is risky but has a positive expected value at all levels of reporting. In our model the risk-neutral taxpayer reports a level of taxable income such that the benefit of reduced *certain* tax payments is equated to the cost of increased *expected* penalties at the margin. Risk aversion clearly will lead to increased reporting to reduce the exposure to risk. In Section 2.2 we present a variant of the taxpayer's objective employing preferences with constant absolute risk aversion (CARA). In the discussion of our experimental findings we employ the CARA framework to illustrate how increasing risk aversion increases predicted reporting of taxable income. However, the implications of the taxpayer services on behavior can be shown much more clearly in a model that simply assumes the taxpayer minimizes expected compliance costs as we do here.

Denote the tax rate as t and let β be the rate the taxpayer must pay on any liability revealed by an audit. This rate includes payment of additional taxes due plus an additional fine, thus $\beta > t$. The probability the taxpayer is audited is a function of her reported income and deductions, $p(R, D) = c + \gamma G(y - R + D)$. Here $G(y - R + D)$ is a function that captures the change in audit probability relative to the baseline probability c due to the difference between reported taxable income and the reference level, y. This function is scaled by $\gamma > 0$ which parameterizes the responsiveness of the audit probability to the deviation of reported taxable from the reference level. We assume $G(0) = 0$, and that $G' > 0, G'' \geq 0$ for any (R, D) that yields an audit probability between 0 and 1. That is, if reported taxable income equals the reference level y the audit probability is c, and the audit probability decreases with reported taxable income.

The decision problem for a risk-neutral taxpayer can be represented in this general framework as

[1] $$\min_{R,D} t(R - D) + p(R, D)\beta \left[\int_R^b (x - R)f(x)dx + \int_j^D (D - \delta)h(\delta)d\delta\right].$$

The first term in the objective function is the tax on reported income and the second term is the expected additional payments resulting from the possibility of an audit. This optimization problem represents a taxpayer confronted with three distinct sources of uncertainty: taxable income, x, deductions, δ, and whether an audit will occur represented by the probability function $p(R, D)$. The primary effect of liability information services in this framework is to change the taxpayer's belief regarding the distributions from which income and deductions are drawn, $f(x)$ and $h(\delta)$.

An interior solution will hold with $R \in (a, b)$ and $D \in (j, k)$ so long as the baseline audit probability c, penalty rate β, and reference point y are sufficiently large to ensure it is not optimal to report zero taxable income. Given an interior solution the first order conditions identifying cost-minimizing reporting and deductions are respectively

[2] $t + \gamma G'(y - R + D)\beta \left[\int_R^b (x - R)f(x)dx + \int_j^D (D - \delta)h(\delta)d\delta\right]$

$-\beta[c + \gamma G(y - R + D)] \int_R^b f(x)dx = 0$, and

[3] $t - \gamma G'(y - R + D)\beta \left[\int_R^b (x - R)f(x)dx + \int_j^D (D - \delta)h(\delta)d\delta\right]$

$-\beta[c + \gamma G(y - R + D)] \int_j^D h(\delta)d\delta = 0.$

Discussion of second order conditions for a minimum are discussed formally in the online appendix. A sufficient condition for the general satisfaction of the second order condition is that γ is sufficiently small, i.e. the response of the audit probability to reported taxable income is not too severe.

The general framework above shows the tradeoffs that the taxpayer faces in this context. Reporting more taxable income increases the tax payment, but reduces the expected penalty through two channels: a reduced probability of being audited and reduced expected penalty conditional on being audited. Optimal choices equate the increased tax payment and reduced expected penalty at the margin. The remainder of this section characterizes how optimal choices respond to changes in the reporting environment. All proofs appear in the online appendix.

PROPOSITION 1: The optimal reported income increases and the optimal reported deduction decreases when the baseline audit probability, c, increases.

PROPOSITION 2: The optimal reported income increases and the optimal reported deduction decreases when γ increases, i.e. when the audit probability becomes more responsive to reported taxable income.

It also useful to note that if the distributions of income and deductions are identically and symmetrically distributed around their respective means, denoted $\bar{x}$ and $\bar{\delta}$ respectively, then the incentive to underreport income relative to $\bar{x}$ is exactly symmetric with the incentive to over-report deductions relative to $\bar{\delta}$.

PROPOSITION 3: If income and deductions are identically and symmetrically distributed around their respective means, so $f(x-\bar{x}) = h(\delta-\bar{\delta})$ for all x and δ, then underreporting of income is equal to over-reporting of deductions: $\bar{x} - R^* = D^* - \bar{\delta}$.

The following proposition characterizes how optimal reporting responds to a shift in the location of either the distribution of income or deductions in the case when the distributions of income and deductions are identically and symmetrically distributed around their respective means.

PROPOSITION 4: Suppose income and deductions are identically and symmetrically distributed around their respective means, therefore $f(x-\bar{x}) = h(\delta-\bar{\delta})$ for all x, δ. Let R^*, D^* be the optimal reported income and deductions. Further suppose the location of the density from which income is drawn shifts by μ, such that income is drawn from $\tilde{f}(x)$ on $[a+\mu, b+\mu]$, with mean $\bar{x}+\mu$, and with the density from which deductions are drawn unchanged and all other factors constant. Denote the optimal reported income and deductions in this case $\tilde{R}, \widetilde{D}$. Then $R^* < \tilde{R} < \tilde{R} + \mu$; that is, the shift in the density of income increases the reported income, but by less than μ. Underreporting of income, relative to the respective means, thus increases. Reported deductions increases, $\widetilde{D} > D^*$, thus over-reporting of deductions increases. Analogously, if the density from which deductions are drawn shifts by μ holding the density of income and other factors constant, then the reported deduction increases but by less than μ and reported income increases. Over-reporting of deductions and underreporting of income both decrease.

To understand the effect of a change in the variance (uncertainty) of the distribution of either income or deductions it is helpful to start by considering the case when the audit probability is simply exogenous. In that case (corresponding to $\gamma = 0$ in our model) the optimal report on each line is independent of the other and is implicitly defined by the expressions $\frac{t}{c\beta} = \int_R^b f(x)dx = 0$ and $\frac{t}{c\beta} = \int_j^D h(\delta)d\delta$, respectively. These expressions indicate that the ratio of the tax rate to the product of the audit probability and the penalty rate, $\frac{t}{c\beta}$, determines the optimal report on each line as a share of the total distribution. Consider the case when the distributions of income and deductions are identically symmetric around their respective means. Then if $\frac{t}{c\beta} > \frac{1}{2}$ underreporting of income and over-reporting of deductions (relative to the expected value of each distribution) is optimal; if $\frac{t}{c\beta} < \frac{1}{2}$ over-reporting of income and underreporting of deductions is optimal; if $\frac{t}{c\beta} = \frac{1}{2}$ the optimal report is the expected value of each distribution. In this case, if the variance of either distribution is reduced the optimal report moves toward the expected value. In particular, for $\frac{t}{c\beta} > \frac{1}{2}$, if the variance of income is reduced it becomes optimal to underreport less, and if the variance of deductions is reduced is becomes optimal to over-report less.

Now consider the endogenous audit model, but for simplicity consider only one line, income.[9] The interesting question is: can the presence of an endogenous audit mechanism change the result that the optimal report moves toward the expected value when the variance is reduced? In particular, might a reduction in uncertainty *increase* underreporting of taxable income? To see

[9] Assuming only one line is equivalent to assuming income and deductions are identically and symmetrically distributed around their respective means at all times, which of course implies equivalent uncertainty between the two lines.

that the answer is yes, consider an initial circumstance in which it is optimal to underreport taxable income, but by an arbitrarily small amount. In that case, the incentive to increase reported income toward the mean of the distribution when uncertainty is reduced in the random audit setting is negligible. However, the reduction in uncertainty reduces the expected penalty conditional on an audit occurring (holding the report constant). A lower expected penalty reduces the incentive to avoid an audit and thus reduces the incentive to report taxable income in the endogenous audit mechanism.

PROPOSITION 5: Consider the model with a single reporting line, income. If audits are random, $\gamma = 0$, and the initial optimum is characterized by underreporting of taxable income, then if income uncertainty is reduced the optimal report increases. However, when the audit probability is endogenous, $\gamma > 0$, and the initial optimum is characterized by underreporting of taxable income but underreporting is sufficiently small, then if income uncertainty is reduced the optimal report decreases.

2.1. Model structured for experimental application

In order to generate specific predictions regarding the effects of information services on behavior in our experimental framework we will impose additional structure on the model. We assume uncertainty regarding income and deductions is represented by uniform distributions, in particular that $x \sim U[a, b]$ and $\delta \sim U[j, k]$. We also assume that the probability of audit is a linear function of the taxpayer's reported taxable income: $p(R, D) = \gamma(y - R + D) + c$.

With this added structure the taxpayer's decision problem can be represented as

[4] $$\min_{R,D} t(R - D) + (\gamma(y - R + D) + c)\beta \left[\frac{(b-R)^2}{2(b-a)} + \frac{(D-j)^2}{2(k-j)}\right].$$

The first order conditions identifying cost-minimizing reporting and deductions are respectively

[5] $t - \beta\gamma\left[\frac{(b-R)^2}{2(b-a)} + \frac{(D-j)^2}{2(k-j)}\right] - c\beta\left(\frac{b-R}{b-a}\right) - \frac{\gamma\beta(y-R+D)(b-R)}{(b-a)} = 0$, and

[6] $t - \beta\gamma\left[\frac{(b-R)^2}{2(b-a)} + \frac{(D-j)^2}{2(k-j)}\right] - c\beta\left(\frac{D-j}{k-j}\right) - \frac{\gamma\beta(y-R+D)(D-j)}{(k-j)} = 0$.

This model generates point predictions regarding how a taxpayer's behavior will respond to changes in the enforcement environment, including changes in the information available to the taxpayer. Holding constant the expected value of these audit parameters, better information regarding them (i.e. a reduction in uncertainty about either parameter value) has no effect. Therefore, both parameters in this model can be treated as expected values. For example, the taxpayer's optimal choices are identical if she is certain the intercept is 10% or if she believes it is equally likely to be 5% or 15%, thus having an expected value of 10%.

To this point we have modeled an audit as applying to both line items, income and deductions. We also want to characterize a setting where an audit may apply to only one of the line items, and a taxpayer is uncertain about which line will be audited if one in fact occurs. Let θ be the probability that income is audited (conditional on an audit occurring), with $1-\theta$ then the probability that deductions are audited. It will also be convenient for the discussion that follows to add a parameter s that multiplies the audit probability function. Incorporating these yields the following optimization problem:

[4'] $\min_{R,D} t(R-D) + s(\gamma(y-R+D)+c)\beta\left[\theta\frac{(b-R)^2}{2(b-a)} + (1-\theta)\frac{(D-j)^2}{2(k-j)}\right]$.

The formalization of the results in this framework are provided in the online appendix. We show that if the audit probability is doubled but only one line will be audited with either line being equally likely, i.e. when $s=2$ and $\theta=\frac{1}{2}$, then optimal reporting is unchanged from the baseline. However, consider the case when the audit probability is doubled and the taxpayer has some information regarding which line is more likely to be audited. If income and deductions are identically and symmetrically distributed around their respective means then this scenario leads to increased compliance on the line with the higher likelihood of audit and reduced compliance on the line less likely to be audited, as is intuitive.

The final question we address with the theoretical model is the effect of a service guarantee that insulates the taxpayer, conditional on her reporting a certain level of income and/or deductions, from paying anything other than the additional taxes owed (i.e. no penalty is assessed) if an audit occurs. Clearly if a service guarantee is provided it is not optimal to report more income or fewer deductions than required to receive the guarantee. Therefore, if cost-minimizing behavior otherwise entails reporting more income or fewer deductions than required to receive the guarantee, the presence of the guarantee changes optimal behavior. To address whether and under what conditions this occurs, denote the income and deduction reports that solve the optimization problem characterized by equations [5] and [6] to be R^*, D^*, and the levels that receive the guarantee R^G, D^G, with $R^* < R^G$, $D^* > D^G$. Then reporting R^G, D^G results in a lower expected cost for the taxpayer if the following condition holds:

$$[7] \qquad t(R^* - R^G - D^* + D^G) + (\gamma(y - R^* + D^*) + c)\beta\left[\frac{(b-R^*)^2}{2(b-a)} + \frac{(D^*-j)^2}{2(k-j)}\right]$$
$$-(\gamma(y - R^G + D^G) + c)t\left[\frac{(b-R^G)^2}{2(b-a)} + \frac{(D^G-j)^2}{2(k-j)}\right] > 0.$$

The first term in [7] is negative and is the additional tax payment incurred when reporting R^G, D^G rather than R^*, D^*. The second term is the expected additional payments resulting from possible audit (including taxes due and penalties) given reports R^*, D^*, which is positive, and the last term is the expected taxes that may be found due if an audit occurs given that the guarantee does apply. Thus the difference in the last two terms represents expected *penalties* resulting from possible audit, but not additional taxes that may be found. It is optimal to report more income and fewer deductions to obtain the guarantee if the added tax payment is less than the savings from the elimination of expected penalties net of expect tax liability that may still be found by audit. This condition will hold when the gap between the reporting levels required to obtain the guarantee and the levels that would otherwise be optimal is sufficiently small.

2.2. Modeling the impact of risk aversion

To this point, we have modeled a risk-neutral taxpayer who seeks to minimize her expected payments. However, it is of course possible that taxpayers, and our experimental subjects, are risk-averse decision makers. Increased reporting of taxable income reduces a taxpayer's exposure to an uncertain cost (penalties resulting from audit) while increasing a certain cost (tax payments on reported taxable income). Therefore, risk aversion clearly increases reporting of taxable income relative to risk neutral behavior. We employ a model of expected utility maximization with constant absolute risk aversion (CARA) preferences to inform the extent to which the predictions of our model are altered by plausible levels of risk aversions.

We consider a taxpayer with CARA preferences over payoffs of the form $u(\pi) = -exp[-\alpha\pi]$, where the payoff, π, is the individual's income less taxes and penalties. Note that α parameterizes the degree of risk aversion. With income denoted I the individual's expected utility maximization problem is:

[8] $$\max_{R,D} \left(1 - p(R,D)\right)(-exp\{-\alpha[I - t(R-D)]\}) + p(R,D)\int_a^b \int_j^k -exp\left\{-\alpha\left[I - t(R-D) - \beta\left(max(0, x-R) + max(0, D-\delta)\right)\right]\right\}\frac{1}{(b-a)(k-j)}dxd\delta$$

Note that the (R, D) pair that maximizes the consumer's expected utility given CARA preferences does not vary with I because the individual's aversion to risk is constant with respect to her income. Therefore uncertainty regarding I does not affect the optimal choice of (R, D). Although this risk-averse framework does not yield convenient analytical solutions, we can solve computationally to find optimal reporting decisions in our experimental treatments for any degree of risk aversion.

3. Experimental design

3.1 Experiment details

Parameters used for the experiments are reported in Table 1. All amounts are denominated in lab dollars. Lab dollars are converted to US dollars at the end of the session at the rate of 900 lab dollars to one US dollar. A summary of treatment variables, and what is varied within and across sessions, is presented in Table 2. The online appendix includes representative written instructions, and computer screenshots. We first describe common features of all experiment sessions, and then discuss specific treatments. An experimental session consists of 24 decision rounds arranged into four series of six rounds each. All decisions are made using a computer, and participants are paid for each decision round. At the beginning of each series the participants complete a task based on Gill and Prowse (2012), which replicates the exertion of effort outside the laboratory.[10] Participants are placed into three income groups according to relative performance.

In each decision round, participants self-report their tax liability by choosing income and deduction amounts to report, then face the possibility of audit and penalties for underreporting tax liability. Participants have 90 seconds to make these decisions. On the left side of the tax reporting screen appears initial information about the subject's actual income and allowable deduction amounts, and information on the audit process. In the middle of the screen, when available, appears information from liability and audit services.

[10] Participants confront a large number of slider bars and, within sixty seconds, are asked to move as many slider bars as possible to their middle position.

The tax form appears on the right side of the screen. Final tax liability is the difference between earned income and deductions claimed (i.e., taxable income), multiplied by a tax rate of 50%. Participants are free to alter their entries on the tax form up until they file or until the tax form times out. As they adjust their entries they can update their tax form by clicking on a "Do the Math" button. By clicking this button, the reported taxable income and tax payment is displayed below the tax form. For simplicity as well as experimental control, it is not possible to receive – legitimately or through evasion – a tax refund.

Tax liability and endogenous audit probabilities are uncertain to the taxpayer. In terms of liability, the tax reporting screen displays a range of possible income and deduction amounts, respectively, and each amount within each range has an equal chance of being the actual amounts (consistent with the uniform distributions employed in Section 2.1). Participants are required to report income and deduction amounts that are contained within these ranges. The primary purpose of this restriction is to place persons in all income groups on equal footing in terms of the range of possible tax underreporting that is allowed. Otherwise, since tax liability increases with income, those with higher incomes would naturally be able to underreport significantly more. Also, this avoids having to vary the audit function across income groups to account for the differing ranges of possible tax payment amounts. One way to think of the restrictions is that they represent what is known – based on matching documentation – by the tax agency. For example, for a participant with a range of possible income amounts between 2000 and 3000, the tax agency could be considered to know about 2000 in income, such that unmatched income is in the 0 to 1000 range.

An endogenous audit process is conducted independently for each participant. Specifically, as described in Section 2.1, we implement a linear audit probability function. The function is anchored on the maximum probable taxable income, such that if one reports this highest amount they face an expected audit chance equal to the intercept of the function, denoted by c, in the full audit case.[11] This corresponds to the theoretical model when the reference point is $y = b - j$. In the full audit case, for every one-dollar decrease in taxable income, the audit probability increases by an amount equal to the slope of the function, denoted by γ, in the full audit case. This linear relationship extends across the range of possible tax payment amounts, and values for c and γ are chosen to avoid extreme probabilities of 0 and 1.

To describe the audit function, participants are presented with a table (on screen) that maps different tax payment amounts with corresponding audit probabilities.[12] The intercept of the audit function is a random variable, and the audit probability table provides the range of possible audit probabilities associated with different tax payment amounts. Subjects are also provided with the rate of change – the tax payment increase associated with a one percentage point decrease in the audit probability. This information is given to make clear that the audit function is linear and continuous.

Information provided by assistance services appear on the tax reporting screen, prior to the audit process. The focus in the study is on the effects of information services rather than the determinants of information service uptake. Therefore, the notion of information services is present in all treatments, although the actual information provided by the services varies exogenously across treatments. To better induce the notion of information services, regardless of whether they provide any information, subjects always see initial information about their income

[11] The maximum probable tax payment is defined as tax rate multiplied by the difference of the maximum probable income (i.e. upper bound of income range) and the minimum probable deduction (i.e. lower bound of deduction range).
[12] To increase the transparency of the audit function, we elected to frame audit chances as a function of the *tax payment*, rather than as a function of *taxable income*.

and deductions, and the audit process (as described above). This information is displayed on the left side of the screen. When available, the information provided by service(s) is included as additional information, rather than – for instance – a replacement of the initial information set. When first viewing the tax reporting screen, subjects see the initial information set automatically. In the middle part of the screen where service information (may) later appear, subjects see a message, e.g., "tax information service is being requested…please wait". After eight seconds, information from the service (if any) replaces the waiting message(s).

After the tax form is filed an audit determination screen appears. A graphic on the screen consists of three balls in a box, the balls alternate colors (white and blue), and when the balls stop changing color the participant is audited (blue) or not (white). If the player is selected for an audit, any unpaid taxes (based on the actual income and allowable deduction amounts) are discovered and collected along with the penalty. As described above, the actual audit probability is conditional on the submitted tax payment.

After the audit determination, a participant is provided with a summary screen that reveals her actual income and deduction amounts, what she reported on the tax form and a detailed breakdown of earnings from the round.[13] Earnings are determined as the difference between actual income, taxes paid, and any audit costs (unpaid taxes and audit penalty).

3.2 Experiment treatments

Our design exogenously varies two factors within-subject and three factors between-subject. Within a session, each subject encounters each of four tax liability information service

[13] In actuality it is not generally the case that uncertainty over liability is resolved after the tax reporting cycle has ended, especially in the absence of an audit. The focus in the study is on a static tax compliance game and so this feedback is provided in order for earnings calculations to be transparent.

conditions and two audit intensity levels (i.e. two audit function slopes). The order tax liability information conditions are encountered varies randomly across participants in the same session, with the condition switching at the start of each new series. The four liability service conditions are: no information; more precise information about both income and deductions; more precise information about income only; and, more precise information about deductions only. When applicable, the liability service reduces by two-thirds the range of probable income and/or deduction amounts. This is done in a deliberate way. In particular, the service reveals whether the actual amount appears within the bottom, middle, or top one-third of the initial distribution, with the actual amount having an equal chance of being any amount within the tighter range.[14]

Participants within the same session are randomly placed in either the low or high audit intensity regime for the first two series, and then all participants switch to the alternative audit intensity setting for the remaining two series. The two audit intensity levels are determined by varying the audit function slope, γ, to equal 0.0001 (low) or 0.00015 (high). The expected value of the audit intercept parameter, c, is fixed at 0.10. The distribution of c is uniform on the interval [0.05 to 0.15].

There are three factors varied between subjects, each with two levels, giving rise to eight between-subject treatments. Specifically, we vary: whether or not a liability service guarantee is available; whether or not an audit information service is available; and whether, upon audit, one ("partial audit") or both ("full audit") line items are audited. The service guarantee, when active, provides amnesty from audit penalties. The guarantee is naturally tied to tax liability services: it is only in effect when the liability service reveals better information about income and/or deductions. To invoke the guarantee, the participant must report at least as much tax liability as

[14] If we instead just tightened the distribution without ever changing expected value, this would be predictable by participants and the service would cease to have value.

that associated with the midpoint of the tighter income (deduction) interval provided by the service. For instance, if the service reveals that the actual deduction lies within the [0, 333] range, the guarantee is invoked if the participant reports 167 or less. In the event of an audit, although no penalties are levied, the taxpayer is responsible for any unpaid taxes discovered.

When the audit information service is active, it resolves the uncertainty regarding the audit chance associated with a particular tax payment. Without this service, given that the intercept of the audit function is randomly determined, participants see a 10 percentage-point range of possible audit chances for a given tax payment. With the service, the audit intercept is revealed, and participants see the exact audit probability associated with different tax payment amounts. In the partial audit setting, described in more detail below, the service further reveals with 80% accuracy whether an audit would target reported income or reported deductions.

In the full audit setting, for which we set $s = 1$, upon audit both reported income and reported deduction amounts are checked for unpaid taxes. If unpaid taxes are discovered, participants must pay all unpaid taxes along with a penalty (in the absence of a service guarantee). Under the low audit intensity setting, the audit probability is between 5% and 15% when one reports the maximum possible taxable income and between 25% and 35% when one reports the minimum possible taxable income. For the high audit intensity regime, audit chances also begin at 5% to 15% but rise more sharply, leading to a range of 35% to 45% when one reports the minimum possible taxable income amount.

In the partial audit setting, by setting $s = 2$ the audit chances are exactly doubled relative to the full audit case. Under partial audits, either the reported deduction or the reported income is selected for audit, each with an equal chance. Only the selected item is checked for unpaid taxes and associated penalties.

3.3 Testable hypotheses

Our experimental design allows for tests of several hypotheses. We state below as formal hypotheses the theoretical predictions that correspond with the specific parameters used in the experiment.

> **Hypothesis 1**. Tax liability services, holding expected liability constant, increases reported taxable income.
>
> **Hypothesis 2**. The introduction of a service guarantee increases reported tax liability.
>
> **Hypothesis 3**. In the absence of audit information services, doubling the audit probabilities while simultaneously switching to partial audits (in this two line-item setting) has no effect on tax reporting.
>
> **Hypothesis 4**. In the partial audit case, a service that with a high degree of accuracy predicts which line item is targeted for audit will drive individuals to report a higher liability for the targeted item and less liability for the item not targeted.
>
> **Hypothesis 5**. Increasing the intercept of the audit function increases reported taxable income.
>
> **Hypothesis 6**. Increasing the slope of the audit function increases reported taxable income.

Hypotheses 5 and 6 correspond directly with Propositions 1 and 2, respectively. The remaining hypotheses derive from the theory presented in Section 2.1 coupled with our experimental design. For example, theory indicates the effect of tax liability services that reduce uncertainty regarding taxable income is ambiguous in general, but our treatments were designed to generate a predicted increase in reported taxable income as indicate in Hypothesis 1. A sampling of theory point predictions for a taxpayer in the middle income group, under the full audit scenario, are presented in Table 3.[15] The audit settings (intercept and slope) vary across rows and liability information service settings vary across columns. The top panel of the table

[15] Table predictions also correspond with partial audit treatments, with the exception of cases where the audit service is in effect.

presents predictions based on risk neutrality and the bottom panel presents predictions under risk aversion assuming CARA preferences (see Section 2.2). For the latter, we set $\alpha = 0.3$.[16] In the discussion that follows, we focus on predictions under risk neutrality unless otherwise noted.

In the absence of any liability or audit information services, theory predicts that individuals will report less than their expected taxable income. Thus, on average, the prediction is that taxpayers will evade, which reflects the stylized fact from field studies. For example, with $\gamma = 0.0001$ and $c = 0.10$, a middle income taxpayer is predicted to report 978 in taxable income, which is 522 lab dollars lower than her expected, actual taxable income. Tax liability services – holding expected liability constant, which occurs when the service reveals taxable income in the middle range – are predicted to increase significantly reported tax liability (Hypothesis 1).[17] For instance, with the same audit settings, providing better information on either income or deductions increases optimal reporting by 78. Providing better information on both increases reporting by 220. Predictions still indicate underreporting, and as such the introduction of a service guarantee increases reported tax liability (Hypothesis 2). As highlighted by the shaded cells in the table, in many but not all cases the taxpayer is predicted to report a higher taxable income in order to invoke the guarantee.

Unambiguous in theory, in the absence of audit information services, doubling the audit probabilities while simultaneously switching to partial audits (in this two line-item setting) has no effect on tax reporting (Hypothesis 3). In the partial audit case, a service that accurately predicts which line item is targeted for audit will drive individuals to report a higher liability for

[16] In the risk elicitation task, both the median and modal response was to select the sure bet ($2) in the first six scenarios and the lottery in the last four. This implies α is in the interval [0.2028, 0.4236]. To apply this estimate we first convert amounts from lab dollars to USD.

[17] As noted earlier in the theory section, the effect of tax liability services on reported tax liability are ambiguous in general, but our models in Section 2.1 (risk neutral) and Section 2.2 (risk aversion) yields a predicted reporting level for any parameters. We have structured the experiment such that reducing the variance of tax liability—as results from the liability services—increases reported liability.

the targeted item and less liability for the item not targeted (Hypothesis 4). For instance, with $=0.0001$ and $c = 0.10$, and no liability information service, the taxpayer reports 1739 in income and 761 in deductions without the audit service. When the audit service predicts that the partial audit is more likely to check income, she reports 1982 in income and claims a deduction of 1000. As illustrated by this, when the service reveals that income (deductions) will not be the target of any audit, in most cases it becomes optimal to report at or near the highest possible deduction (lowest possible income) amount; for the line item that is targeted, in most cases it becomes optimal to report near the expected true amount.

As unambiguously predicted by the theory, increasing (decreasing) the audit function slope or intercept increases (decreases) reported taxable income (Hypotheses 5 and 6). For example, in the absence of liability information services and with $c = 0.10$, increasing the audit slope from $\gamma = 0.0001$ to $\gamma = 0.00015$ increases optimal reporting by 237. When an audit service reveals the intercept of the audit function, this leads to a wide range of cost-minimizing reports; for instance, with the low audit slope, and no liability services, reporting ranges from 834 ($c = 0.05$) to 1107 ($c = 0.15$) lab dollars.

Last, comparing the top and bottom panels of Table 3 highlights the effects of risk aversion. As can be gleaned from the table, as expected, taxpayers report a higher taxable income under risk aversion in all cases. Hypotheses 2 to 6 continue to hold under risk aversion for our experimental design. However, as reporting under risk aversion is more truthful in the absence of services, or when enforcement is weak, the effects of providing services or increasing enforcement are dampened. The effects of providing liability information services, as previously discussed, are ambiguous in theory. As illustrated by Table 3, although the point predictions for our experiment are unambiguous for the case of risk neutrality, whether information services (holding expected liability constant) increase or decrease reported tax liability for risk averse taxpayers depends on the specific parameter values. Under risk aversion, the effect of liability information services is approximately zero when averaged across all parameter settings.

3.4 Participants and procedures

The experiments were conducted with undergraduate students at the University of Tennessee and Virginia Commonwealth University. Experiments at both locations used identical software, written instructions, and protocols. Recruiting was accomplished using the Online Recruiting System for Experimental Economics (ORSEE) developed by Greiner (2015). The participant databases were built using recruitment posters and flyers, in-class announcements, and course email lists. The experiment was computerized, programmed and conducted with the experiment software z-Tree (Fischbacher, 2007).

An experimental session proceeds in the following fashion. Each participant sits at a computer, and is not allowed to communicate with other participants. An experiment moderator welcomes everyone for their participation, explains that earnings are based on decisions in the experiment, decisions are anonymous, and that experimenter deception is not permitted. Then, the software is initialized and on-screen instructions first guide participants through a set of risk elicitation tasks modeled after Holt and Laury (2002). The experiment moderator answers any questions prior to when decisions are made.

The instructions for the tax experiment are then conveyed by a set of printed instructions that are read aloud to ensure both common knowledge and that the participants at each site received exactly the same instructions. The first practice round (unpaid) takes place as the moderator reads the instructions. This allows participants to see where various pieces of information were located on the screen and to help facilitate understanding of the decision task. Prior to the second practice round, the moderator helps participants work through example scenarios to help make sure that participants understood: how decisions determine the reported tax payment; how the tax payment determines the audit chance; the outcome of the audit process; and how earnings are calculated. Participants undertake a second unpaid practice round, and then have a final opportunity to ask questions.

The participants are informed that all decisions are private; the experimenter is unable to observe the decisions, and the experimenter does not move about the room once the session starts to emphasize the fact that the experimenter is not observing the participants' compliance decisions. This reduces, to the extent possible, peer and experimenter effects. All actions that participants take are made on their computer.

The experiment proceeds for 24 paid decision rounds, although the actual number of rounds is not pre-announced nor is the length of a series.[18] After the final decision round, participants learn of their earnings from the risk elicitation exercise. Participants are then directed to complete both a demographic and taxpayer attitude questionnaire. After the questionnaire is completed, participants are called up to the front of the room individually and paid their earnings in cash.

The last column of Table 2 presents the subject numbers for each of the between-subjects treatments. All treatments are replicated across the two lab locations, and for each treatment, there are four to six sessions and 84 to 96 total participants. Overall, there are 715 participants. The number of subjects per session varied widely due to differential show-up rates and variation

[18] We carried out several pilot sessions to finalize the instructions, software, exchange rates and number of decision periods. As the only difference in the last pilot session was the experiment length (20 periods overall, and five per series), we include this data in the analysis.

in when sessions were scheduled. As the only participant interactions are through the income group determination task, the number of participants in a session should have little if any effect on decisions or outcomes. Average earnings were $32.02 with a range of $14.25 to $47.25. Sessions lasted between 90 and 130 minutes.

4. Results

Descriptive statistics for the experiment data are reported in Table 4. Fifty-two percent of participants can be characterized as risk averse based on the risk elicitation task, 42% of participants are female, and 60% indicate being employed. Given the experimental design, the expected actual taxable income is approximately 1500, and participants can underreport (in expectation) between -1000 to 1000 lab dollars. The mean reported taxable income is 1444, which suggests a modest level of underreporting. As illustrated by Table 3, the level of underreporting is far less than what is predicted by the risk neutral model, which is a common finding in past experimental research on individual income tax compliance.

Table 5 presents tests of the six hypotheses using various data subsamples. For within-subject tests, we use paired *t*-tests. For between-subject tests, we use two-sample *t*-tests that allow for unequal variances. In all cases, test statistics are calculated using individual-specific means. For a clean test of the effects of liability information services (Hypothesis 1), we compare reported taxable income in the presence versus absence of the service using Treatment 1 (T1) data for the full audit setting and Treatment 5 data for a parallel test in the partial audit case. In contrast to the predictions of the risk neutral model, for the full audit setting, liability services have a negative and significant impact of 89.74 lab dollars (this result, and the possibility that it may be explained in part by risk aversion of subjects, is discussed further in Section 5). For the partial audit case, there is a null effect. Including a liability service guarantee increases reporting by 116.69 in the full audit case, supporting Hypothesis 2. The guarantee thus offsets the perverse impact of providing services. The guarantee has no mean effect in the partial audit case, and perhaps this result is due to liability services having no impact in the absence of a guarantee.

As predicted by theory (Hypothesis 3), in the absence of liability information services, there is no significant difference in reported taxable income across full and partial audit settings. Reporting is significantly less under full audits, averaging 154.06, when liability services are implemented. This result reflects the fact that liability services only have a discernable effect in the full audit case. In the partial audit setting, when the audit informs the participant which line item is more likely to be audited, this encourages additional tax underreporting on the non-targeted item as consistent with Hypothesis 4. For instance, when the audit targets reported income, this increases the reported deduction by 116.65 relative to the case where no targeting information is provided. We note, however, that this result must be interpreted with caution as information on the intercept of the audit function and audit targeting is provided simultaneously. Splitting the data to compare observations for which the intercept of the audit function is above the mean (i.e. $c > 0.10$) or not, we find that reported taxable income is higher by 69.31 for the latter group in the full audit case. Increasing the slope of the audit function has a positive and statistically significant effect for both the full and partial audit cases, supporting Hypothesis 6.

Overall, the theory has some predictive power, but many results for the partial audit setting in particular are inconsistent with expectations. The simple tests here of course ignore much of the heterogeneity in the experimental design (e.g. different income groups, treatment interactions), and do not control for participant characteristics. We now estimate a series of econometric models to gain additional insight. Overall, the econometric results support the findings from the hypothesis tests above, and further reveal important treatment interaction effects.

4.1 Econometric analysis

The econometric analysis focuses on the two decision variables in the experiment – reported income and reported deduction – as well as their difference, reported taxable income. The endogenous audit probability is a function of reported taxable income, which creates an association between the two decision variables. We note that the acquisition of services in this study is not a choice variable. Instead, the design exogenously determines the presence or absence of the various services, thus avoiding potential endogeneity issues related to the uptake decision.

In the econometric models, to capture the effects of tax liability information assistance we include the indicators *Income and Deduction Info*, *Income Info Only* and *Deduction Info Only*. The omitted category is no liability information. As the partial resolution of uncertainty can impact expected values (see Proposition 4), we include the continuous variables *Expected Income Change* and *Expected Deduction Change*, which measure the deviations from expected value based on the initial distributions and those provided by the information services. With the latter control variables, the liability service indicators are capturing the effects of reducing uncertainty. To measure the effects of the liability service guarantee, we include the indicator *Guarantee Available*. With this indicator we are capturing the "intent to treat" effect; whether one reports appropriate amounts to actually invoke the guarantee is potentially endogenous. The effects of receiving the audit service, which resolves uncertainty about the intercept of the audit function, is captured though inclusion of the continuous covariates *Audit Intercept Increase* and

Audit Intercept Decrease, which simply measure differences from the expected value of the initial distribution and what is revealed by the service. The indicator *Audit Service* equals one for the audit service treatment. Given the audit change variables, its coefficient can be interpreted as the change in reported liability for a setting where the audit service simply reveals that the expected value of the initial distribution is the actual value. The interactions *Audit Service* × *Liability Service* and *Audit Service* × *Guarantee Available* capture possible service interaction effects. Finally, we include indicators to capture differences due to the audit slope parameter and income group assignment.

Included along with the treatment-related covariates are variables capturing individual characteristics, a linear time trend, two controls for order effects given the audit intensity and the presence/absence of a liability service vary within-subject, and an indicator for experiment lab location. To analyze the three tax reporting outcomes, we estimate ordinary least squares regressions using the experiment panel data. To control for possible heteroskedasticity and within-subject serial correlation, we compute robust standard errors with clustering at the participant level. Further, robust *t* and *F* statistics are used when evaluating hypotheses.

Table 6 and Table 7 present regressions using data from the full audit and partial audit treatments, respectively. Additional regressions are provided in the online appendix.[19] Focusing first on the full audit treatment results for reported taxable income (Model 1), we confirm that there is an unanticipated and perverse effect of liability information services: reducing uncertainty about one's tax liability actually *decreases* reported taxable income. Although the

[19] The appendix provides regressions that parallel those reported here, but without the additional control variables (e.g. risk aversion, gender, etc.). We find that treatment variables that were statistically significant in the prior models remain so here, with the same signs and similar magnitudes. Although the frequency of "corner solution" outcomes is modest, we alternatively estimated the regressions using a two-limit Tobit. The range of possible reported income and deduction amounts is constrained, and 21.7% of income reports, 14.4% and 4.1% of taxable income reports correspond with either the maximum or minimum possible amounts. The estimated marginal effects (evaluated at the sample means of the data), are virtually indistinguishable from the OLS coefficients.

effect of liability services is theoretically ambiguous, with the chosen parameters there is an expected large and *positive* effect of the liability service on reporting for risk neutral taxpayers. As discussed in the theory section, unlike in the case of random audits where the effect of reducing uncertainty is unambiguously positive, reducing the variance (i.e. decreasing the range of probable actual amounts) serves to decrease the expected penalty conditional upon audit as extreme bad outcomes are ruled out. This, in turn reduces the effectiveness of the endogenous audit. Although, as mentioned, this countervailing effect is dominated theoretically with our parameters, the opposite appears to be true behaviorally, which is consistent with risk aversion being a factor. Another surprise is that the magnitude of the effect does not depend on whether uncertainty reduction occurs for only one line item rather than both.

As predicted by theory, when the liability service reveals a higher expected tax liability (higher income or lower deductions), reported taxable income increases. The magnitude of the effect suggests a reporting increase of approximately 50 cents for every one-dollar increase in expected liability. While intuition may suggest that this effect should be 1-to-1, recall that an increase in expected liability means that – holding reported liability constant – one now faces a lower audit probability which of course decreases the incentive to report. This perhaps “weak” response is rather intuitive. Consider what may happen when a taxpayer who donates very little to charity learns that others in her income group donate much more. If she believes that those in her income group face the same audit scrutiny, she would be inclined to report more in donations with little fear of raising an audit flag.

The effect of making available a liability service guarantee is positive and significant, as expected. Importantly, this cancels out the negative effect of the liability service, and even leads to a net increase in reporting. The effects of having the audit information service is asymmetric in the following sense: when the intercept is drawn to be above expected value this has a positive but small effect, but when drawn to be below expected value this motivates a relatively large decrease in reporting. There are interesting effects of having both liability and audit information assistance. In particular, when both service types occur, there is no *overall* impact of liability service information and no impact of having an available guarantee. To see this, note that the interaction *Audit Service* × *Liability Service* is positive and significant. The estimated effect (76.97) is nearly equal in magnitude yet opposite in sign to the effect of having the liability service only (-71.58 to -76.90). The interaction *Audit Service* × *Guarantee Available* is negative and significant. This effect (-109.14) is likewise similar in magnitude yet oppositely signed relative to the guarantee effect when only a liability service is in play (111.68). One possible explanation for both interaction effects is that participants may more heavily focus on the information provided by the audit information service when both are activated.

The effect of increasing the endogenous audit intensity (slope) is an increase in reporting, as expected. This finding is consistent with parallel results from random audit experiments when audit probabilities are increased. The effects of income group membership largely reflect the fact that those with higher incomes owe more in taxes. What is informative, however, are the particular *levels* of the effects. The middle and high income groups face higher taxable incomes of 500 and 1000 respectively. Thus, as both of the coefficients are statistically different from these respective amounts, and are lower, this suggests an important income effect: those with higher incomes report less relative to expectations. Another way to state this effect is that those with higher incomes underreport their taxes more relative to the (expected) truth. That those with higher incomes evade more is consistent with field evidence (Andreoni, Erard and Feinstein, 1998). This type of income effect has further been a consistent finding in the analysis of data from related experimental work, even with experimental designs where the prediction is for those with higher incomes to actually evade less (Vossler and McKee, 2017). As a possible explanation, assuming that risk aversion is a behavioral driver in the context of this tax evasion "gamble", it may be that subjects' utility functions follow decreasing absolute risk aversion such that an increase in wealth leads to an increase in tolerated risk.

Examination of Model 2 (reported income) and Model 3 (reported deduction) reveals similar directional effects of treatment variables in terms of reported tax *liability*. That is, reporting more income has the same effect on liability as does an equal decrease in deductions. There are however a few interesting patterns not discernable from the analysis of taxable income. First, the observed effects of reducing tax liability uncertainty, and having a guarantee available, appear to have been driven almost entirely by changes in reported income. This could be due to increased salience of the tax reporting decision, as it appears first on the tax form. (A related observation is that the goodness-of-fit is dramatically higher for the income reporting model.) Second, when the tax liability information alters expectations regarding income or deductions, this impacts both income and deduction reporting. Reporting on the line item that has not changed in expected value is less dramatic.

Now we turn discussion to the analysis of partial audit treatments, interpreting the regression results reported in Table 7.[20] With the exception of the effects of the audit information service – as it also provides better information about which line item will be audited – theory predicts equivalence across the partial and full audit settings as we implemented them. There are a few qualitative similarities in results: increasing the audit intensity (slope) increases reported liability; changes in expected liability triggered by better liability information motivates

[20] In these regressions, we exclude the service interaction effects, which are jointly insignificant in all three models.

reporting changes in the anticipated direction; and middle and high income group members report less relative to their liability. There are further some noticeable differences. Reducing uncertainty over liability has a null effect on reported taxable income, although remains significant and negative for reported income. There is no impact of further providing a service guarantee. Most prominent is that when the audit service reveals with a high degree of accuracy (80%) which item will be checked upon audit, this nudges taxpayers to report more liability for this item, but also to report less for the other item. This effect is asymmetric: the decrease in reporting is roughly 2.5 times greater on the unlikely targeted amount relative to the increase in reporting on the targeted amount. Overall, the audit service decreases reported taxable income. Theory instead predicts an offsetting effect. The audit service simultaneously resolves uncertainty on the intercept of the audit function. This has an influence on the two reporting items, but the overall effect is only weakly significant. This could reflect the fact that information on which item would be the subject of audit is more salient.

Last, we discuss to the effects of the additional control variables. Risk aversion increases reporting, as expected, but only in the full audit setting. This could be driven by the fact that the range of possible audit penalties, given uncertainty over liability, is reduced when at most one reported amount is checked for accuracy. Consistent with previous tax compliance experiments, females report more in taxes. On average, participants in the VCU lab report more in taxable income. Last, as the game is repeated participants reduce reported taxable income.[21]

[21] We further examined whether there is heterogeneity in treatment effects due to lab location, risk aversion, or income group assignment. To explore lab effects, we expanded Model 1 and Model 4 to allow for a full set of interactions between variables capturing treatment effects and the lab location indicator. We fail to reject that the interactions are jointly significant, suggesting that the main distinction across labs is a shift in mean reported taxable income. Using a parallel approach, we also fail reject that there are important differences in treatment effects due to either risk aversion or income group assignment. Although there is no significant evidence of interaction effects, we suspect that our experimental design may be underpowered for detecting these.

5. Discussion

A key finding of our study is that, in an endogenous audit environment that may better approximate tax compliance settings in the field, when taxpayers have better information regarding their actual tax liability this does not necessarily decrease evasion. This is true in theory, where the effect of better information is ambiguous. Our experiment was designed with parameter values such that risk-neutral theory predicts evasion should decrease, but we instead find that it either increases or remains unchanged (depending on treatment settings). The findings could be driven in part by risk aversion and the salience of extreme outcomes. For a given level of evasion, improved information significantly reduces maximum possible penalties in our experimental framework. This reduces the risk of very negative outcomes from an audit, and risk-averse taxpayers may respond with increased evasion. This is evident in the point predictions in Table 3, where the effects of information services are significantly muted by risk aversion relative to risk neutrality, and in some case the degree of risk aversion represented in the table does predict increased evasion when the information service is provided.

We find that when information services reveal to a participant that her tax liability is higher (lower) than expected she reports more (less), but the increase is not 1-to-1 with the increased expected liability. This finding has implications for classification of taxpayers into peer groups (i.e. audit classes). In particular, this suggests that if a taxpayer has higher expected liability than is typical for her group, and therefore she can engage in greater evasion without triggering a high likelihood of audit because her report will not appear unusually low, she will take advantage of this by increasing evasion. Moreover, *the regulator's knowledge* of a taxpayer's likely true liability is important for deterring evasion. The less effectively the regulator is at grouping taxpayers with similar liability to form peer groups the more this creates opportunities for evasion.

Our main finding in the partial audit setting is that individuals respond to the salience of some aspect of their report being audited with high probability. When the audit information service provides a signal of which line is likely to be audited, participants do report more on the targeted entry but respond primarily by evading more on the line item unlikely to be audited. This suggests that if information services are increasingly able to identify tax line items that rarely receive audit scrutiny, perhaps because the auditor finds it difficult or costly to do so, this information will lead taxpayers to exploit this “enforcement gap”.

One innovation in this study is the development and experimental examination of a simple and transparent endogenous audit rule. Existing theory work in this area has examined endogenous audit rules where audit chances are based on peer evaluations. While such mechanisms may explain well the enforcement realities (or possibilities) for small groups, they are unlikely to capture the individual income tax compliance setting, where the number of regulated agents is very large. Our audit rule captures a setting where a taxpayer’s chance of audit increases with her expected evasion, and as such the mechanism targets the worst offenders, but importantly is independent of the tax reporting of others. The simplicity of the mechanism provides an opportunity to explore the effects of interventions well beyond what is done here.

Given that most prior experimental work in this area has focused on (exogenous) random audits it is important to identify and assess tradeoffs. Many tax agencies that employ endogenous audits, including the IRS, do not freely advertise their audit procedures, especially not with much detail. If one looks at the IRS web page devoted to explaining the audit process (Internal Revenue Service, 2006), even well-educated people are unlikely to learn anything definitive.

Perhaps as a result, most taxpayers are highly uncertain about actual audit probabilities or how changes in their reporting behavior alter the chances they get audited or penalized. Given this uncertainty, it will be natural for some to think that audits are completely random, and others – for instance those using sophisticated tax software or those reading about procedures in books written by former tax agency employees – are likely to view the process as endogenous. With both taxpayer "types" in the field, in our opinion using either type of underlying audit process in the lab is likely to be insightful. One theoretical advantage of the linear endogenous audit function we use is that it avoids the corner solution outcomes (i.e. full compliance or maximal evasion) that frequently arise with the random audit model under the assumption of risk neutrality. This, in turn, allows for a richer variation in predicted outcomes, which is often desirable when identifying treatment effects.

Last, as a first step in analyzing the effects of taxpayer assistance services in a setting with endogenous audit rules, we focused on a setting where taxpayers had unbiased priors, and services provided more accurate information. Our theory and experimental design can be extended in a straightforward way to explore cases where either priors and/or the information set provided by liability or audit information services are biased. For instance, empirical evidence suggests that professional tax preparers are prone to favor the taxpayer (in terms of decreasing reported liability) when interpreting tax regulations. Even if taxpayers are aware of biases associated with tax preparers or third-party software, this nevertheless may have important behavioral effects – taxpayers may elect to proceed under a veil of ignorance or otherwise rationalize underreporting in such settings.

Online Appendix

Supplementary material to this article can be found online at….

References

Allingham, Michael G. and Agnar Sandmo, 1972. Income tax evasion: A theoretical analysis. *Journal of Public Economics* 1 (4), 323-338.

Alm, James, 2012. Measuring, explaining, and controlling tax evasion: Lessons from theory, experiments and field studies. *International Tax and Public Finance* 19, 54-77.

Alm, James, Kim Bloomquist, and Michael McKee, 2015. On the external validity of policy experiments in the lab: Examples from tax compliance experiments. *Economic Inquiry* 53, 1170-1186.

Alm, James, Todd L. Cherry, Michael Jones and Michael McKee, 2010. Taxpayer information assistance services and tax compliance behavior. *Journal of Economic Psychology* 31, 577-586.

Alm, James, Mark B. Cronshaw and Michael McKee, 1993. Tax compliance with endogenous audit selection rules. *Kyklos* 46, 27-45.

Alm, James and Michael McKee, 2004. Tax compliance as a coordination game. *Journal of Economic Behavior and Organization* 54 (3), 297-312.

Andreoni, James, Brian Erard, and Jonathan Feinstein, 1998. Tax compliance. *Journal of Economic Literature* 36 (2), 818-860.

Ayers, Frances, Betty R. Jackson, and Peggy Hite, 1989. The economic benefits of regulation: evidence from professional tax preparers. *The Accounting Review* LXIV (2), 300-312.

Beck, Paul J. and Woon-Oh Jung, 1989. Taxpayer compliance under uncertainty. *Journal of Accounting and Public Policy* 8, 1-27.

Bloomquist, Kim M., Ed Emblom, Drew Johns and Patrick Langetieg, 2013. Estimates of the Tax Year 2006 individual income tax underreporting gap. *IRS Research Bulletin*, Publication 1500, 69-78.

Cason, Timothy N., Lana Friesen and Lata Gangadharan, 2016. Regulatory performance of audit tournaments and compliance observability. *European Economic Review* 85, 288-306.

Cason, Timothy N. and Lata Gangadharan, 2006. An experimental study of compliance and leverage in auditing and regulatory enforcement. *Economic Inquiry* 44, 352-366.

Clark, Jeremy E., Lana Friesen, and R. Andrew Muller, 2004. The good, the bad, and the regulator: An experimental test of two conditional audit schemes. *Economic Inquiry* 42, 69-87.

Erard, Brian and Chih-Chin Ho, 2001. Searching for ghosts: Who are the nonfilers and how much tax do they owe? *Journal of Public Economics* 81, 25-50.

Evans, Mary F., Scott M. Gilpatric, and Lirong Liu, 2009. Regulation with direct benefits of information disclosure and imperfect monitoring. *Journal of Environmental Economics and Management* 57 (3), 284-292.

Fischbacher, Urs, 2007. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics* 10 (2), 171-178.

Gill, David and Victoria Prowse, 2012. A structural analysis of disappointment aversion in a real effort competition. *American Economic Review* 102 (1), 469-503.

Gilpatric, Scott M., Christian A. Vossler, and Michael McKee, 2011. Regulatory enforcement with competitive endogenous audit mechanisms. *RAND Journal of Economics* 42 (2), 292-312.

Greiner, Ben, 2015. Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association* 1, 114 – 125.

Harris, Louis and Associates, Inc., 1988. 1987 Taxpayer opinion survey. Conducted for the U.S. Internal Revenue Service, Internal Revenue Service Document 7292, Washington, DC, January 1988.

Holt, Charles A. and Susan K. Laury, 2002. Risk aversion and incentive effects. *American Economic Review* 92(5), 1644-1655.

Internal Revenue Service, 2006. The Examination (Audit) Process. Accessed July 2018: https://www.irs.gov/newsroom/the-examination-audit-process.

Phillips, Mark, 2012. Reconsidering the deterrence paradigm of tax compliance. *IRS Research Bulletin*, Publication 1500, 99-106.

Snow, Arthur and Ronald S. Warren Jr, 2005. Tax evasion under random audits with uncertain detection. *Economics Letters* 88, 97-100.

Vossler, Christian A. and Michael McKee, 2017. Efficient tax reporting: The effects of taxpayer information services. *Economic Inquiry* 55 (2), 920-940.

Table 1

Experiment parameters.

Parameter / construct	Value(s)
Income	Low Income Group: 1000 to 2000 Middle Income Group: 1500 to 2500 High Income Group: 2000 to 3000
Itemized Deduction	0 to 1000
Audit Probability Function	Endogenous, with probability equal to: $s[\gamma(Max\ Taxable\ Income - Reported\ Taxable\ Income) + c]$
Audit Slope, γ	0.0001 or 0.00015
Audit Intercept, c	0.05 to 0.15
Full Audit	$s = 1$
Partial Audit	$s = 2$
Penalty Rate, β	200%
Tax Rate, t	50%

Note: all indicated ranges correspond with uniformly distributed random variables

Table 2

Experiment treatments.

Treatment	Tax Liability Service	Audit Intensity	Audit Information Service	Liability Service Guarantee	Audit Type	Participants
1	Varies within session		No	No	Full	84
2	Varies within session		Yes	No	Full	88
3	Varies within session		No	Yes	Full	88
4	Varies within session		Yes	Yes	Full	96
5	Varies within session		No	No	Partial	93
6	Varies within session		Yes	No	Partial	96
7	Varies within session		No	Yes	Partial	84
8	Varies within session		Yes	Yes	Partial	86
					Total:	715

Table 3

Selected theoretical predictions, Full Audits.

Audit parameters		Liability Information Services						
γ	c	None	Income only (low)	Income only (middle)	Income only (high)	Income and Deduction (low)	Income and Deduction (middle)	Income and Deduction (high)
Risk Neutrality								
0.00010	0.05	834	681	928	1166	560	1000	1333
0.00010	0.10	978	798	1056	1306	634	1198	1583
0.00010	0.15	1107	899	1167	1428	696	1278	1833
0.00015	0.05	1116	905	1152	1389	698	1244	1666
0.00015	0.10	1215	982	1238	1485	744	1309	1833
0.00015	0.15	1306	1051	1315	1571	784	1364	1910
Risk Aversion								
0.00010	0.05	1088	837	1101	1358	614	1176	1622
0.00010	0.10	1202	930	1203	1471	676	1258	1796
0.00010	0.15	1302	1011	1293	1568	728	1324	1898
0.00015	0.05	1284	1004	1265	1517	728	1290	1802
0.00015	0.10	1366	1068	1337	1598	770	1346	1892
0.00015	0.15	1440	1126	1402	1670	804	1394	1958

Notes: table entries are predictions of reported taxable income for taxpayers in the middle income group. Predictions for the high (low) income group can be obtained by adding (subtracting) 500. The "low", "middle" and "high" scenarios coincide with the cases where the information service identifies the true liability to fall within the bottom, middle and upper third of the uncertainty interval(s), respectively. Following Proposition 3, predictions for cases where the liability service provides better information on only the deduction are identical to the "income only" cases. Shaded cells represent situations where, when the option is available, the taxpayer optimally reports to invoke the liability service guarantee.

Table 4

Data description.

Variable Name	Description	Mean	S.D.
Dependent variables			
Reported Taxable Income	'Reported Income' minus 'Reported Deduction'	1444.108	554.332
Reported Income	Income reported on tax form	1982.476	498.089
Reported Deduction	Deduction reported on tax form	538.368	287.268
Experimental Treatment variables			
Income and Deduction Info	=1 if liability info. service partially resolves uncertainty on both income and deduction	0.250	0.433
Income Information Only	=1 if liability info. service partially resolves uncertainty on income only	0.250	0.433
Deduction Information Only	=1 if liability info. service partially resolves uncertainty on deduction only	0.250	0.433
Liability Service	= 1 if any liability information service is in effect	0.750	0.433
Expected Income Change	Change in expected income, revealed by tax info service	2.061	192.666
Expected Deduction Change	Change in expected deduction, revealed by tax info service	-1.256	192.077
Guarantee Available	=1 if liability info. service guarantee is available	0.372	0.483
Audit Service	=1 if audit information service treatment	0.510	0.500
Audit Intercept Increase	Reported increase in audit function intercept relative to expected value, 0 to 5%; =0 if 'Audit Intercept Decrease' >0 or if 'Audit Service' = 0	0.646	1.340
Audit Intercept Decrease	Reported decrease in audit function intercept relative to expected value, 0 to 5%; =0 if 'Audit Intercept Increase' >0 or if 'Audit Service' = 0	0.628	1.315
Audit Targets Income	=1 if audit info. service reports an 80% chance that income is targeted for audit; =0 if 'Audit Service' =0 or 'Partial Audit'=0	0.124	0.330
Audit Targets Deduction	=1 if audit info. service reports an 80% chance that deduction is targeted for audit; =0 if 'Audit Service' =0 or 'Partial Audit'=0	0.131	0.337

High Income Group	=1 for high income group; \$1500 to \$2500	0.319	0.466
Middle Income Group	=1 for middle income group; \$2000 to \$3000	0.331	0.471
High Audit Slope	=1 if high audit slope, γ=0.00015	0.500	0.500
Partial Audit	=1 if Partial Audit treatment	0.504	0.500
Additional control variables			
Risk Averse	=1 if participant selected the safe option at least six times in the risk elicitation task	0.520	0.500
Female	=1 if participant is female	0.423	0.494
Employed	=1 if participant employed full or part-time	0.602	0.489
Lab Location	=1 if participated at VCU lab	0.297	0.457
Liability Order	=1 if tax liability service in effect at start of session	0.728	0.445
Audit Order	=1 if high audit slope in effect at start of session	0.504	0.500
Round	Round in experiment, 1 to 24	12.514	6.890

Table 5

Hypothesis tests, Reported Taxable Income.

Hypothesis	Subsample 1	Subsample 2	Difference of Means
H1: Liability info. services	T1 | Liability Service = 1	T1 | Liability Service = 0	–89.74**
H1: Liability info. services	T5 | Liability Service = 1	T5 | Liability Service = 0	30.83
H2: Liability info. guarantee	T3 | Liability Service = 1	T1 | Liability Service = 1	116.69*
H2: Liability info. guarantee	T7 | Liability Service = 1	T5 | Liability Service = 1	–46.69
H3: Full versus partial audits	T1 | Liability Service = 0	T5 | Liability Service = 0	–33.50
H3: Full versus partial audits	T1 | Liability Service = 1	T5 | Liability Service = 1	–154.06**
H4: Audit targeting effect[a]	T6 | Audit Targets Income = 1	T5	23.75
H4: Audit targeting effect[a]	T5	T6 | Audit Targets Income = 0	108.12**
H4: Audit targeting effect[b]	T6 | Audit Targets Income = 1	T5	116.65**
H4: Audit targeting effect[b]	T5	T6 | Audit Targets Income = 0	14.59
H5: Audit intensity, intercept	T2, T4 | Audit Intercept Increase > 0	T2, T4 | Audit Intercept Increase ≤ 0	69.31**
H6: Audit intensity, slope	T1, T2, T3, T4 | High Audit Slope = 1	T1, T2, T3, T4 | High Audit Slope = 0	52.58**
H6: Audit intensity, slope	T5, T6, T7, T8 | High Audit Slope = 1	T5, T6, T7, T8 | High Audit Slope = 0	56.46**

*Notes: : * and ** denote differences that are statistically different from zero at the 10% and 5% significance levels, respectively, using t-tests based on individual-specific means. [a] Comparison based on reported income. [b] Comparison based on reported deduction.*

Table 6 Tax reporting models: Full Audit treatments.

	Model 1: Reported Taxable Income	**Model 2: Reported Income**	**Model 3: Reported Deduction**
Liability Service Effects			
Income and Deduction Info	−76.90** (32.74)	−50.30** (20.42)	26.60 (20.46)
Income Information Only	−72.41** (30.82)	−37.73* (19.58)	34.69* (20.21)
Deduction Information Only	−71.58** (30.72)	−62.77** (19.80)	8.81 (20.08)
Expected Income Change	0.55** (0.03)	0.69** (0.02)	0.14** (0.02)
Expected Deduction Change	−0.51** (0.03)	0.13** (0.02)	0.63** (0.02)
Liability Service Guarantee			
Guarantee Available	111.68** (42.74)	72.21** (25.48)	−39.46 (25.91)
Audit Service Effects			
Audit Service	−1.86 (37.95)	−5.68 (26.68)	−3.82 (25.56)
Audit Intercept Increase	10.80** (4.49)	3.54 (3.36)	−7.26** (2.96)
Audit Intercept Decrease	−22.35** (4.48)	−9.83** (3.31)	12.53** (2.79)
Service Interaction Effects			
Audit Srvc. × Liability Srvc.	76.97* (41.42)	56.27** (27.59)	−20.77 (27.24)
Audit Service × Guarantee	−109.14* (60.11)	−78.63** (36.70)	30.53 (34.98)
Other Experiment Treatments			
High Audit Slope	60.22** (13.39)	33.24** (9.33)	−26.98** (8.46)
High Income Group	927.88** (29.14)	974.27** (19.80)	46.39** (17.53)
Middle Income Group	443.59** (25.90)	469.32** (17.02)	25.73* (15.26)
Additional Control Variables			
Risk Averse	78.81** (30.15)	35.51* (18.69)	−43.30** (17.88)
Female	58.10* (31.32)	25.27 (18.74)	−32.83* (18.45)
Employed	−25.64 (30.46)	−15.51 (18.46)	10.14 (18.44)
Lab Location	81.65** (32.26)	26.12 (20.11)	−55.53** (19.14)
Liability Order	−33.18 (33.26)	−13.83 (20.55)	19.35 (19.74)
Audit Order	−46.54 (29.92)	−24.79 (18.61)	21.74 (17.62)
Round	−4.13** (0.99)	−3.21** (0.70)	0.93 (0.61)
Constant	991.87** (50.65)	1534.20** (33.58)	542.32** (30.36)
Number of Observations	8427	8427	8427
F	98.90**	204.10**	79.00**
R^2	0.519	0.696	0.220

*Notes: * and ** denote estimates that are statistically different from zero at the 10% and 5% significance levels, respectively. Standard errors (in parentheses) are clustered at the participant-level.*

Table 7 Tax reporting models: Partial Audit treatments

	Model 4: Reported Taxable Income	**Model 5: Reported Income**	**Model 6: Reported Deduction**
Liability Service Effects			
Income and Deduction Info	−22.89 (17.00)	−31.53** (11.99)	−8.63 (11.49)
Income Information Only	−25.76 (17.06)	−37.37** (12.19)	−11.60 (12.50)
Deduction Information Only	−11.23 (17.99)	−26.23** (12.76)	−15.00 (12.11)
Expected Income Change	0.62** (0.02)	0.74** (0.02)	0.12** (0.01)
Expected Deduction Change	−0.52** (0.02)	0.15** (0.02)	0.66** (0.02)
Liability Service Guarantee			
Guarantee Available	8.02 (25.23)	16.89 (15.66)	8.88 (15.63)
Audit Service Effects			
Audit Targets Income	−53.78* (28.11)	34.38* (17.82)	88.16** (20.71)
Audit Targets Deduction	−49.20* (28.31)	−94.21** (20.97)	−45.01** (17.19)
Audit Intercept Increase	−4.81 (3.48)	−5.61** (2.78)	−0.80 (2.62)
Audit Intercept Decrease	−6.24* (3.76)	−1.21 (2.93)	5.04* (3.05)
Other Experiment Treatments			
High Audit Slope	48.69** (10.73)	21.25** (7.98)	−27.44** (7.49)
High Income Group	910.35** (25.78)	943.03** (16.54)	32.68** (16.12)
Middle Income Group	403.77** (21.85)	447.50** (14.32)	43.73** (14.40)
Additional Control Variables			
Risk Averse	25.61 (24.96)	7.93 (15.61)	−17.68 (15.61)
Female	105.86** (25.82)	71.13** (16.30)	−34.73** (15.86)
Employed	−28.03 (24.52)	−7.23 (16.08)	20.80 (15.42)
Lab Location	84.13** (26.66)	58.40** (16.34)	−25.73 (16.63)
Liability Order	−42.03 (28.22)	−18.60 (16.98)	23.43 (18.13)
Audit Order	−32.00 (26.32)	−4.79 (16.35)	27.21* (15.92)
Round	−3.52** (0.81)	−2.80** (0.59)	0.72 (0.54)
Constant	1078.74** (41.83)	1573.32** (27.98)	494.58** (28.9)
Number of Observations	8559	8559	8559
F	113.78**	262.38**	97.19**
R^2	0.577	0.710	0.253

*Notes: * and ** denote estimates that are statistically different from zero at the 10% and 5% significance levels, respectively. Standard errors (in parentheses) are clustered at the participant-level.*