Optimal Auditing in the Presence of Third-Party Information

Simon Boserup Jori Pinje

May 31, 2010

Abstract

Tax evasion is almost certainly the most widespread form of economic crime: as such, it is of interest in its own right. However, derived from this behavior is also the added distortion of introducing a wedge between the statutory tax system and the "real" or effective tax system. Such distortions, in turn, spur behavior not intended by the statutory tax system, for example by feeding back into labor market choices. Thus, what matters for distributive concerns is not the statutory tax system but the effective average tax rates induced by the audit system which, current theory suggest, may be substantially different. Tax evasion may be a significant source of distortions since a taxpayer, by underreporting income, may lower his expected tax liability and thus his effective average tax rate defined as the ratio of expected tax and penalty payments to true, *ex ante* unobserved, income.

We establish Scotchmer's conjectured relationship in our data set. We find strong evidence that, within small population fractiles, based on the empirical distribution of third-party reported income, effective average tax rates are regressively biased. Conversely, in the variation between fractiles, we find strong evidence of progressive bias. Finally, this is bourne out in the pooled data as an *overall* progressive bias. Our contribution relative to the literature on tax evasion is 1) to establish Scotchmer's conjecture empirically using only raw data and minimal assumptions and 2) to show that this conjecture is strongly featured in a fully specified, state-of-the-art, data-calibrated, workhorse model of optimal auditing and tax evasion. The covariance structure of true income and effective average tax rates in our model is highly robust and we predict that similar empirical relationships would be found in data from any tax auditor that, as the Danish tax agency does, employs a strong signal in predicting true incomes.

1 Introduction

Tax evasion is almost certainly the most widespread form of economic crime: as such, it is of interest in its own right. However, derived from this behavior is also the added distortion of introducing a wedge between the statutory tax system and the “real” or effective tax system. Such distortions, in turn, spur behavior not intended by the statutory tax system, for example by feeding back into labor market choices. Thus, what matters for distributive concerns is not the statutory tax system but the effective average tax rates induced by the audit system which, current theory suggest, may be substantially different. Tax evasion may be a significant source of distortions since a taxpayer, by underreporting income, may lower his expected tax liability and thus his effective average tax rate defined as the ratio of expected tax and penalty payments to true, *ex ante* unobserved, income.

Therefore, an important normative question in the literature on optimal tax auditing is whether, given identical statutory tax rates, effective average tax rates are the same for high-income and low-income individuals. This remains an unresolved issue. Theoretical contributions agree that as true income increases, *ceteris paribus,* so will the difference between the average tax rate intended by the statutory tax system and the effective average tax rate induced by the audit system. In other words, effective tax rates are *regressively biased*. However, in a seminal paper, Scotchmer (1987) points out that tax authorities often have access to data on population characteristics as well as tax returns when conducting audits – if these population characteristics are correlated with true income and the tax authorities use these data in the audit system, it may be that, *overall*, effective average taxes are *progressively biased*.

Whether this is a real phenomenon is necessarily an empirical question since it depends crucially on the distribution of income and the other population characteristics employed by the audit system. As of yet, studies of this kind have not been possible as tax authorities usually are not willing to reveal the guiding principles of their audit systems or what data, if any, they use as signals of true income. Importantly, Scotchmer's conjecture cannot be established in pooled data alone, since it depends crucially on the way in which *between-audit-group variation* dominates *within-audit-group variation.* However, using a recent, unique, Danish data set from Kleven, Knudsen, Kreiner, Pedersen, and Saez (2010), which includes *10,470 individual taxpayer observations* on reported incomes/taxes, post-audit incomes/taxes, income reports made by third parties such as employers and banks, which of taxpayers were flagged for audit by the Danish audit system, as well as other variables, we are able to do exactly this. The unique detail of this data set allows us to calculate approximate effective average tax rates for each individual and separate the covariance of this variable with true income both *within* a set of taxpayers with approximately the *same* level of third-party reported income and *between* sets of taxpayers with *different* levels of third-party reported income.

We establish Scotchmer's conjectured relationship in our data set. We find strong evidence that, within small population fractiles, based on the empirical distribution of third-party reported income, effective average tax rates are regressively biased. Conversely, in the variation between fractiles, we find strong evidence of progressive bias. Finally, this is bourne out in the pooled data as an *overall* progressive bias. Our contribution relative to the literature on tax evasion is 1) to establish Scotchmer's conjecture empirically using only raw data and minimal assumptions and 2) to show that this conjecture is strongly featured in a fully specified, state-of-the-art, data-calibrated, workhorse model of optimal auditing and tax evasion. The covariance structure of true income and effective average tax rates in our model is highly robust and we predict that similar empirical relationships would be found in data from any tax auditor that, as the Danish tax agency does, employs a strong signal in predicting true incomes.

Methodologically, we benefit from two main sources. First, Kleven et al. (2010), analyzes third-party reported income information as a predictor of tax evasion using a highly detailed, stratified random data set of tax returns and third-party information of Danish taxpayers. They find that third-party reported income is a very strong predictor and that it reveals approximately 95% of all income of Danish taxpayers. The tax evasion rate on income reported by third parties is only 0.3%, while for self-reported income the tax evasion rate is 37%. They also analyze the predictive power of other population characteristics such as gender, age, and occupation and, while not insignificant, it does not nearly compare to that of third-party reported information. The Danish tax agency, SKAT, leverages third-party reported information substantially in their efforts to discourage tax evasion.

This suggests that a properly specified theoretical model of optimal auditing must have this feature. To do this, we interpret third-party reported information as a generalized form of *audit groups* – groups which allow a tax agency to condition detection strategies on population characteristics to leverage information – which is an important concept in the existing literature. In this terminology, Scotchmer's conjecture is that effective average tax rates, while being regressively biased *within audit groups*, are progressively biased *between audit groups* to such an extent that pooled data will exhibit progressively biased effective average tax rates.

The second methodological input is that of Erard and Feinstein (1994). They analyze an equilibrium model of optimal auditing and tax evasion within a particular audit group. We generalize their model to describe optimal auditing both within and between audit groups. We model an optimizing tax agency subject to a budget restriction playing a tax evasion/detection game against a population of taxpayers heterogeneous in true income, income reported by third parties, and a simple but necessary behavioral characteristic, honesty.

We calibrate the model's key parameters to data. The main difficulty of applying our model is the necessity of utilizing numerical methods in the determination of equilibrium. Despite the model's parsimony, compared to the complexity of the problem at hand, the covariance structure of the output fits the qualitative features of our observed data well.

We now proceed to the main body of the paper. Section 2 contains a brief review of the relevant literature. Section 3 outlines the Danish tax system, describes the main features of the data, and establishes Scotchmer's conjecture empirically. Section 4 describes our model and the calibration of key parameters. Section 5 outlines the numerical strategy and establishes the correspondence of Scotchmer's conjecture and the model-generated output and Section

6 concludes. Section 7 provides details related to simulations and Section 8 documents the numerical implementation (using MATLAB and STATA).

Terminology

To spare the reader unnecessary repetition, we employ two short-hand conventions throughout the text:

1. The average tax rate implied by the statutory tax system will simply be referred to as the *statutory tax rate* or, symbolically, τ. In the context of a proportional tax system, we will sometimes refer simply to the marginal tax rate, t.

2. The *ex ante* effective average tax rate implied by the audit system (given the statutory tax system) will simply be referred to as the *effective tax rate* or, symbolically, τ^{eff}.

2 Literature Review

Economic literature on tax evasion dates back to Allingham and Sandmo (1972). Their model describes the tax evasion decision of taxpayers faced with a fixed probability of detection and a fine to be paid in case of detection. This seminal work has inspired a large literature on taxpayer behavior.[1] A natural question to ask in light of their model, is how to select the optimal audit strategy in order to maximize a tax agency's objective, usually net revenue given by tax revenue plus fines less audit costs. Early contributions are Reinganum and Wilde (1985, 1986a,b) and Scotchmer (1987) followed notably by Sanchez and Sobel (1993), Erard and Feinstein (1994) and others.

Although the models of optimal auditing differ in many dimensions, for example restricting the budget of tax agencies, allowing the tax agencies to commit to an audit strategy or not, or adding the notion of inherently honest taxpayers, they all have the same qualitative equilibrium property: the optimal audit probability is decreasing in reported income – either continuously decreasing or in a step-wise manner.[2]

The intuition behind this result is simple. The tax compliance game played by the tax agency and taxpayers is essentially a screening problem in which high-income taxpayers can increase their expected payoff by imitating low-income

[1] Making the fine proportional to taxes evaded instead of income evaded (Yitzhaki, 1974), adding endogenous labor supply (Pencavel, 1979, and others), considering a repeated tax compliance game (Engel and Hines, 1999), and many other features have been proposed. See Andreoni, Erard, and Feinstein (1998) and Slemrod (2007) for excellent acounts.

[2] The literature on optimal tax auditing can be divided into two broad groups, those assuming that the tax agency can commit to an audit policy and those who do not. The former uses the principal-agent tools developed by Myerson (1981) and the latter uses Bayesian-Nash methods. While the predictions of the latter group with respect to reporting behavior seem more realistic, e.g. see Andreoni et al. (1998), it is notable that their predictions with respect to the regressivity of effective tax rates within audit groups are the same.

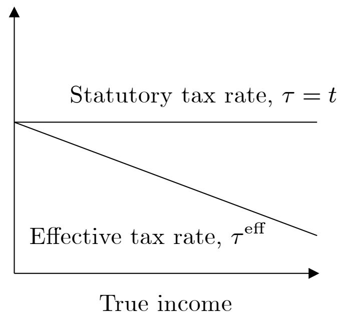

Figure 1. Regressive Bias in the Case of a Proportional Tax, t, and No Audit Groups.

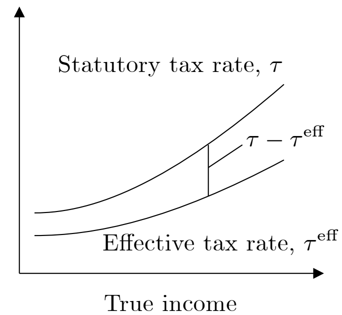

Figure 2. Regressive Bias in the Case of Progressive Taxation and No Audit Groups.

taxpayers. To dissuade tax evasion, the tax agency should apply high audit probabilities to the lowest income reports and lower probabilities to high income reports. Facing a budget restriction, the tax agency ends up auditing none of the high income reports and only some of the low income reports; if this was not the case, an evading taxpayer could lower the probability of getting caught by filing an even lower tax return, which cannot be optimal for the tax agency. Hence, rather than eliminating tax evasion altogether, the goal becomes to discourage very low reports by high-income individuals.

The decreasing audit probability has the direct implication that *ex ante* effective tax rates are regressively biased compared to the statutory tax system. For example, given a constant proportional tax rate, the effective tax rate is decreasing in true income. On average, high-income taxpayers can lower their tax liability via tax evasion more than low-income taxpayers and still face a lower audit risk. Hence, the effective tax rate must be higher for low-income taxpayers. This is illustrated in Figure 1 in the simple case of proportional taxation. The statutory tax rate is constant while the effective tax rate decreases in true income. In the case of progressive taxation, shown in Figure 2, regressive bias is characterized by an increasing difference between the statutory tax rate and the effective tax rate.

This result is very robust, but as Scotchmer (1987) points out, tax agencies do not consider only reported income when choosing their audit strategy. If they have access to more information, e.g. profession, age, or third-party reported income, they can split taxpayers into audit classes or groups, improving the efficacy of deterrence efforts if the grouping provides some signal of taxpayers' true income. Scotchmer suggests that standard models of optimal auditing should be thought of as describing the strategy of the tax agencies *within* audit groups. She demonstrates how applying audit groups to the optimal auditing problem may turn the regressive bias into a progressive one.[3]

[3] Macho-Stadler and Perez-Castrillo (2002) also find that the regressive bias may be overturned, when approaching audit classes by letting the tax authorities receive a signal that is informative of the taxpayer's true income. Special to their analysis is that they include an

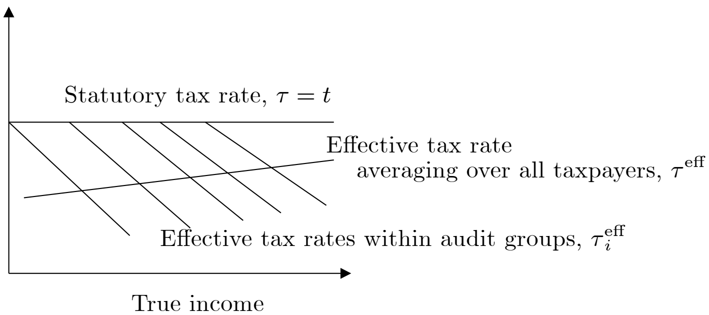

Figure 3. Regressive Bias Within Audit Groups and Progressive Bias Pooling All Taxpayers in the Case of a Proportional Tax, t.

To illustrate the effect of including audit groups, assume that the classification into audit groups conveys a signal of the true income. Then each audit group will span a smaller set of true incomes than the entire set of taxpayers. Averaging over the effective tax rates of individual taxpayers, the effective tax rate is no longer decreasing in true income, cf. 3. In fact, it may be increasing – a progressively biased tax schedule. The mechanism driving the result is that some low-income taxpayers benefit from becoming high-income individuals within their audit group, while some high-income taxpayers now become low-income taxpayers within their audit group. The reclassification changes the risk of being audited and, hence, the effective tax rate, which results in the slope depicted in Figure 3. The regressive bias is a *within*-audit-group property, while the variation in effective tax rates *between* audit groups works in the direction of a progressive bias. The overall effect when pooling all taxpayers together is a combination of the two.

To separate within- and between-variation empirically, one needs detailed data. Kleven et al. (2010) obtains such data in cooperation with the Danish tax agency, SKAT. SKAT performed thorough auditing of a stratified random sample of Danish wage earners' and self-employed taxpayers' tax returns for the year 2006. This data set is unique in that it includes income reports made by third parties, such as employers, banks, etc. on the income earned by the taxpayers. They find that third-party reports, which the Danish tax agency relies heavily on in its tax enforcement strategies, are very effective in the effort to deter tax evasion. In Denmark, 95% of all income is third-party reported and they find that the tax evasion rate is only 0.3% on income subject to third-party reporting while it is 37% on self-reported income. Socioeconomic factors such as e.g. gender, marital status, home ownership are also correlated with tax evasion and, therefore, true income, but the magnitudes are very small in comparison to factors related to the information provided by third-parties.

information asymmetry, such that the taxpayer does not know the realization of the signal, only the distribution of it.

3 Tax Collection in Denmark

SKAT's tax collection efforts rely heavily on information reports by third parties. During some year t incomes are earned and by the end of January in year $t+1$ SKAT receives information reports from employers, banks and other entities. By mid-March, SKAT sends out pre-populated tax returns based on third-party information and other information that they possess about the taxpayers, such as the taxpayers' residence and workplace for calculating commuting allowances. Subsequently, the taxpayers have until May 1 to correct their tax return which can be done by contacting SKAT via telephone, e-mail or internet. In case of no corrections, the pre-populated tax return counts as final.

After the deadline, SKAT's computerized system processes the tax returns, supplying audit flags to returns that the system finds likely to contain errors. The system is entirely deterministic and does not as such assign a probability of audit. After the tax returns have been processed, tax examiners assess the flagged returns and decide whether or not to initiate an audit based on the severity of the different kinds of flags, local knowledge and auditing resources.

If an audit discovers underreporting the taxpayer may pay the taxes owed immediately or postpone the payment at an interest. If the tax examiner views the underreporting as deliberate, the tax agency may impose a fine according to a nonlinear fining scheme depending on the assessed intentionality of the misreporting.[4] In practice, such fines are rare since it is hard to prove that the underreporting is not an honest mistake. In very severe cases, underreporting is punished with imprisonment of up to one and a half years or in some cases up to eight years.

Data

The data originates from an experiment conducted by SKAT in the years 2006–2008, originally analyzed in Kleven et al. (2010). The experiment involved a stratified random sample of 17,764 self-employed individuals and 25,020 employees and recipients of benefits in Denmark. In the present study, we narrow our focus to a subsample of non-treated employees and recipients of benefits and their incomes in the fiscal year 2006. The sample is a stratified random sample of 10,470 selected Danish taxpayers.[5] For each taxpayer, SKAT conducted an unannounced audit after the deadline for changing the tax return (May 1, 2007). The tax audits were comprehensive in the sense that SKAT examined all items on the tax return, demanding documentation for all items about which SKAT

[4] For intentional violations the fine is two times the tax evaded (exceeding 30,000 DKK) and one times the tax evaded below 30,000 DKK. For severe negligence the fine is one times the tax evaded (exceeding 30,000 DKK) and one half times the tax evaded below 30,000 DKK.

[5] Note the randomness of our sample as opposed to tax compliance data obtained from the regular audits that is heavily biased by over-sampling taxpayers who are likely to have misreported their income in either direction. The sampling strategy involved a stratification on tax return complexity. So-called "light" taxpayers with low-complexity tax returns were under-sampled while "heavy" taxpayers with high-complexity tax returns – characterized by having foreign source income – were over-sampled.

did not possess information (SKAT, 2009). SKAT made a significant effort to have tax examiners perform homogeneous audits by e.g. organizing training workshops and distributing detailed audit manuals. The audits took up 21% of the resources devoted to tax audits in 2007 (Kleven et al., 2010).

Of course, it is unlikely that the tax examiners find all hidden income, such as that stemming from cash-only businesses and other black market activities. We focus our attention on the detectable part of tax evasion given the methods available to SKAT and thus denote our empirical counterpart of true income "detectable income". In what follows, we will write true income when in fact we mean detectable income.

For each taxpayer, we have income and tax records on the pre-populated tax return, the final return as potentially changed by the taxpayer, and the post-audit return. In addition, we possess information on the generated audit flags that would normally constitute a basis for selecting taxpayers for audits. We use the pre-populated tax returns as our notion of third-party information.[6]

The Tax System and Tax Compliance in Denmark

The Danish income tax system (in 2006) operates with many different measures of income. Here, we will provide the headlines; for an overview see Table 1 and Kleven et al. (2010) for details. Labor market income, i.e. salary, fringe benefits and other earned income, are taxed proportionally by a labor market tax of 8%, while an earned income tax credit (EITC) of 2.5% is provided for labor market income up to 292,000 DKK.[7] Capital income is a net concept, and different tax rates apply depending on whether net capital income is positive or negative. For most taxpayers net capital income is negative due to interest payments on mortgages. State taxes (bottom, middle and top tax) are levied on the so-called personal income, which, in addition to positive net capital income, consists of labor market income plus social transfers, and pensions less labor market taxes and some pension contributions. The state taxes constitute a progressive tax scheme with a personal allowance and three brackets. Local taxes (county and municipality) are levied on "taxable income" which is similar to the state tax base except that it allows for negative net capital income deductions and other deductions such as transport allowances. In this way, Denmark has a version of the Nordic dual income tax[8] only applying for negative capital income which is taxed at a flat rate while positive capital income is taxed progressively just as regular income. Stock income (dividends and capital gains) is subject to a

[6] There are, however, some instances where the pre-populated returns contain some elements of self-reporting or where the third-party information arrives too late to be included in the pre-populated returns. For example, taxpayers can report changes to SKAT before the pre-populated returns are produced. If they do so by telephone these changes will be counted as performed by SKAT and, thus, part of the third-party information used to pre-populate the tax return even though they are in fact self-reported. But in most cases the difference between items on the final return and the pre-populated return reflect the entire self-reporting behavior.

[7] Approx. 53,000 USD (1 USD = 5.5 DKK).

[8] For a discussion of the Nordic dual income tax., see e.g. Nielsen and Sørensen (1997).

Table 1. An Overview of the Danish Tax System, 2006.

Tax	Tax base	Bracket (DKK)*	Rate (pct.)
Labor market tax	Labor income	none	8.0
EITC	Labor income	up to 292,000	2.5
Bottom tax	Personal income + $\max(\text{capital income}, 0)$	38,500–	5.5
Middle tax	— // —	265,500–	6.0
Top tax	— // —	318,700–	15.0†
Local taxes	Taxable income (= pers. income + cap. income – deductions)	38,500–	33.3‡
Stock income tax	Stock income	0-44,300; 44,300-	28.0; 43.0

*1 USD = 5.5 DKK.
†The top tax rate may be lowered by the "tax ceiling" that limits the sum of state taxes (bottom, middle and top) and local taxes (excl. church taxes) to 59%. In the average municipality the tax ceiling lowers the top rate by 0.08 percentage points.
‡In the avg. municipality and county incl. optional church tax of on avg. 0.74.
Source: Kleven et al. (2010).

two-rate scheme with the high rate setting in at 44,300 DKK.

Table 2 presents some descriptive statistics of the sample on major income components. The table shows sample means with standard errors of means in parentheses – all numbers are calculated accounting for the stratification scheme. Column (1) presents pre-audit figures measured at the deadline, May 1, and column (5) shows figures reported by third-parties. Self-reported figures (the difference between (1) and (5)) is shown in column (6). Negative figures mean that taxpayers on average adjust the number downwards to less than what third-parties have reported. Columns (2)–(4) describe how the figures in (1) were adjusted by the tax examiners during the audit. Columns (3) and (4) split the audit adjustments into positive (meaning underreporting) and negative (meaning overreporting) adjustments, while column (2) holds the average net adjustment, i.e. the sum of (3) and (4).

The top panel of Table 2 shows figures on net income, an income concept not used in the tax system but still informative, and total taxes. The former is defined as the sum of personal income, capital income, stock income, self-employment income and foreign income less deductions. Pre-audit net income is on average a little less than 200,000 DKK with a significantly positive net adjustment from SKAT of almost 1,700 DKK. The positive net adjustment reflects an asymmetry in the reporting behavior with underreporting being more than ten times as high as the overreporting. Third-party reported net income is slightly higher than pre-audit net income mainly due to deductions not included in the third-party reports, implying a negative self-reported net income.

The bottom panel of Table 2 features a decomposition into main income components. The asymmetry in the over- and underreporting found for net income is noticeable for all components. Another interesting fact appears when looking at net adjustments as a percentage of the means. The greatest relative amount of underreporting is found on items the least subject to information reporting and vice versa. Self-employment income tops the list with a net adjustment of 20.4% of the mean self-employment income level followed by stock income (6.6%), foreign income (3.8%) and the rest being between 0.4% and 1.5% in numerical terms. The incomes with the highest share of underreporting are probably also those subject to the most complex tax code, but on average reporting behavior is "in favor" of the taxpayer.

Progressive Bias in the Danish Tax System

Given the calibration exercise to come, it is natural to begin by checking the raw data for evidence of regressive or progressive bias in effective tax rates.

We investigate this in a subset of the population of taxpayers in the data. Specifically, we look at a subset of 890 taxpayers, which we define as potential evaders[9] since for non-evaders the effective tax rate is always the statutory tax rate.

We use SKAT's audit flag system as a proxy for the audit probability function. We denote this the *flag function* and let it take the value of the fraction of flags assigned within an audit group. Audit groups are defined as 1/40th fractiles of the distribution of third-party reported income. The flag system selects a rather large proportion of taxpayers for audit so we will need to scale the flag function so that it corresponds to a realistic budget constraint.[10] We denote this scaled flag function $f(i)$, where $i \in \{1, \ldots, 40\}$ is the index of a particular audit group, ranked by the level of third-party reported income.

In this manner we can calculate the effective tax rate[11] for each individual and each audit group as

$$\tau^{\text{eff}}\left(i, T, \tilde{T}, Y\right) = \frac{f(i)\,T + (1 - f(i))\,\tilde{T}}{Y},$$

[9] See the discussion of parameter Q in Section 4. In anticipation of the discussion, potential evaders are defined by not having their entire income reported by third parties and by not having mistakenly overreported their income.

[10] During normal operations in SKAT, the flag system is utilized in an ad hoc manner: the generated flags automatically opens a tax return for audit – however, whether this particular return is audited at all, in part, or in full depends, to our knowledge, mostly on the particular tax examiner's experience, local knowledge and his ability to identify an *ex ante* suspicious composition of reported income. Of course, SKAT must also satisfy a budget constraint with respect to resources available for audits, but there is no overarching system for prioritizing the audit of a particular return with a regard to maximizing revenue or compliance. Therefore, we resort to a uniform normalization of the flag function to approximate an effective audit function. We chose the normalization such that the average audit probability is the same as in the benchmark calibration of our model.

[11] In this paper, we take the view that the proper definition of the effective tax includes all expected payments, i.e. includes expected penalties of evasion. However, since we do not know what a proper θ would be in this context, we perform the calculation of τ^{eff} setting $\theta = 0$. This does not change the results discussed in this subsection. If we set θ to any reasonable value, e.g. 1.2 as in Erard and Feinstein (1994), or an intermediate value the same relationships are found in the data and the statistical significance of the estimated coefficients are of the same order.

Table 2. Tax Compliance in Denmark, income year 2006

	Pre-Audit Income	Audit Adjustment			Third-Party Reported Income	Self-Reported Income
		Net adj.	Under-reporting	Over-reporting		
	(1)	(2)	(3)	(4)	(5)	(6)
Net Income	193,277	1,664	1,825	-161	193,498	-222
	(1,906)	(480)	(479)	(22)	(1,908)	(105)
Total Tax	63,178	636	695	-59	63,304	-126
	(841)	(246)	(246)	(9)	(843)	(40)
Income components						
Earnings	156,127	672	683	-11	156,096	31
	(2,275)	(203)	(203)	(6)	(2,275)	(50)
Personal Income	209,232	1,137	1,195	-58	209,396	-164
	(1,950)	(480)	(479)	(17)	(1,954)	(62)
Capital Income	-10,884	142	198	-56	-10,975	90
	(272)	(27)	(24)	(11)	(270)	(37)
Deductions	9,264	-143	-213	70	8,949	314
	(178)	(28)	(26)	(11)	(175)	(40)
Stock Income	3,612	239	262	-24	3,484	128
	(546)	(40)	(39)	(10)	(543)	(66)
Self-Employment	103	21	23	-2	103	0
	(60)	(8)	(8)	(1)	(60)	–
Foreign Income	479	-18	6	-25	440	39
	(92)	(19)	(4)	(19)	(90)	(17)

Notes: The sample contains 10,740 taxpayers denoted as employees or recipients of public transfers (unemployed, pensioners, etc.). Due to the stratification strategy employed by SKAT, the sample contains 74.6% "heavy" taxpayers (i.e. with high-complexity tax returns) and 25.4% "light" taxpayers, while the population has 32.6% heavy taxpayers and 67.4% light taxpayers.
Net income is defined as personal income + capital income – deductions + stock income + self-employment income + foreign income. Standard errors of means in parentheses. All estimates are population weighted.
All amounts in DKK (1 USD = 5.5 DKK).

where T is taxes due on true income, $\tilde{T}$ is taxes due on reported income, and Y is true income for an individual in audit group i. Denote $\tau - \tau^{\text{eff}}$ the effective tax rate bias. We define the average tax rate, τ, in the usual way, $\tau = T/Y$.

First, we will check whether the relationship between $\tau - \tau^{\text{eff}}$ and Y is positive *within* audit groups (i.e. for fixed i). This will serve as evidence of the regressive bias predicted by prevailing theory since the difference between the statutory tax rate and the effective tax rate would be higher for high-income taxpayers in group i relative to low-income taxpayers. Note that, *ex ante*, this is not necessarily to be expected since, in models of optimal auditing, regressive bias within audit groups is generated by the decreasing profile of the audit function. In contrast, in this calculation it is assumed that the flag function is constant within an audit group. Furthermore, the declining profile of the audit function stems directly from an assumption of revenue maximization on part of the tax agency but the declared objective of SKAT is actually *compliance*.

Nevertheless, there are clear indications of regressive bias within audit groups. Despite the fact that each audit group contains only 22–23 observations, in most audit groups the correlation of $\tau - \tau^{\text{eff}}$ and Y is significantly positive using both OLS standard errors and Huber-White heteroscedasticity robust standard errors. For three audit groups the OLS coefficient is negative but in these instances the coefficient is not significantly different from 0.[12] All in all, if the 40 regression coefficient are restricted to be identical, we cannot reject the simple hypothesis that the joint slope is positive at virtually any confidence level.

Conversely, Scotchmer conjectures that effective tax rates may be progressively biased *between* audit groups, in which case this effect may dominate the within group variation such that pooled data displays progressive bias. If this is the case, $\tau - \tau^{\text{eff}}$ is decreasing in third-party reported income since the difference between the statutory tax rate and the effective tax rate will be higher for low audit groups than for high audit groups.

As Figure 4 shows, there is a clear negative correlation in the between-group data,[13] indicating progressive bias. The simple OLS coefficient is significantly negative at virtually any confidence level.

Finally, Figure 5 bears out Scotchmer's conjecture in full. Using pooled data, we estimate a clear negative relationship between the effective tax rate bias and true income. In this case also, the OLS coefficient is significantly negative at virtually any confidence level.

Thus, the data seems to correspond quite well to Scotchmer's conjecture that effective tax rates may be regressively biased within audit groups but progressively biased between groups.

[12] In addition, the OLS slope coefficients for a further 5 and 6 audit groups are not statistically significant at the 5% level using OLS and Huber-White std. errors, respectively.

[13] The observations for the between-groups analysis are calculated as the expectation of $\tau - \tau^{\text{eff}}$ and Y, respectively, for each group.

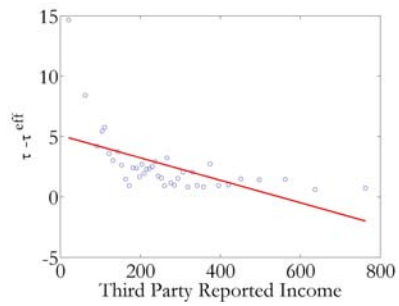

Figure 4. Progressive Bias in the Data: Between Audit Group Variation. The red line is the best linear fit of the data and the coefficient is significantly negative using both standard OLS std. errors and Huber-White std. errors.

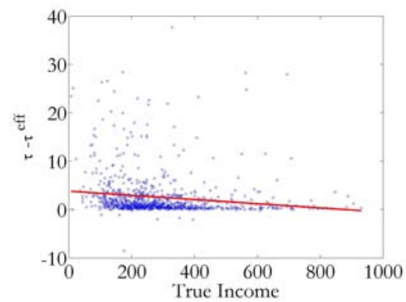

Figure 5. Progressive Bias in the Data: Pooled Variation. The red line is the best linear fit of the data and the coefficient is significantly negative using both standard OLS std. errors and Huber-White std. errors.

4 Theory: A Model of Income Tax Auditing Subject to Information Returns

We build our model on Erard and Feinstein (1994) with a generalization to incorporate third-party income reports. Although current theory is well-equipped to analyze behavior *within* an audit group, we wish to make statements about *aggregate* reporting behavior as well as the tax agency's overall response.

As shown by Kleven et al. (2010), third-party reported income is by far the most powerful predictor of reporting behavior available, making it an ideal candidate for defining audit groups. However, as this variable, like true income, is intuitively best understood as a continuous variable, we allow the tax agency to choose audit functions contingent on the third-party information of a particular taxpayer and interpret each *level* of third-party reported income as an audit group.

Individual Reporting Behavior

Individual taxpayers have true taxable incomes y and file tax returns $\tilde{y}$. Part of true income is reported by third parties. This amount, z, is common knowledge for all parties, i.e. it is a perfect signal. Therefore, we write $y = u + z$, where u is residual income which can be positive or negative as it includes both e.g. wages and deductions not reported by third parties. u is *ex ante* unknown and can only be ascertained by the tax agency by conducting a costly audit. We denote the reported residual x, such that $x = \tilde{y} - z$. The simultaneous distribution of income and third-party reports is common knowledge among taxpayers and the tax agency.

Two facts of the data defy a purely rationalistic modelling of evasion behavior, cf. Table 3: First, a proportion $40.3\% \cdot 18.2\% = 7.3\%$ of taxpayers have non-zero residual income, but do not evade taxes – Second, a proportion 81.8% of taxpayers have zero residual income and do not claim unwarranted deductions. Thus, the overall *compliance rate* is $18.2\% - 7.3\% + 81.8\% = 92.7\%$, but we do not know whether the 81.8% are inherently honest of just effectively honest due to practical circumstances. Erard and Feinstein (1994) splits taxpayers into two broad groups, honest and dishonest taxpayers, and assume that these two types differ only in reporting behavior and not in the scope for evasion, i.e. *their true income distributions are the same*, scaled for the relative proportion of each type. This simple pair of assumptions cannot be reconciled with the two above facts. We prefer to remain agnostic as to whether the 81.8% are inherently honest or merely honest due to practical circumstance and keep these taxpayers in a separate group, which we denote *compliant taxpayers.* As we will argue, the large mass of compliant taxpayers at $u = 0$, makes auditing reports $x = 0$ unattractive, which means that, without loss of generality, we can narrow our focus to equilibria in which dishonest taxpayers for all z choose reports no higher than 0.

Thus, the tax agency's problem becomes disentangling reports by dishonest and honest taxpayers while compliant taxpayers affect the problem only through the scale of the densities of the two former groups. In the subset of taxpayers for whom $u \neq 0$, we denote by Q the fraction of honest taxpayers. Thus, the conditional density of u given z of honest taxpayers is $Q \cdot f_{u|z}(u)$ while for dishonest taxpayers it is $(1 - Q) \cdot f_{u|z}(u)$. We denote by $F_{u|z}$ the conditional distribution function associated with $f_{u|z}$.

We follow Erard and Feinstein (1994) in assuming that taxes are linear in income. Clearly this an abstraction[14] but not an extreme one. As shown by Figure 6, a linear tax function seems to be a reasonable approximation of aggregate tax payments.[15] Figure 6 shows total taxes paid as a function of net income for our sample revealing a roughly linear relationship in which a simple linear regression yields a slope coefficient of approximately 42%.

Finally, while honest taxpayers always report $x = u$, we assume that dishonest taxpayers are risk neutral and maximize expected utility

$$z + p(x|z)\left[(1-t)u - \theta(u-x)\right] + (1 - p(x|z))\left[u - tx\right].$$

[14]It would also be possible to perform the analyses using a full, nonlinear specification of taxes. The main cost would be that of computational intensity as it will then be necessary to adjust the numerical algorithm to account for discontinuities in the differential equation describing the equilibrium. These will arise when the marginal tax rate changes abruptly as an individual tax payer, by adjusting his tax return, slides from one tax bracket to another. We would not expect the conclusions of this paper would to be substantially affected by this change.

[15]An average marginal tax of 42 percent also corresponds quite well to SKATs own estimate of the 2009 marginal tax rate, 43.2 percent, see http://www.skm.dk/tal_statistik/indkomstfordeling/689.html?rel.

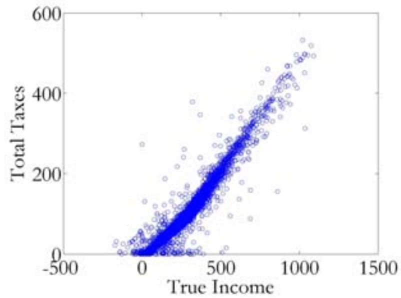

Figure 6. Total Taxes and True Income (in 1,000 DKK) in Denmark, 2006. Income is measured by net income (= personal income + capital income – deductions + stock income + self-employment income + foreign income) and total taxes are those due post-audit.

In optimum, the taxpayer's choice must satisfy the first order condition

$$u = x + \frac{p(x|z) - \frac{t}{\theta+t}}{p'(x|z)}. \tag{1}$$

Given that $p'(x|z)$ is negative and $p(x|z) \leq \frac{t}{\theta+t}$, which will be the case for the optimal $p(\cdot)$, increasing the audit probability will, *ceteris paribus*, lower tax evasion, $u - x$, since the risk of getting caught has risen – the deterrence effect. Ultimately, setting $p(x|z) = \frac{t}{\theta+t}$ eliminates tax evasion entirely. Lowering $p'(\cdot)$ (increasing its absolute value) lowers tax evasion because taxpayers can lower the risk of an audit more for a given decrease in tax evasion. Finally, for fixed p, the model implies that individual tax evasion is increasing in the tax rate, t.

To ensure that a maximum is found, taxpayers must satisfy the second order condition

$$p''(x|z)(x-u) + 2p'(x|z) \leq 0.$$

Optimal Audit Response

The tax agency chooses a continuum of audit schedules $p(x|z)$ for all z. The audit schedule is chosen to maximize the expected revenue (taxes plus fines)

$$\int_{\underline{z}}^{\bar{z}} \left(\int_{\underline{x}}^{\bar{u}} \left(p(x|z) \left[t\mathrm{E}(y|x,z) + \theta \left(\mathrm{E}(y|x,z) - \tilde{y} \right) \right] + (1 - p(x|z)) \, t\tilde{y} \right) \mathrm{d}F_{x|z} \right) \mathrm{d}F_z$$

subject to the budget constraint

$$c \int_{\underline{z}}^{\bar{z}} \left(\int_{\underline{x}}^{\bar{u}} p(x|z) \, \mathrm{d}F_{x|z} \right) \mathrm{d}F_z \leq \int_{\underline{z}}^{\bar{z}} B(z) \, \mathrm{d}F_z \equiv \mathbf{B}$$

where $F_{x|z}$ is the induced conditional distribution function for reported residual income, x, given third-party reported income, z, F_z is the marginal distribution function for z, $B(z)$ is the proportion or density of the overall audit budget, $\mathbf{B}$, allocated to income reports with third-party reported income z. For each x, z, the tax agency must choose p to solve

$$\max_{p} \left(p\left[t\mathrm{E}\left(y|x,z\right)+\theta\left(\mathrm{E}\left(y|x,z\right)-\tilde{y}\right)\right]+(1-p)t\tilde{y}\right)\mathrm{d}F_{x|z}\mathrm{d}F_{z}$$
$$-\lambda c\left[p\ \mathrm{d}F_{x|z}-B\left(z\right)\right]\mathrm{d}F_{z}$$

where λ is the Langrangian multiplier on the budget constraint. This implies a point-wise FOC

$$tE(y|x,z))+\theta E(y|x,z)-\theta\tilde{y}-t\tilde{y})-\lambda c \gtreqless 0 \tag{2}$$

which is greater than, equal to, or less than zero as $p=\frac{t}{t+\theta}$, $p\in\left(0,\frac{t}{t+\theta}\right)$, or $p=0$. The upper bound on the domain of p is the minimal p that induces completely honest reporting at x. We look for equilibria in which the tax agency chooses a mixed strategy, such that (2) holds with equality for all interior solutions x, z.[16]

As mentioned above, our model is a generalization of the model in Erard and Feinstein (1994). Specifically, our model simplifies to theirs if 1) z is identical for all individuals such that $F_{uz}=F_u$, 2) $\log(u)\sim\mathcal{N}\left(\mu,\sigma^2\right)$,[17] and 3) $B(z)\equiv\mathbf{B}$. In this case, the problem becomes that of a partial optimization for a fixed $B(z)$ within an audit group. From Erard and Feinstein we have the following results, restated here for convenience:

Proposition 1 *Fix $z=\hat{z}$. In the unique, within-group equilibrium:*

1. $p(x|z)$, $\frac{\partial p(x|z)}{\partial x}$ *are continuous.*
2. *If $x(u')=x(u'')$ for $u',u''\in[\underline{u},\bar{u}]$, then $u',u''\leq u^{pool}$ and $x(u')=x(u'')=x\left(u^{pool}\right)=\underline{u}$, where $u^{pool}\equiv\max\{u\in[\underline{u},\bar{u}]:x(u)=\underline{u}\}$.*
3. *$x(u)$ is unique and $\frac{\partial x}{\partial u}>0$, $\forall u\in(\underline{u},\bar{u}]$. In addition, $\frac{\partial p}{\partial x}<0$, $\forall x\in[\underline{u},\bar{x}]$ and $p(\bar{x})=0$, where $\bar{x}\equiv x(\bar{u})$.*

Proof. Erard and Feinstein (1994) ■

First, the audit function and its first derivative are continuous – this is useful since we will be using methods of differential equations to identify equilibria. Second, any pooling that occurs must be at the lower bound of the residual income distribution. Third, the best response function of the taxpayers, $x(u)$, is unique and strictly increasing in u on $(\underline{u},\bar{u}]$. This fact follows from (1) and is useful as this implies, by the Implicit Function Theorem, that the inverse function $u(x)$ is unique on $(\underline{u},\bar{x}]$ – in addition, p is decreasing on the domain of reports until it reaches 0 at the report made by the dishonest individual with the highest residual income, $\bar{x}$.

[16] The SOC is $\frac{\partial\mathrm{E}(y|x,z)}{\partial p(x|z)}\geq 0$. In our simulations the solutions we consider will always satisfy this criterion.

[17] I.e. $M=0$ in (3).

Thus, the unique equilibrium of the model is described by the functions $u(x)$ and $p(x|z)$. Once p is determined, the former is implicitly defined as the solution to the taxpayers' FOC and the tax agency chooses p such that (2) holds with equality. The two equations are connected by the tax agency's conditional expectation of taxpayers' true income given the reported income and the third-party information at hand, $\mathrm{E}(y|x,z)$, which is

$$\mathrm{E}(y|x,z) = z + \frac{Qf_{u|z}(x)\,x + (1-Q)\,f_{u|z}(u(x|z))\,\frac{\partial u(x|z)}{\partial x}u(x)}{Qf_{u|z}(x) + (1-Q)\,f_{u|z}(u(x|z))\,\frac{\partial u(x|z)}{\partial x} + \mathbf{1}(x=0)\,M} \quad (3)$$

where $\mathbf{1}(\cdot)$ is the indicator function and the derivative $\frac{\partial u(x|z)}{\partial x}$ is derived from (1) by differentiating implicitly to get $\frac{\partial u}{\partial x} = 2 + \frac{p''(x)(x-u)}{p'(x)}$.[18] The factor M makes an appearance because of the mass of compliant taxpayers at $u=0$. Therefore, $\mathrm{E}(u|x,z)$ is discontinuous in this point which means that the differential equation that describes the equilibrium, c.f. Equation 6 in the Appendix, will have a singularity at $x=0$. However, since the mass of taxpayers in this point consists of only compliant taxpayers (since this mass is infinitely large compared to the density of dishonest taxpayers at $x=0$), it will not be necessary to search for solutions for which $\bar{x} \in (0,\bar{u}]$. Any audit strategy setting $p(0) > 0$ will incur large costs and zero revenue and, due to Proposition 1-1, there can be no equilibrium audit strategy that sets $p(0) = 0$ and $p(x) > 0$ for some $x > 0$.

We are then able to derive a second order differential equation, Equation 6 in the Appendix, which determines the optimal equilibrium responses $p(x|z)$ and $x(u)$ in audit group z using the expressions for $\mathrm{E}(y|x,z)$, $u(x)$, $\frac{\partial u}{\partial x}$ and the tax agency's FOC. However, since some taxpayers pool at the lowest report, to obtain sufficient conditions for equilibrium we need the tax agency's FOC at $x = \underline{u}$ separately as

$$\mathrm{E}(u|x=\underline{u},z) = \frac{Qf_{u|z}(x)\,x + (1-Q)\int_{\underline{u}}^{u^{\text{pool}}} u \cdot f_{u|z}(u)\,\mathrm{d}u}{Qf_{u|z}(x) + (1-Q)\int_{\underline{u}}^{u^{\text{pool}}} f_{u|z}(u)\,\mathrm{d}u} = \frac{\lambda c}{t+\theta} + \underline{u}. \quad (4)$$

As per Proposition 1, the model contains Erard and Feinstein (1994) as a special case when attention is limited to a single audit group in which taxpayers are homogeneous in third-party income reports. To illustrate, Figures 7-10 depict the within-group equilibrium for fixed $B(z)$ at 10%, $\log(u) \sim \mathcal{N}\left(3.42, 0.3^2\right)$ truncated on $[20,44]$, $Q=0.4$, and $t=0.5$.

Figure 7 shows the audit schedule, $p(x|z)$. As stated in Proposition 1-1, it starts in $\underline{u}$, is downward sloping, and terminates in $p(\bar{x}) = 0$. This form balances the need to audit in order to raise revenue with the cost of doing so. The negative slope reflects the need to discourage high-income taxpayers from reporting too low incomes.

[18] In this case, $f_{x|z}(x(u)) = f_{u|z}(u(x))\left|\frac{\partial u(x,z)}{\partial x}\right| = f_{u|z}(u(x))\,\frac{\partial u(x,z)}{\partial x}$ since the SOC implies that $\frac{\partial u}{\partial x} \geq 0$ in interior optimum.

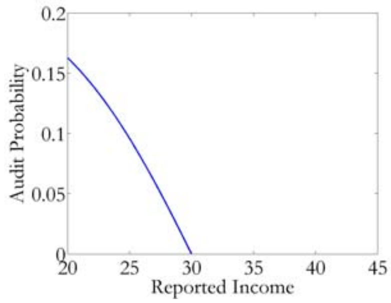

Figure 7. The Optimal Audit Schedule.

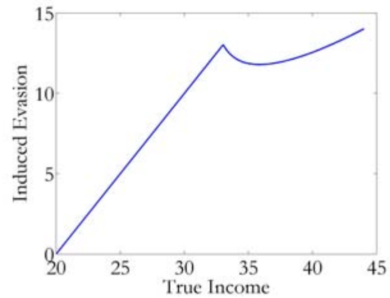

Figure 8. Evasion by True Income for Dishonest Taxpayers.

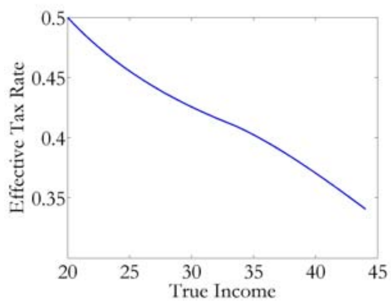

Figure 9. Regressive Bias for Dishonest Taxpayers. In this case, the statutory marginal tax rate is set to $t = 0.5$.

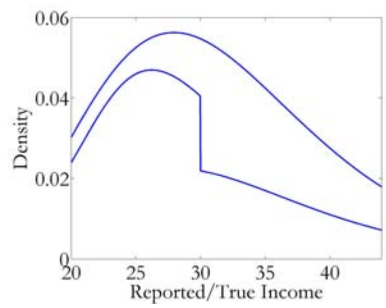

Figure 10. Induced Reporting Behaviour. The lower curve graphs the density of reports by dishonest taxpayers, excluding the mass point at $x = \underline{u}$, while the upper curve graphs the true income distribution.

Figure 8 shows the amount of evasion as a function of income. The linear increase in the first part of the graph reflects pooling of dishonest taxpayers: this comes to be as, for a given audit schedule, there will be some income in $[\underline{u}, \bar{u}]$, u^{pool}, for which the most profitable report is $\underline{u}$ – consequently all taxpayers with residual incomes $u < u^{\text{pool}}$ also report $x = \underline{u}$. Therefore, there will be a point mass in the induced distribution of reports, $f_{x|z}(x)$. After this pooling point, evasion falls rapidly in income until evasion again becomes increasing in income as the probability of detection becomes sufficiently low.[19]

Figure 9 shows the effect of the optimal audit schedule on the effective tax rate. The effective tax rate is calculated as

$$\tau^{\text{eff}} = \frac{p(x) \cdot (ty + \theta(y - \tilde{y})) + (1 - p(x)) \cdot t\tilde{y}}{y}. \quad (5)$$

[19]The extent to which evasion, $u - x$, drops at u^{pool} depends on the shape of the audit schedule and, hence, ultimately on the distribution of u given z.

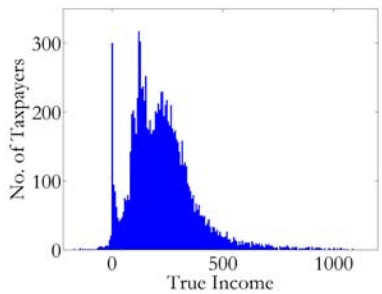

Figure 11. The Distribution of True Income, y, found in our data sample. Income is measured by net income (= personal income + capital income – deductions + stock income + self-employment income + foreign income).

Since the audit schedule is decreasing, the effective tax rate must also be decreasing. Therefore, high-income taxpayers pay significantly less than indicated by the statutory tax rate, which, in the case of Figure 9, is $t = 0.5$.

Figure 10 shows the induced distribution of incomes and reports. The top graph is the original income distribution, which in this case is lognormal. The lower graph shows the distribution of induced reports, i.e. the equilibrium response of all taxpayers to the audit schedule. The right part of the graph is just a scaling of the original income distribution by Q while the left part is a weighted average of reports by honest and dishonest taxpayers. The whole graph is somewhat lower than the original income distribution as there is a mass point of mostly dishonest taxpayers at $\underline{u}$ which is omitted from the figure.

Parametrization

Income Distributions

We use the taxpayer data from Kleven et al. (2010) to construct the income distribution needed in the model. As income measure we choose net income defined as the sum of personal income, capital income, stock income, self-employment income and foreign income less deductions. Figure 11 shows the empirical distribution of true net income in our sample. The distribution looks roughly lognormal.

To fit the simultaneous distribution of z and u, we fit a mixed lognormal distribution.[20] Approximately 82% of all observations have $u = 0$, i.e. these individuals have their entire income reported by third parties. This means that the conditional distributions of $u|z$ consist of a continuous density and a mass point in $u = 0$. The parametrically fitted distributions are estimated

[20] Ideally, one would want to use a non-parametric kernel distribution estimate. Although this is possible it does create some numerical difficulties in that such estimates usually generate troughs of zero density in the tails of the conditional distributions which will cause the numerical algorithm to fail.

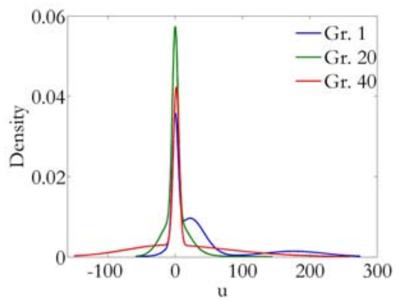

Figure 12. Conditional Densities of u Given z. Show for the lowest, the middle and the top audit group using 40 audit groups. The densities are estimated excluding the mass point at $u = 0$. The residual income, u, is measured in 1,000 DKK.

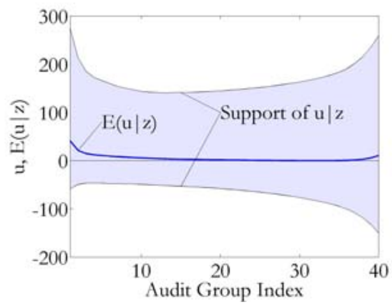

Figure 13. The Support of u Across Audit Groups. The conditional densities of $u|z$ are truncated at the 0.5 and 99.5 percent fractiles of the unrestricted conditional distributions. The residual income, u, is measured in 1,000 DKK.

excluding the observations in the mass point. Relative to the distribution of true net income, y, in Figure 11, the conditional distributions (excl. the mass point) are not lognormal and they display much more symmetry, in general. Also, the conditional distributions have very little variance as compared to the distribution of net income, attributable to the fact that third-party reports convey substantial information as to the value of the taxpayers' true income. Figure 12 depicts three conditional distributions of u given z in the lowest, middle and top part of the domain of z.

For numerical reasons, we need to truncate the domain of the conditional distributions where the densities are negligible. This is necessary to prevent the ratio $\frac{f_{u|z}(x)}{f_{u|z}(u(x))}$ in Equation 6 in the Appendix from diverging to zero or infinity when $f_{u|z}(x)$ or $f_{u|z}(u(x))$ approach 0, respectively. We do this by truncating the unrestricted conditional densities at the 0.5% and 99.5% fractiles. This means that the supports of the resulting distributions will vary in z. This is illustrated in Figure 13 along with the conditional mean of u given z. We do not view the truncation as a serious limitation of the model as we truncate the conditional distribution where there is, essentially, no density.

Generally, the conditional distributions look very similar across audit groups, but there are subtle differences. The conditional mean of $u|z$ is U-shaped; decreasing in z from approximately $41,000$ DKK to just under 0 DKK with a slight increase to approximately $11,000$ DKK at the end. The wider support of $u|z$ in the extremities of z in Figure 13 reflects that taxpayers with very low or very high levels of third-party reported income generally have more complex income portfolios explaining the greater amount of income for which SKAT does not have third-party information. The middle group are to a much higher degree "typical" taxpayers with less residual-income variance.

Honesty

A key model parameter is the fraction of honest taxpayers, Q. In order to determine an appropriate value of this parameter, we must account for the fact that, in reality, some taxpayers seem to make reporting mistakes. For example, in the data some reports are adjusted downward by the auditor which means that, in the absense of an audit, the taxpayer would have payed more than intended by the statutory tax system.

We approach the problem in the following manner. First, we assume that no taxpayer will try to evade taxes on income that is reported by a third party (this assumption is bourne out in the data as shown in Table 3). Then we separate the taxpayers into two groups, one containing those whose true income is entirely reported by third parties so that $y = z + u$ and $u = 0$, and the other containing those with some residual income not subject to third-party reporting, $u \neq 0$. The second group is then separated according to whether or not the audit led to a change in their reported income, i.e. whether or not $x \neq u$. In other words, we are classifying taxpayers into groups of compliant ($u = x = 0$), inherently honest ($u \neq 0$, $x = u$) and dishonest taxpayers ($u \neq 0$, $x \neq u$).

Table 3. Calibration of Q, the Fraction of Honest Taxpayers

	Entire income reported by third-parties	Some income not reported by third-parties
Total reports	8783	1957
# underreported	0	904
# correct	8773	758
# overreported	10	295
Fraction of correct reports	.999	.403
	(.0002)	(.0219)
# not underreporting	8773	1053
Fraction not underreporting	.999	.541
	(.0002)	(.0221)
# honest taxpayers*	–	1348
Fraction honest taxpayers*	–	.689
	–	(.0205)

*Assuming that unintentional underreporting is as frequent as unintentional overreporting. I.e. # honest taxpayers = 295 + 758 + 295 = 1348. The extra honest taxpayers relative the category of "not underreporting" are distributed into sample strata proportionally to the sample strata sizes.
Notes. There are $N = 10,740$ taxpayers in the sample. Standard errors in parentheses. All shares and standard errors are calculated subject to the stratification scheme. This explains the discrepancies between the no. of taxpayers and the share of taxpayers estimated in a given category since the no. of taxpayers is simply counted in the sample.

Table 3 shows this decomposition. First, note that among taxpayers whose entire income is reported by third parties, only 0.1% misreport their income, and any misreports are in fact taxpayers overreporting their income. Among the 18.2% of the taxpayers that have some of their income not reported by third parties, 40.3% report correctly. We can define honest taxpayers in several ways. The simplest is to include only those reporting correctly. This definition fails to acknowledge the fact that some taxpayers make reporting errors. However, modelling reporting errors is beyond the scope of this paper. A revenue maximizing tax agency cares not whether revenue is collected from dishonest taxpayers who intentionally underreport or honest taxpayers who do so by mistake. We classify overreporting taxpayers as honest and underreporting taxpayers as dishonest. Thus, the number of honest taxpayers is the sum of those reporting correctly and those overreporting by mistake, which corresponds to $Q = 54.1\%$. The residual, 45.9%, consists of both dishonest taxpayers and taxpayers underreporting by mistake, whom we cannot distinguish.[21]

Penalty

Our model, building on Erard and Feinstein (1994), incorporates a penalty for tax evasion that is proportional to the amount of income evaded. As in Erard and Feinstein (1994), we set the proportionality factor θ equal to 1.2 meaning that tax evasion is fined with 1.2 times the evaded income on top of what is owed in taxes. Fines in the Danish context are calculated on the basis of *taxes* evaded. In the case of deliberate tax evasion, the fine is calculated as one times the evaded taxes under 30,000 DKK and two times the evaded taxes exceeding 30,000 DKK. As a back-of-the-envelope calculation, assume setting the proportional tax rate at 42%. Evading taxes amounting to 100,000 DKK, equivalent to concealing 238,000 DKK, will result in a fine of 170,000 DKK, or approximately a share of 0.71 of the income evaded.[22] Compared to this, a θ of 1.2 is high. So what is a reasonable level? On the one hand, there is also the possibility of imprisonment of up to 18 months (and in very severe cases even eight years) and additionally the humiliation of being found guilty of tax evasion. On the other hand, fines for tax evasion are rare and imprisonment even more so. The level of the penalty has a bearing on the amount of tax evasion of dishonest taxpayers in the model, but we believe that the impact of the chosen level on the implications of adding explicit audit group to the theory are of the second order. In addition, as we assume that taxpayers are risk neutral, θ will effectively also capture the risk aversion presumably present among the taxpayers, arguing for a higher value than the monetary penalty alone would justify.[23]

As a normalization, we set the cost of an audit, c, to 1. Thus, $\mathbf{B}$ can be interpreted as the percentage of the population of taxpayers (or, for $B(z)$, of an audit group) excluding compliant taxpayers that is selected for audit. We set $\mathbf{B}$ equal to 4% which will result in an audit coverage of approximately 0.7% of

[21] Alternatively, if one assumes that unintentional underreporting is as frequent as unintentional overreporting – symmetry of mistakes – the fraction of honest taxpayers is 68.9%. Although varying the size of Q has quantitative effects on the output of the calibrated model, this does not qualitively change the correlation of true income and effective tax rates.

[22] For tax evasion dubbed "severe negligence" the fines are set as one half times evaded taxes up to 30,000 DKK and one times the tax evaded exceeding that.

[23] Although the argument for any particular value of θ is tenuous, the conclusions regarding bias of effective tax rates will hold for any positive θ. Since we are not engaged in positive statements about the tax system, the value of θ is not very important.

the entire population of taxpayers.[24]

5 Results

Simulation Strategy

An individual solution to Equation 6 in the Appendix, $\left(p, \frac{\partial p}{\partial x}\right)$, that corresponds to a particular z is found numerically using methods of Ordinary Differential Equations (ODE) as initial value problems.[25] The solver is initialized using $p(\bar{x}) = 0$ and $p'(\bar{x}) = \left(\frac{t}{t+\theta}\right) / (\bar{u} - \bar{x})$, where $\bar{x} \equiv x(\bar{u})$. Thus, starting at the end-point of the equilibrium-path audit probabilities, a numerical solver finds values in steps until $\underline{u}$ is reached. This ensures that the taxpayers' as well as the tax agency's first and second order conditions are met for reports $x \in (\underline{u}, \bar{x}]$. However, since a positive mass of taxpayers are pooling their reports at $x = \underline{u}$, the expectation $\mathrm{E}(u|x, z)$ is not differentiable in this point. Therefore, we check that the tax agency's FOC is met in the pooling point separately after finding some candidate solution, cf. Equation 4.

The difficulty in identifying equilibria in this model stems from *a priori* indetermination of λ and $\bar{x}$: we must satisfy $\mathrm{E}(u|x = \underline{u}, z) - \underline{u} = \frac{\lambda c}{t+\theta}$ which depends on both variables, λ directly but also via u^{pool} which also depends on $\bar{x}$ through $p(\underline{u})$ and $p'(\underline{u})$. Our solution method searches the space of possible $(\lambda, \bar{x})$ for candidate solutions, for each checking whether the tax agency's optimization constraints are satisfied on the entire domain of x, until satisfactory solutions are found.

While mathematically and intuitively third-party reported income, z, is naturally understood to be a continuous variable described by the simultaneous distribution of u and z, we determine the optimal allocation of the total audit budget on the domain of z by constructing a representative grid of values on the domain of z and maximizing total revenue as a function of $B(z)$. Of course, given the rather large number of representative audit groups, care must be taken to ensure that the optimum is global by using a global maximization algorithm.

Simulation Results

Audit Intensity

The optimal profile is depicted in Figure 14 together with the value of $B(z)$ that ensures that $p(\bar{x})$ becomes 0 at exactly $\bar{x} = 0$. Since the distribution of reported residual-income, x, conditional on the third-party reports, z, contains

[24] Kleven et al. (2010) report an audit rate for the entire Danish population of 4.2%. But this covers audits at varying breadths and depths and should more be thought of as an upper bound. To compare, the US audit rate for individual taxpayers was between 0.5 and 1.0% in the fiscal years 2000-2009 according to Internal Revenue Service (2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009).

[25] We employ an ODE solver developed in Shampine (2009).

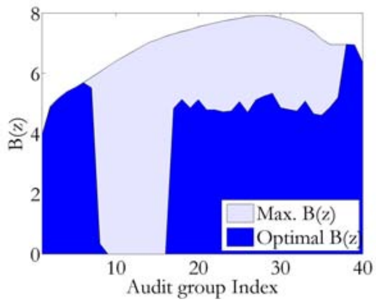

Figure 14. The Optimal Budget Allocation Across Audit Groups. The budget is allocated such that 4 percent of all taxpayers, excl. those with u = 0, are audited. As 81.8 percent of the taxpayers in the sample have u = 0, the total budget amounts to auditing approximately 0.7 percent of all taxpayers. The percentages shown denote the share of taxpayers within an audit group (excl. u = 0) selected for audits.

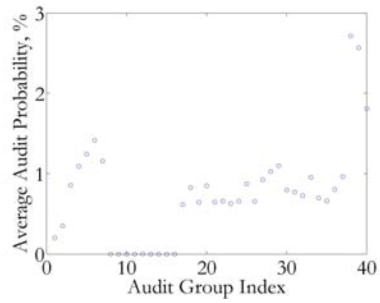

Figure 15. The Average Audit Probability Across Groups. The audit intensity, as measured by the probability of evasion being detected, matches the pattern of the budget allocation across groups.

a substantial mass point of taxpayers that are "honest due to circumstances" in $x = 0$, it can never be optimal for the tax agency to apply a positive audit probability here. Thus, the mass point provides a natural upper limit to the values of $B(z)$ in the optimal budget allocation. The optimal budget allocation shown in Figure 14 allocates maximum budget to the very low and high audit groups, and zero or intermediate values to the middle groups.

Within audit groups, the audit schedules are qualitatively the same as that found in the standard Erard and Feinstein (1994) model, cf. Figure 7, Section 4. By averaging the audit schedules within audit groups, we get the between-groups optimal audit schedule shown in Figure 15. This schedule, however, does not look anything like the within-audit-group schedules reflecting instead the allocation of audit resources across audit groups.

Evasion

Figure 16 shows the relationship between average evaded income and true income for dishonest taxpayers. Within audit groups, the relationship found is identical to that in Erard and Feinstein (1994), cf. Figure 8 in Section 4, except that the relationship is strictly monotonic and positive for audit groups not subject to audit pressure since all dishonest taxpayers in those groups pool at the lowest possible report, $x = \underline{u}$. In the pooled output, we find an overall positive relationship between true income and evaded income. However, there is a small interval in which evasion is decreasing. This reflects the difference in opportunities to evade taxes of the different audit groups proxied e.g. by the

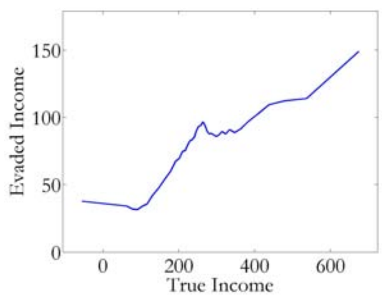

Figure 16. Evaded Income by Dishonest Taxpayers. Shown as a function of true income (both in 1,000 DKK). Since within-audit-groups' evasion curves overlap on the net income scale, we average across audit groups to produce this average evasion diagram.

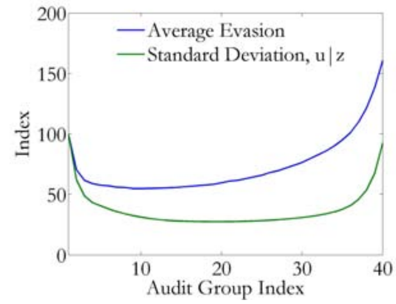

Figure 17. Average Evasion and $\sigma_{u|z}$ Between Audit Groups. Both are indexed to 100 at audit group 1. The correlation coefficient is 0.72.

standard deviation of $u|z$. Intuitively, the higher the standard deviation of true residual-incomes, the more difficult it becomes for the tax agency to distinguish whether a relatively low level of reported residual-income is a sign of evasion or not, and the less risky evasion becomes. As illustrated in Figure 17, there is a very close correspondence between the standard deviations and the average evasion across audit groups. Both are U-shaped and high for the very low and the very high audit groups and the correlation coefficient is 0.72.

This result has implications for empirical studies of tax evasion conditioning on post-audit income. First, it is important to condition on the audit groups actually employed by the tax agency (or a proxy thereof as e.g. third-party reported information), especially when using micro-level data, such as in Clotfelter (1983). Feinstein (1991) does not condition on audit groups and finds no hard evidence of a relationship between income and evasion. Second, income is not necessarily the variable of interest in such empirical studies: What we should expect to find is a causal relationship from the opportunity to evade taxes to the level of evasion. The implicit hypothesis of studies trying to find some relationship between income and evasion is that, apart from the relationship between risk aversion and income, income and the extent of evasion possibilities are positively correlated. As evidenced by Figure 16 this is not necessarily the case, at least not on the entire domain of y. In a setting where the tax agency holds third-party information on taxpayers' income, we should expect to find that a variable expressing evasion possibilities, such as the standard deviation of $u|z$ would explain much of the variation in evasion leaving little explanation power to the level of net income. In Klepper and Nagin (1989) this idea is also pursued, showing a correlation on line item returns between noncompliance and IRS' ability to establish and punish noncompliance using summarized TCMP data. Among other things, they condition on third-party reporting (a dummy

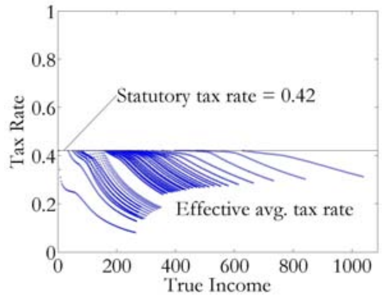

Figure 18. The (ex ante) Effective Tax Rate within Audit Groups. Shown as a function of true income. Output is diplayed for positive incomes and tax rates.

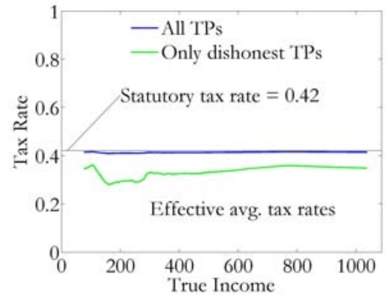

Figure 19. Aggregate, Pooled Variation in the Effective Tax Rate. Shown as a function of true income. Calculations are performed on positive incomes and tax rates.

variable indicating whether a line item is subject to information reporting or not), proxies for the complexity of the line item (and hence the deniability in case misreporting is found) and, importantly, the variance of true income within each line item. The most significant result of their analysis is, similarly, the positive correlation between the variability of true income and noncompliance. That their results are consistent with ours follows from Figure 17 and the fact that the correlation coefficient between evasion and $\sigma_{u|z}$ between audit groups is 0.72 in the model-generated output, while in the data it is 0.76.

Effective Tax Rate Bias

We calculate the bias of effective average tax rates in a manner similar to that in Section 3, except, as noted, we follow the literature and include expected penalties in the calculation, c.f. Equation 5. Recall, in Section 3 we illustrated the bias in effective average tax rates using $\tau - \tau^{\text{eff}}$ since post-audit average tax rates vary across individuals. As the marginal tax rate is constant in the model, $\tau = t$, we illustrate the effective average tax rate bias in the same manner as in the stylized Figures 1 and 3 in Section 2 rather than Figures 4 and 5.

Figures 18 and 19 show the model-generated effective average tax rates within groups and averaging pooled data, respectively. Figure 19 displays a positive relationship when pooling taxpayers, which reflects the dominant effect of progressive bias between audit groups, in line with Scotchmer's conjecture and the relationships found in the raw data. Although effective average tax rates within groups are regressively biased, this relationship is reversed in the pooled output by the progressive bias between audit groups.

In fact, this qualitative result is surprisingly robust within the context of the model. For all parametrizations we have explored, effective average tax rates have exhibited the same properties as in the figures above.

6 Concluding Remarks

This paper is all about information. First, we clearly show the necessity, from the researcher's point of view, of obtaining sufficiently detailed data for studies of tax evasion and deterrence – without it, the within- and between-variation of tax evasion cannot be separated and conclusions about individual behavior become tenuous at best. Second, we highlight the theoretical importance of considering third-party reported information in analyses of tax evasion and deterrence – without it (or another proxy for audit groups) we cannot generate the subtle covariance structure of income and tax evasion and its effect on effective tax rates.

Specifically, we use data and model-generated output to confirm Scotchmer's (1987) conjecture. First, using raw data and minimal assumptions we find strong evidence for a regressive bias within audit groups and progressive bias between audit groups, resulting in an overall progressive bias in pooled data. Second, we find that these empirical findings are consistent with state-of-the-art theory of optimal auditing and evasion, once we allow for the tax agency's use of third-party reported information and make the population of taxpayers heterogeneous in true income, income reported by third parties, and honesty.

The prediction that effective tax rates exhibit this qualitative pattern seems robust: it is generated by our realistically complex model, calibrated to data, as well as in Scotchmer's simple model in which the assumptions are importantly different. Furthermore, we find the same pattern in data collected in a setting in which the tax agency's stated objective is somewhat different from that used in current models. We are confident that similar empirical relationships would be found in data from any tax auditor that, as SKAT does, employs a strong signal in predicting true incomes.

A natural objection to the model we employ is the lack of general equilibrium considerations, for example feedback into labor market choices. As experience shows (e.g. Pencavel, 1979), adding such features to the model complicates the analysis substantially, which in our setting may be prohibitive. In any case, this issue is beyond the scope of this paper. However, we do believe that the value-added of pursuing this line of research is small. Rather, we believe that efforts should be directed towards reaching a consensus on a reduced-form model of optimal auditing and evasion, which, in turn, can be applied in other theoretical settings.

Our model seems to indicate a somewhat higher level of effective tax rate bias than seen in the raw data. Since we are mainly interested in the sign of the correlations of income and effective tax rates, we do not judge this as a failure of the model *per se* – rather, this is an indication that quantitative conclusions should be drawn on the basis of a more elaborate model. In particular, such a model should allow for more uncertainty in taxpayers' optimization with respect to audit risk, taxes owed (i.e. allowing for reporting errors), and probability of being fined a particular amount conditional on income underreported or taxes evaded. Further, such a model should be sure to incorporate precisely the objective function of the particular tax agency in question – e.g. whether the goal is revenue maximization or compliance – as well as allow for risk aversion on part of dishonest taxpayers. Note that all of these extensions on the taxpayer side will tend to make them less inclined to evade. Thus, we can reasonably expect that any such extensions will make the model correspond even closer to the observed data.

The broad implications for the empirical literature on tax evasion is the need to control explicitely for the difference in between- and within-variation of income and tax evasion. For example, it is impossible to falsify the theoretically predicted increasing relationship between income and tax evasion within audit groups using pooled data alone – this may explain the hitherto mixed results in the empirical literature.

A logical next step for the literature is to estimate a structural model including third-party information to qualify existing model tests, which so far have relied exclusively on survey and experimental data. The model in this paper is well suited as a point of departure for such an excercise.

7 Appendix

A Simple Robustness Check

To check in the most basic way whether our results are particular to the Danish incomes distribution, we will consider a version of the model presented in Section 4 in which there are ten possible levels of third-party reported income and true incomes are uniformly distributed within audit groups. In this case, we will let true income be distributed on the interval $[20, 220]$.

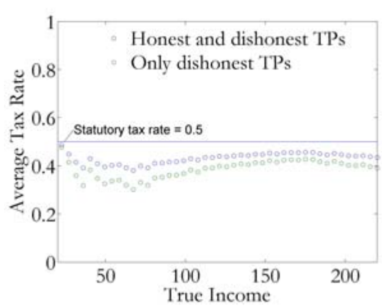

Figure 20. Scenario 1: Effective Tax Rate with 10 Audit Groups.

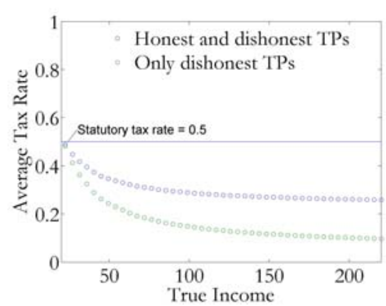

Figure 21. Scenario 2: Effective Average Tax Rate without Audit Groups.

As Figure 20 illustrates, progressive bias of effective tax rates in pooled data is also a feature of this simply specified model using a counterfactual income distribution. To compare, Figure 21 illustrates the case in which the TA ignores third-party information.

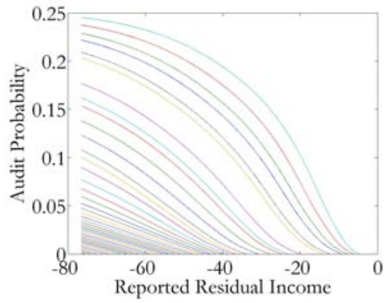

Figure 22. Examples of Optimal Audit Functions, $p(x|z)$. Audit function are shown for a particular audit group and for varying budget allocations. Increasing the budget results in higher optimal audit functions and a higher $\bar{x}$.

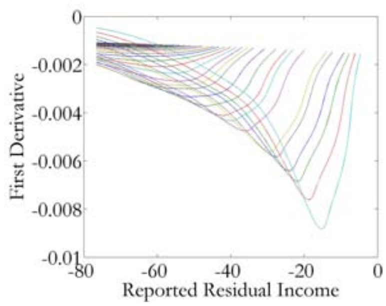

Figure 23. Examples of Optimal Audit Functions: $p'(x|z)$.

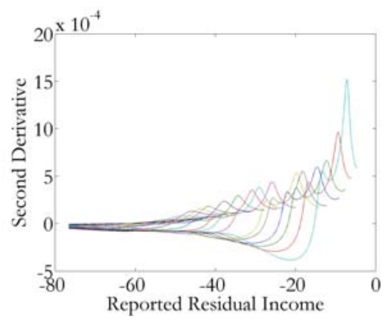

Figure 24. Examples of Optimal Audit Functions: $p''(x|z)$.

A Set of Optimal Audit Schedules Within an Audit Group

The second order differential equation is obtained by combining (1), (2), (3) and the expression for $\frac{\partial u}{\partial x}$ to get

$$p''(x) = \left(\frac{Qf_u(x)\left(\frac{\lambda c}{\theta+t}\right)}{(1-Q)\, f_u(u(x)) \left[\frac{p(x)-\frac{t}{\theta+t}}{p'(x)} - \frac{\lambda c}{\theta+t}\right]} - 2 \right) \times p'(x)^2 \left(\frac{t}{t+\theta} - p(x)\right)^{-1} \tag{6}$$

suppressing z for convenience. Here we are concerned with equilibrium conditions, so we can safely ignore the point mass, M.

To illustrate, we provide plots of a set of optimal audit functions, within an audit group, which vary on $\bar{x}$ and therefore B:

Bibliography

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

Andreoni, J., B. Erard, and J. Feinstein (1998). Tax compliance. *Journal of Economic Literature 36*(2), 818–60.

Clotfelter, C. (1983). Tax evasion and tax rates: An analysis of individual returns. *The Review of Economics and Statistics 65*(3), 363–373.

Engel, E. M. R. A. and J. Hines, James R (1999). Understanding tax evasion dynamics. *NBER Working Papers 6903*.

Erard, B. and J. S. Feinstein (1994). Honesty and evasion in the tax compliance game. *RAND Journal of Economics 25*(1), 1–19.

Feinstein, J. (1991). An econometric analysis of income tax evasion and its detection. *The RAND Journal of Economics 22*(1), 14–35.

Internal Revenue Service (2000). *Data Book 2000, Publication 55B*. Washington, DC.

Internal Revenue Service (2001). *Data Book 2001, Publication 55B*. Washington, DC.

Internal Revenue Service (2002). *Data Book 2002, Publication 55B*. Washington, DC.

Internal Revenue Service (2003). *Data Book 2003, Publication 55B*. Washington, DC.

Internal Revenue Service (2004). *Data Book 2004, Publication 55B*. Washington, DC.

Internal Revenue Service (2005). *Data Book 2005, Publication 55B*. Washington, DC.

Internal Revenue Service (2006). *Data Book 2006, Publication 55B*. Washington, DC.

Internal Revenue Service (2007). *Data Book 2007, Publication 55B*. Washington, DC.

Internal Revenue Service (2008). *Data Book 2008, Publication 55B*. Washington, DC.

Internal Revenue Service (2009). *Data Book 2009, Publication 55B*. Washington, DC.

Klepper, S. and D. Nagin (1989). The anatomy of tax evasion. *Journal of Law, Economics, and Organization 5*(1), 1–24.

Kleven, H. J., M. B. Knudsen, C. T. Kreiner, S. Pedersen, and E. Saez (2010). Unwilling or unable to cheat? Evidence from a randomized tax audit experiment in Denmark. *NBER Working Paper 15769*.

Macho-Stadler, I. and J. D. Perez-Castrillo (2002). Auditing with signals. *Economica 69*(273), 1–20.

Myerson, R. (1981). Optimal auction design. *Mathematics of operations research 6*(1), 58–73.

Nielsen, S. B. and P. B. Sørensen (1997). On the optimality of the Nordic system of dual income taxation. *Journal of Public Economics 63*, 311–329.

Pencavel, J. H. (1979). A note on income tax evasion, labor supply, and nonlinear tax schedules. *Journal of Public Economics 12*(1), 115–124.

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

Reinganum, J. F. and L. L. Wilde (1986a). Equilibrium verification and reporting policies in a model of tax compliance. *International Economic Review 27*(3), 739–60.

Reinganum, J. F. and L. L. Wilde (1986b). Settlement, litigation, and the allocation of litigation costs. *The RAND Journal of Economics 17*(4), 557–566.

Sanchez, I. and J. Sobel (1993). Hierarchical design and enforcement of income tax policies. *Journal of Public Economics 50*(3), 345–69.

Scotchmer, S. (1987). Audit classes and tax enforcement policy. *American Economic Review 77*(2), 229–33.

Shampine, L. F. (2009). Vectorized solution of ODEs in Matlab with control of residual and error. `http://faculty.smu.edu/shampine/`.

SKAT (2009). Borgernes efterlevelse af skattereglerne - indkomståret 2006. Report. `www.skat.dk`.

Slemrod, J. (2007). Cheating ourselves: The economics of tax evasion. *Journal of Economic Perspectives 21*(1), 25–48.

Yitzhaki, S. (1974). A note on 'income tax evasion: A theoretical analysis'. *Journal of Public Economics 3*(2), 201–202.