

Hierarchical design and enforcement of income tax policies

Isabel Sánchez

Universidad Carlos III de Madrid, Calle Madrid, 126, 28903 GETAFE, Madrid, Spain

Joel Sobel*

Department of Economics, University of California, San Diego, La Jolla, CA 92093-0508, USA

Received October 1991, revised version received January 1992

A hierarchical model of tax compliance is studied in which the government selects a tax policy and then delegates the responsibility to collect taxes to the IRS. There is a fixed distribution of income in the economy. Individuals differ by their income, which is known neither to the government nor to the IRS. We completely characterize the solution to the IRS's revenue-maximizing problem. When taxpayers are risk neutral, the optimal auditing policy divides the reported incomes into at most three groups. We show that the government provides a smaller budget to the IRS than the IRS would wish.

1. Introduction

A characteristic of bureaucracies is a division of functions between different agencies. In hierarchical structures higher levels of the organization design policies that are implemented by lower entities. If the agency that enforces policy has access to more information or has different goals than the level of the organization that creates the policy, then conflicts of interest arise. As a result, the behavior of the hierarchical organization may be quite different from the behavior of an organization in which a single entity both makes and enforces policy decisions.

In this paper we analyze a hierarchical model of tax compliance designed to describe the conflict between the government, which has the responsibility of selecting a tax policy, and the auditor, which has the job of enforcing the policy. We will refer to the auditor as the IRS, after the Internal Revenue

Correspondence to: I. Sánchez, Universidad Carlos III de Madrid, Calle Madrid, 126, 28903 GETAFE, Madrid, Spain.

*This paper contains material that originally appeared as Chapters 2 and 3, 'Principal-agent models of income tax compliance' and 'Hierarchical design and enforcement of income tax policies' of Sánchez (1987). We thank Jim Andreoni, Suzanne Scotchmer, and referees for helpful comments. Sobel thanks the NSF and the Sloan Foundation for financial support.

Service, which collects taxes in the United States. We choose to divide the roles of designing and enforcing tax policy to add realism to existing models of the tax enforcement procedure. The government delegates the responsibility to collect taxes because it is unable to carry out the enforcement activity on its own. The division of responsibility arises if the task of enforcing tax policies requires information and expertise that the government does not have. We assume that the government cannot fully observe the actions of the IRS, so there is a potential conflict between the incentives of the IRS and the government. We investigate the consequences of this conflict.

In the model of this paper there is a fixed distribution of income in the economy. We ignore the distortions on labor supply produced by taxation. Observable attributes of the taxpayers sort them into audit classes. Within each class, individuals differ by their income, which is known to neither the government nor the IRS. The IRS is able to conduct audits to learn a taxpayer's income. Audits are costly, so the IRS asks taxpayers to report their income voluntarily and then audits only a fraction of the reports. Given any reported income, the IRS selects a probability of conducting an audit; if it does not conduct an audit, then the taxpayer pays an amount dictated by the tax structure designed by the government. If the IRS does conduct an audit, then the taxpayer pays an amount that depends on the reported income and the amount (equal to the true income) found by the audit. While the IRS is able to formulate the auditing policy, the government determines the actual tax structure and the budget that the IRS uses to conduct audits.

We have two main findings. First, we completely characterize the solution to the IRS's revenue-maximizing problem. Our model is based on Scotchmer (1986). Our results extend her treatment of the auditor's problem. When taxpayers are risk neutral and the distribution of income in the economy is fixed and known to the government, the optimal auditing policy divides the reported incomes into at most three groups. The lowest group is audited with sufficiently high probability that only the individuals with the lowest income in their class make these reports; no individual dishonestly reports an income this low. Reports in the middle group are audited with positive probability but not frequently enough to deter cheating. Reports in the highest group are not audited at all.

Second, we show that the government generally provides a smaller budget to the IRS than the IRS would wish. At the optimal auditing budget, an extra dollar allocated to auditing generates more than one dollar of revenue.

We also describe distortions that arise because the government cannot observe income without cost. We give conditions under which an increase in the cost of auditing leads to a decrease in spending for public goods. The result, while intuitive, is not true in general. High income individuals have more opportunities to evade taxes. Consequently, as noted by Scotchmer (1986, 1991), effective tax rates are more regressive than the nominal rates.

When the cost of auditing goes up, the government may increase public good spending in order to redistribute income.

Recent contributions to the literature on optimal taxation have investigated the possibility of tax evasion on the design of optimal income tax policies under the assumption that income is not completely observable. Allingham and Sandmo (1972) and Srinivasan (1973) present the first models of tax evasion. These papers and the work of Kolm (1973), Yitzhaki (1974) and Christiansen (1980), among others, assume that the probability of discovering a tax evasion does not depend on information obtained from tax returns. Instead, it is assumed that a fixed proportion of tax cheaters are discovered.

More recent work examines the strategic interaction between taxpayers and the tax collection agency. The control of tax evasion differs from that of other delinquent activities in that the enforcement agency can use information provided voluntarily by the taxpayer. Enforcement and penalty policies can be based on reported income and any other information correlated with unknown income which in turn is influenced by the auditing policy.

Reinganum and Wilde (1985), Scotchmer (1986), Border and Sobel (1987) and Mookherjee and Png (1989) treat models in which the IRS is able to commit to an auditing schedule that depends on reported income. Although these papers differ in their level of generality and their assumptions on preferences, all assume that the auditor has direct control over the tax structure. That is, they implicitly assume either that there is a single agency in charge of the design and implementation of the tax policy or at least that there is no conflict of interest among the agencies that decide on the different instruments.

Crémer, Marchand and Pestieau (1990) use a model that is identical to ours. They do not characterize the solution to the IRS's problem, but they limit attention to cutoff strategies; the results in section 3 of this paper give conditions when there is no loss of generality in doing so. They go beyond this paper and discuss what happens to the optimal auditing policy when the government can divide workers into income classes.

Some other papers have studied incentives in a hierarchical setting. Demski and Sappington (1987) present a model of hierarchical enforcement in which Congress must design incentives that encourage a self-interested regulator to acquire the information needed to induce a monopolist to control costs. They contrast the predictions of their model with one where the regulator acts to maximize social welfare. Tirole (1986) discusses incentives in a hierarchical organization where the actor in charge of the organization cannot control contracts between actors at lower levels. He discusses the impact that cooperative behavior within levels of the hierarchy has on optimal incentive schemes. In our framework, it may be in the

interest of a taxpayer to cooperate with the auditor in order to reduce taxes. Such behavior would arise if we allowed individuals to bribe the IRS. However, we assume that bribes are not feasible.

The next section describes the basic model. Section 3 characterizes the solution to the revenue-maximizing problem of the auditor. Section 4 uses the results of section 3 to describe the qualitative properties of the government's optimal tax and auditing policies. Section 5 compares the solution to the three-layer incentive problem of section 5 with the policies that are optimal when the government need not delegate auditing authority to the IRS. Section 6 discusses some of the limitations of our model.

2. The model

We consider a population of individuals characterized by their income, i , which is private information. The distribution of income is supported on the interval $[l, h]$; $F(i)$ represents the fraction of the population with income less than or equal to i ; we assume that $F(\cdot)$ is continuously differentiable and that $f(i) = F'(i) > 0$ on (l, h) . At times we also assume that the hazard rate associated with $F(\cdot)$, $\gamma(i) = [1 - F(i)]/f(i)$, is strictly decreasing.

The government must raise money in order to provide public goods. It does so by designing a tax policy. Since we assume that income is private information, the tax policy must include some procedures for determining an individual's income. The government cannot do this directly but uses an auditor (the IRS) with access to a technology that can discover an individual's true income. It costs $c > 0$ to conduct an audit. The tax policy of the government consists of a tax schedule $t(\cdot)$ and a budget B that the IRS uses to conduct audits. The IRS treats both the tax schedule and its budget as parameters in its optimization problem. In section 3, where we analyze the IRS's problem, we assume that $t(\cdot)$ is strictly increasing and differentiable.¹ In section 4, where the government chooses the tax structure, we assume that $t(\cdot)$ is an increasing affine function.

Each taxpayer is required to file a report to the IRS. The IRS decides whether or not to audit the report. If her report is not audited, then the taxpayer pays the mandated tax associated with the report. Otherwise the IRS discovers the individual's true income. The individual must then pay taxes corresponding to her true income and also a penalty proportional to the amount that she underreports; the penalty rate, π , is fixed outside the model. With these assumptions, if $t(\cdot)$ is the tax function, then an individual with income i who reports an income of r and is audited pays

¹We make the differentiability assumption for convenience only; monotonicity of $t(\cdot)$ implies that it is differentiable almost everywhere, which is all that is needed for our results.

$t(r) + (1 + \pi)[t(i) - t(r)]$ if $r < i$ and $t(r)$ otherwise. There are no rewards for overreporting income.

The IRS decides which reports to audit. We assume that the probability of conducting an audit depends only on the amount of income reported and that the IRS cannot exceed its budget in expectation.

We now describe the preferences of the actors in our model. Taxpayers are risk neutral and seek to minimize their expected tax payments. The IRS picks an auditing policy that maximizes the revenue it expects to collect. We also discuss the case in which the IRS maximizes net revenue collected. The government selects the auditing budget of the IRS, the amount of spending on public goods, and the tax structure.

It is natural to ask why the government delegates the responsibility to auditing to the IRS. Implicitly we are assuming that it is too expensive for the government to perform the auditing function by itself and that the IRS has the expertise needed to do the job. There is still the possibility that the government could create an incentive scheme that induces the IRS to do something other than maximize the expected revenue collected from auditing. Sánchez and Sobel (1991) present an information structure under which the best thing that the welfare-maximizing government can do is ask the IRS to maximize the expected revenue raised from audits. This result requires the special assumption that the government could observe only the total amount of money that the IRS spends on auditing. It is natural to believe that there are other objectives that the IRS might pursue.

The government sets its tax policy to maximize a function of the after-tax distribution of utility. Formally, we assume that individuals have identical utility functions, $U(A, R) = A + \phi(R)$, defined over after-tax income, A , and the level of public good spending, R . $\phi(R)$ is the monetary value to an individual of R dollars of public expenditure. $\phi(\cdot)$ is strictly increasing, strictly concave, and there is a positive level of expenditure, R^* , such that $\phi'(R^*) = 1$. If v denotes the after-tax distribution of income, then the government maximizes $\int W[U(A, R)] dv$, where $W(\cdot)$ is a twice continuously differentiable function such that $W'(\cdot) > 0 > W''(\cdot)$. The concavity of $W(\cdot)$ reflects the government's concern for income distribution.

3. The IRS's problem

In this section we characterize the solution to the auditor's problem. Given a tax function, $t(\cdot)$, and a budget, B , we find the audit probabilities that maximize the expected revenue collected by the IRS given that it stays within its budget. We also describe the policy of the IRS that maximizes net revenue.

Let $E(i, r; p(\cdot)) = t(r) + (1 + \pi)p(r)[t(i) - t(r)]$ be the expected tax of an individual with income i who reports r when $p(\cdot)$ is the auditing function.

Let $r(i)$ denote the actual report of an individual with income i and let $T(i; p(\cdot)) = E(i, r(i); p(\cdot))$ be the expected tax payment of this individual. The IRS's problem is to find the function $p(\cdot)$ to solve:

$$\max \int_l^h T(i; p(\cdot)) dF(i) \text{ subject to } \int_l^h cp(r(i)) dF(i) \leq B,$$

where, for each i , $r(i)$ solves

$$\min_{r \in [l, i]} E(i, r; p(\cdot)). \quad (\text{RI})$$

We begin with a simple preliminary result.

Lemma [Scotchmer (1986)]. No taxpayer reports more than her income. If $p(r) > 1/(1 + \pi)$, then no taxpayer with income $i > r$ will report income r . A taxpayer reports honestly only if $p(r) \geq 1/(1 + \pi)$ for all $r \in [l, i]$.

Proof. The first claim follows because $t(\cdot)$ is strictly increasing and there are no rebates given to a taxpayer when an audit reveals that income has been overreported. A taxpayer with income $i > r$ who reports r pays expected taxes of $t(r) + (1 + \pi)p(r)[t(i) - t(r)]$, which by assumption is greater than $t(i)$ if and only if $p(r) > 1/(1 + \pi)$. Hence if $p(r) > 1/(1 + \pi)$, then the individual would be better off reporting honestly than reporting r , but if $p(r) < 1/(1 + \pi)$, the individual would gain by lying.

Lemma 1 demonstrates that an audit probability equal to $1/(1 + \pi)$ is sufficient to encourage honest reporting. Since it is costly to audit, the IRS will never choose to audit at a higher probability. In order to guarantee that the IRS's problem has a solution, we assume that taxpayers only distort their income if they strictly gain by doing so. Otherwise the IRS would audit some reports with probability slightly greater than $1/(1 + \pi)$ in order to make honesty the unique best response for some taxpayers. From now on we restrict attention to audit rules $p(\cdot)$ that satisfy $(1 + \pi)p(r) \leq 1$ for all r .

Proposition 1 characterizes the solutions to the taxpayer's problem. The result is well known in the literature on incentives. We use arguments developed by Myerson (1981) and Maskin and Riley (1984) in their analyses of the optimal auction design problem. Morton (1988) treats the same type of problem in an analysis of auditing for employee fraud. The problems studied in these papers differ from ours formally only because of the budget constraint faced by the auditor. Samuelson (1984) characterizes the solution to a problem that contains the same type of constraint; he obtains qualitative

results identical to Proposition 2. The proof of Proposition 1 is little different from the arguments given in these papers, so we omit it.

Proposition 1. Given any $p(r)$, if $r(i)$ maximizes taxpayer i 's expected net income, then

$$p(r(i)) \text{ is nonincreasing in } i, \tag{1}$$

and

$$T(i; p(\cdot)) = T(l; p(r(l))) + \int_l^i p(r(s))(1 + \pi)t'(s) ds. \tag{2}$$

Conversely, if (1) and (2) hold, then $r = r(i)$ maximizes the expected net income of an individual with income i over all reports in the set $\{r(j): j \in [l, h]\}$.²

Proposition 1 allows us to simplify the auditor's problem. Using (2) we can rewrite the objective function of the IRS:

$$\begin{aligned} \int_l^h T(i; p(\cdot)) dF(i) &= \int_l^h [T(l; p(\cdot)) + \int_l^i p(r(s))(1 + \pi)t'(s) ds] dF(i) \\ &= \int_l^h (1 + \pi)t'(i)\gamma(i)p(r(i)) dF(i) + T(l; p(\cdot)), \end{aligned}$$

where $\gamma(i) \equiv [1 - F(i)]/f(i)$ is the hazard rate and the last equality follows from changing the order of integration. Consequently, the optimization problem of the IRS becomes:

$$\begin{aligned} \max \int_l^h (1 + \pi)t'(i)\gamma(i)p(r(i)) dF(i) + T(l; p(\cdot)) \\ \text{s.t. } p(r(i)) \text{ nonincreasing in } i, \quad p(r(i)) \in [0, 1/(1 + \pi)], \\ \int_l^h cp(r(i)) dF(i) \leq B, \quad \text{and for each } i, \quad r(i) \text{ solves (RI)}. \end{aligned}$$

We can simplify the problem further. Since taxpayers with income l can do

²(1) and (2) do not restrict $p(r)$ for reports made by no individual. As a result, (1) and (2) are not sufficient to conclude that $r = r(i)$ maximizes the expected net income of an individual with income i over all reports. If (1) and (2) hold, then we can modify $p(r)$ on $\{r: r \neq r(j) \text{ for all } j \in [l, h]\}$, so that $r = r(i)$ is taxpayer i 's best report for all i .

no better than report their income honestly, $T(l; p(\cdot)) = t(l)$ independent of the auditor's policy. Also, (2) tells us how to find $r(i)$ knowing $p(r(\cdot))$ provided that $(1 + \pi)p(r(i)) < 1$. Since we have $r(i) = i$ whenever $(1 + \pi)p(r(i)) = 1$ by Lemma 1, the IRS solves its optimal auditing problem by finding a function $q^*(\cdot)$ to solve:

$$\begin{aligned} & \max \int_l^h (1 + \pi)t'(i)\gamma(i)q(i) dF(i) \\ & \text{s.t. } q(i) \text{ nonincreasing in } i, \quad q(i) \in [0, 1/(1 + \pi)], \quad \text{and} \quad \int_l^h cq(i) dF(i) \leq B. \end{aligned} \tag{Q}$$

This problem is a straightforward optimization problem in the unknown function $q(\cdot)$; $q(i)$ is the probability that an individual with income i is audited. Propositions 2 and 3 describe the solution to this problem. Since the basic qualitative results are well known [for example, see Myerson (1981), Maskin and Riley (1984), Morton (1988), or Samuelson (1984)], we omit the proof. We are not familiar with a complete treatment of the problem in the literature, so we refer the interested reader to Sánchez and Sobel (1991) for technical details.

Proposition 2. (a) *If $t'(i)\gamma(i)$ is strictly decreasing, then the auditor's problem (Q) of maximizing net revenue subject to a budget constraint has a unique solution. The expected revenue-maximizing auditing function $p^*(\cdot)$ satisfies*

$$p^*(i) = \begin{cases} 1/(1 + \pi), & \text{if } i < a, \\ 0, & \text{if } i \geq a, \end{cases}$$

where $F(a) = B(1 + \pi)/c$ if $B(1 + \pi) < c$ and $a = h$ otherwise.

(b) *There always exists a solution to the auditor's problem of the form:*

$$p^*(i) = \begin{cases} 1/(1 + \pi) & \text{if } i \in [l, a), \\ d, & \text{if } i \in [a, b), \text{ where } l \leq a \leq b \leq h \text{ and } d \in (0, 1/(1 + \pi)), \\ 0, & \text{if } i \in [b, h] \end{cases}$$

The corollary describes the qualitative implications of Proposition 2.

Corollary 1. *There exists an optimal auditing policy in which the taxpayers are divided into at most three groups. Each member of the lowest income group reports her true income and is audited with probability $1/(1 + \pi)$. All individuals in the middle group make the same report, which is less than their income, and*

are audited with probability $d \in (0, 1/(1 + \pi))$. Individuals with the highest income all report the smallest income that is never audited.

Given the auditing function described in Proposition 2, the Corollary follows from (2).

Proposition 3. *If the auditor acts to maximize the expected value of taxes and penalties collected net of auditing costs, then there always exists a solution to the auditor's problem of the form:*

$$p^*(i) = \begin{cases} 1/(1 + \pi), & \text{if } i < a, \\ 0, & \text{if } i \geq a. \end{cases}$$

The solution is unique if $t'(i)\gamma(i)$ is monotonically decreasing.

In case (a) of Proposition 2 the solution to the auditor's problem is unique and has a simple form. The IRS audits only the low reports and does so frequently enough to encourage honest reporting from taxpayers with incomes below a . All higher reports are never audited. If an individual with income greater than or equal to a reports that her income is a , then she pays $t(a)$ in taxes. If $t'(i)\gamma(i)$ is strictly decreasing, then the monotonicity constraint on $q(\cdot)$ does not bind at the optimum. If taxes are linear, the condition reduces to the decreasing hazard rate assumption frequently made in agency models.

If $t'(i)\gamma(i)$ is not strictly decreasing, then we cannot guarantee that the solution of the IRS's problem is unique. Nor can we guarantee that it has a simple cut-off form. Part (b) of Proposition 2 states that optimal audit functions need not be too complicated.

One notable property of the solution is the monotonicity of the audit function. Border and Sobel (1987) show that this property also holds when the government is able to pick taxes and penalties and directly controls the auditing process [(Mookherjee and Png (1989) show by example that monotonicity may fail to hold when taxpayers are risk averse)]. When taxes are increasing in reported income the monotonicity of the audit function has a straightforward intuition: low reports are the most attractive ones. The IRS must audit these reports frequently in order to discourage large underreports of income. The intuition does not take into account the possibility that the penalty for tax evasion depends on the amount of income underreported. If the marginal penalty for underreporting income is increasing, then the auditing rules need not be decreasing in reported income. As an example, imagine that the economy consisted primarily of two income groups, rich and poor, with the lowest income rich individual having a significantly

greater income than the wealthiest poor person. Suppose that the judicial system has established penalties that heavily (as a fraction of income) punish a rich person when an audit reveals that she has reported a poor person's income, but only levies a small (proportional) fine if income is only slightly misrepresented. It may be the case that the IRS maximizes revenue by auditing reports of income approximately equal to the poorest rich person. Heavy penalties and light auditing could be enough to discourage the rich from reporting a poor person's income. The same possibility arises if the IRS is able to divide taxpayers into rich and poor at fairly low cost. The optimal policy might include a preliminary screening to divide the rich from the poor, followed by a relatively heavy auditing of individuals who make the lowest reports in each income class.

Another feature of the optimal auditing policy is that it leads to effective taxes that are more regressive than the nominal taxes. Cr mer, Marchand and Pestieau (1990) make a similar observation, while Scotchmer (1991) demonstrates it generally. The essential idea is that a high income individual has a greater opportunity to hide income than a low income individual. For example, if $t(i) = \tau i$ and $\gamma(\cdot)$ is strictly decreasing, then the effective tax is equal to τi if $i \leq a$ and τa otherwise, so a tax function that is nominally proportional is effectively regressive at high incomes.

Another property of the optimal auditing policy follows from the form of (Q). An increasing affine transformation of the tax function multiplies the objective function in (Q) by a positive constant. Therefore such a change does not affect the strategy of the IRS. In particular, if the government uses an affine tax function, then the auditing policy of the IRS depends only on its budget and not on the marginal tax rate.

Propositions 2 and 3 state that we can simplify the auditor's problem by restricting attention to the small class of feasible audit functions that take on at most one value other than 0 and $1/(1 + \pi)$. Hence the budget-constrained auditor solves:

$$\max \int_l^a t'(i)\gamma(i) dF(i) + d(1 + \pi) \int_a^b t'(i)\gamma(i) dF(i) \tag{BC}$$

$$\text{s.t. } l \leq a \leq b \leq h,$$

$$0 \leq d \leq 1/(1 + \pi), \tag{3}$$

and

$$F(a)c/(1 + \pi) + [F(b) - F(a)]cd \leq B. \tag{4}$$

If the budget constraint (4) does not bind, then the auditor has enough

resources to audit everyone with probability $1/(1+\pi)$. If either of the constraints in (3) bind or if $a=b$, then the IRS audits with probability $1/(1+\pi)$ as many of the lowest returns that it can afford. Otherwise the first-order conditions of (BC) imply that at an interior solution

$$t'(a)\gamma(a) = t'(b)\gamma(b) = \int_a^b t'(i)\gamma(i) dF(i) / [F(b) - F(a)]. \quad (5)$$

Since the first equation in (5) could not hold if $t'(\cdot)\gamma(\cdot)$ is strictly monotonic, we know that the optimal policy of auditor takes on a value different from 0 or $1/(1+\pi)$ only if $t'(\cdot)\gamma(\cdot)$ fails to be monotonic. It is not difficult to construct examples in which the auditor's optimal policy must take on three values.

The net revenue maximizing auditor need only look for a cutoff value a^* that solves:

$$\max_a \int_a^b [(1+\pi)t'(i)\gamma(i) - c] dF(i). \quad (\text{NR})$$

The first-order condition for (NR) is

$$t'(a)\gamma(a) = c/(1+\pi). \quad (6)$$

The monotonicity assumption in Proposition 3 implies that the objective function in (NR) is strictly concave. Since (6) characterizes the solution to (NR) when $t'(\cdot)\gamma(\cdot)$ is strictly decreasing, it can be used to describe how the revenue-maximizing amount of auditing changes when the tax function or the distribution of income changes. We note the obvious properties: increases in the marginal tax rate or upward translations of the income distribution increase the amount of auditing; increases in c or decreases in π reduce the amount of auditing.

Examination of (Q), (BC), or (NR) shows that the auditor's problem depends on c and π only through the ratio $c/(1+\pi)$ [replacing d by $d'=(1+\pi)d$ makes this apparent in (BC)]. Raising penalties for under-reporting effectively reduces the cost of an audit.

The sharp characterization that we provide in this section depends strongly on the form of the objective function of the IRS and the risk neutrality of the taxpayers. If individuals preferences are linear and the auditor's objective function is linear in $T(i; p(\cdot))$, then the techniques of this section apply. We can say nothing about the optimal audit function if either taxpayers are risk averse or the auditor's objective function is nonlinear.

4. The optimal hierarchical policy

We are now ready to describe the solution to the hierarchical tax design problem. We concentrate on the case where the government sets affine taxes and the hazard rate is strictly decreasing; we discuss in detail the special case in which taxes must be linear. We make two different assumptions on the control that the government exercises on the IRS. In the first part of the section we assume that the government selects the tax function and the IRS's budget; the IRS then selects the auditing policy that maximizes expected revenue subject to a budget constraint. At the end of the section we describe how the analysis changes when the government does not control the budget of the IRS. In this case we assume that the IRS maximizes net revenue.

When $t(\cdot)$ is affine and the hazard rate is decreasing, Propositions 2 and 3 guarantee that the IRS's problem has a unique solution, which takes a particularly simple form. The IRS uses its budget to audit with probability $1/(1+\pi)$ all of the lowest reports. When the IRS has a budget of B , it audits every report in the interval $[l, a]$, where

$$F(a(B))c/(1+\pi) = B \quad (7)$$

[if $B \geq c/(1+\pi)$, then the IRS can afford to audit with sufficient frequency to force honest reporting by everyone; budgets greater than $c/(1+\pi)$ are never optimal]. As we noted in section 3 the IRS's auditing policy is independent of the tax rate. The after-tax income of taxpayers is deterministic, since only honest reports are audited in equilibrium.

We write the tax function $t(i) = \alpha + \tau i$ (α may be negative) and denote by

$$A(i; B) = \begin{cases} i - t(i), & \text{if } i \leq a(B), \\ i - t(a(B)), & \text{if } i > a(B), \end{cases} \quad (8)$$

the after-tax income of an individual with income i . Since the population size has been normalized to unity, the per capita revenue raised by this tax policy is

$$M(B, t) = \alpha + \tau \int_l^{a(B)} i dF(i) + \int_{a(B)}^h \tau a dF(i). \quad (9)$$

The government must pick a tax function, $t(\cdot)$, a budget, B , and a level of public goods expenditure, R , to solve

$$\max \int_l^h W[u(i)] dF(i) \tag{GP}$$

s.t.

$$R + B = M(B, t), \tag{10}$$

$$\alpha + \tau l \leq l, \tau \leq 1, \tag{11}$$

(7), (8), and (9).

where we use $u(i) = \phi(R) + A(i; B)$ to denote the after-tax utility of an individual with income, i . We suppress the dependence of $u(\cdot)$ on the tax function, auditing policy, and R to simplify notation. The constraints in (11) guarantee that nominal taxes are no greater than income and that the marginal tax rate is no greater than one.

We treat the optimization problem as if the government chooses a , τ , and α , with a determining B through (7), and (10) determining R . It is straightforward to write the first-order conditions for this problem. Expressions (12), (13), and (14) are the derivatives of the objective function in (GP) with respect to α , τ , and a :

$$[-1 + \phi'(R)] \int_l^h W'[u(i)] dF(i), \tag{12}$$

$$H(a) \left\{ - \int_l^h W'[u(i)] dF(i) + \int_l^h W'[u(i)] dG(i) \right\} \\ + (a - H(a)) (\phi'(R) - 1) \int_l^h W'[u(i)] dF(i), \tag{13}$$

and

$$\tau \left\{ \phi'(R) D(a, \tau) \int_l^h W'[u(i)] dF(i) - \int_a^h W'[u(i)] dF(i) \right\}, \tag{14}$$

where $H(a) = aF(a) - \int_l^a i dF(i)$, $D(a, \tau) = 1 - F(a) - f(a)c/\tau(1 + \pi)$, and

$$G(i) = \begin{cases} \int_l^i (a - s) dF(s)/H(a), & \text{if } i \leq a, \\ 1, & \text{if } i > a. \end{cases}$$

We begin the discussion by assuming that the first constraint in (11) does not bind. Since we assume that $W'(\cdot)$ is strictly positive, the expression in (12) can be equal to zero only if $\phi'(R)=1$. This result has a clear interpretation. The constant term α of the affine tax structure is a lump-sum transfer. When α is positive, all pay a lump sum to the government; when α is negative, the transfer goes to the agents. Hence the tax schedule is progressive if $\alpha < 0$. Expression (12) states that it should be adjusted in order to set government spending at the level R^* (recall $\phi'(R^*)=1$). At this level of public spending the value of an additional dollar of tax revenue is exactly one dollar. When lump-sum transfers are feasible the inability to observe income does not distort public goods spending.

Expression (13) is the derivative of the objective function of (GP) with respect to τ . The first term of the derivative is the difference of expectations of $W'(\cdot)$ with respect to $F(\cdot)$ and $G(\cdot)$. (13) is positive whenever $\phi'(R) \geq 1$. To see this, note that $F(i) \leq G(i)$ for all i [check this by observing that the density of $G(\cdot)$ crosses that of $F(\cdot)$ exactly once and from above]; that is $F(\cdot)$ first-order stochastically dominates $G(\cdot)$. Since $W'(\cdot) > 0$,

$$-\int_l^h W'[u(i)] dF(i) + \int_l^h W'[u(i)] dG(i) \geq 0. \quad (15)$$

Equality in (15) holds if and only if $u(i)$ is constant, which in turn requires that $\tau=1$ and $a=h$. The solution to the optimization problem must involve $\tau=1$. Again this result is intuitive given the structure of the problem. We have assumed that taxpayers do not choose effort levels. Consequently, a change in the marginal tax rate does not influence the amount of taxable income available to the government. Since the optimal auditing policy of the IRS does not depend on the marginal tax rate, increasing τ while holding the rest of the tax and auditing policy constant serves to redistribute income: The additional revenue obtained by this policy provides benefits to taxpayers that are independent of income while higher income taxpayers pay absolutely more for these benefits.

Finally, we discuss (14), which measures the response of the objective function of (GP) to changes in the level of auditing a . We can use this expression to show formally that the budget provided by the government is always less than the net revenue-maximizing budget. Intuitively this result is simple. The IRS will spend an extra dollar if by doing so it raises an extra dollar of revenue. The government uses its revenue first to pay for auditing and then to redistribute income or pay for public goods. At the level of auditing that is optimal for the IRS, a one dollar increase in budget leads to precisely a one dollar increase in revenue collected. At this point the government is unable to redistribute income since all revenue collected goes to cover the cost of auditing. Consequently there can be no distributional

advantage to auditing at this point. Slemrod and Yitzhaki (1987) identify similar effects.

We can arrive at the same conclusion with algebraic arguments. (14) is negative when $a=h$; hence some taxpayers are not audited in equilibrium. Also, $D(a, \tau)$ must be positive at the solution to the government's problem. If $D(a, \tau)$ were less than or equal to zero, then (14) would be negative, therefore it would be in the interest of the government to reduce the amount of auditing. We know from (6) that the net revenue-maximizing cutoff level for the IRS given the marginal tax rate τ would set $D(a, \tau)$ to zero. Hence the government provides less for auditing than a net revenue-maximizing auditor would wish to spend.

If $R=R^*$, then $f(l)>0$ implies that (14) is negative when $a=l$. While we have no simple conditions that imply that (14) is strictly decreasing in a , the fact that (14) is negative for $a=l$ suggests the possibility that it is optimal for the government to allocate no money for auditing. A closer investigation of the structure of the problem explains why it is not in the interest of the government to raise a slightly above l . If the government is able to raise revenue, $R=R^*$, when $a=l$, then the lump-sum component of taxes is sufficient to raise enough revenue to guarantee support for the optimal level of public good spending. Hence, the only reason to audit is to redistribute income: raising a allows the government to increase the size of the lump-sum transfer. An infinitesimal increase in a from l has a negative redistributive effect, however. The revenue gained by the increase in auditing comes from taxpayers of all incomes equally. Therefore the net income of the lowest individual cannot increase. In fact, the net income must fall because a costly means of raising revenue has replaced the lump-sum tax. This argument requires that the government can raise the optimal amount of revenue, R^* , without auditing – an unlikely case. Examination of (14) shows that if $\phi'(l)$ is sufficiently large (when there is no auditing the most that the government can raise is $\alpha=l$), then (14) will be positive when $a=l$.

We now consider the case where $t(l)=l$ at the optimum. It follows from (12) that the level of spending for public goods is no greater than R^* . Furthermore, (13) still implies that $\tau=1$.³ Intuitively the result follows because, if auditing is costly, it is always advantageous for the government to raise money through increases in the marginal tax rate rather than lump-sum transfers. If $\alpha+\tau l=l$ and $\tau<1$, then the government can increase τ and lower α in a way that leaves tax revenue unchanged, but raises the taxes collected from high income individuals. Since the optimal auditing policy does not depend on τ and α , the government can increase the marginal tax rate without changing the number of evaders. Because the government seeks to

³One can prove the result by letting ξ_1 and ξ_2 be the non-negative multipliers for the constraints in (11), noting that at the optimum the derivative (12) must be equal to ξ_1 and the derivative (13) must be equal to $\xi_1 l + \xi_2$, and that $a-H(a) \geq l$.

equalize after-tax incomes, increasing τ is attractive. (14) implies that if $\phi'(l)$ is sufficiently large, then there will necessarily be positive auditing in equilibrium.

Even if $W'(0) = \infty$, the government may wish to force an individual with income l to pay all of her income in taxes. The form of the social welfare function allows her to be repaid with public goods since, provided that $R > 0$, even a taxpayer with no after-tax income obtains positive utility.

Since the government maximizes a continuous function over a compact set, a solution to its problem exists. We summarize the properties of the solution in the next proposition.

Proposition 4. Assume that the hazard rate is monotonically decreasing and that the government can levy taxes of the form $t(i) = \alpha + \tau i$. The government's optimization problem has a solution. The qualitative properties of the solution include:

- (a) The marginal tax rate is $\tau = 1$ and there are non-negative lump-sum transfers.
- (b) The auditor is underbudgeted.
- (c) The government spends no more than R^* on public goods.
- (d) If $t(l) < l$, then the government spends exactly R^* on public goods.

The proposition does not provide conditions under which an interior solution does or does not exist. Sufficient conditions are easy to provide. If $l > R^*$, then the government can provide the optimal level of spending using lump-sum taxes alone. Indeed, in this case it conducts no audits. If the average income, μ , is less than R^* , then the government could not provide for all of its public goods projects even if it could costlessly extract the entire income of every taxpayer. In this situation the lowest income taxpayers would pay everything to the government. If l is small relative to R^* , then the government would need to conduct audits in order to raise tax revenue.

The discussion above does not depend strongly on the choice of the objective function for the government or on the assumption that the government picks R . The qualitative properties described in Proposition 4 are true if the government's objective is to maximize $\phi(R) + \int_i^h V[A(i)] dF(i)$, where $A(i)$ is the after-tax income of agent i and $V(\cdot)$ is increasing and concave, or a weighted average of individual utilities, say $\phi(R) + \int_i^h \lambda(i) A(i) dF(i)$ [R^* needs to be redefined in parts (c) and (d), but it should still be interpreted as the government's optimal level of spending for public goods]. Parts (a) and (b) of the proposition continue to hold if the government levies taxes to pay for a fixed expenditure on public goods (and, in order to guarantee a solution to the government's problem, that the IRS can collect enough to pay for the necessary public goods).

The results of Proposition 4 may be misleading because they depend on

the government's ability to use lump-sum transfers. When costless lump-sum transfers are available, there is no reason to use public good spending to redistribute income. Since a dollar spent on public goods in excess of R^* is worth less than one dollar, it is wiser simply to redistribute this income using the lump-sum component of taxes. We would not expect parts (c) and (d) of Proposition 4 to hold without the lump-sum transfers.

To confirm this intuition, we discuss the government's optimal policy when it is constrained to use linear taxes, that is $\alpha=0$. We can ignore condition (12). Expression (13) represents the effect of an increase of τ on welfare. At the optimum either (13) is positive and $\tau=1$ or (13) is zero.

First consider the case in which $\tau < 1$ at the optimum. It follows from (15) that $\phi'(R) < 1$ or, equivalently, $R > R^*$. The result contrasts with Proposition 4(c) and (d). If lump-sum transfers are possible, then the government is able to use them to redistribute income. If they are not feasible, then the government may compensate by spending more on public goods. The marginal dollar spent on a public good influences all individuals in the same way, while an increase in τ leads to higher tax collections from high income individuals. So the government is willing to raise taxes to pay for public projects beyond the first-best level R^* .

It is also possible to have $\tau=1$ at the optimum. In the unlikely event that the mean income is less than R^* , there is not enough income in the economy to support the first-best level of spending on public goods. Since $R \geq R^*$ is not possible, it must be that the optimal policy involves $R < R^*$ and $\tau=1$.

As before, there is a possibility that the government audits no returns. To rule this out, we may assume that the government must raise more than l (the maximum level of tax revenue that can be raised without auditing) in order to provide an adequate level of public goods. Assumptions on either $\phi(\cdot)$ or $W(\cdot)$ would be sufficient.⁴

Proposition 5. Assume that the hazard rate is monotonically decreasing and that the government can levy taxes of the form $t(i)=\tau i$. The government's optimization problem has a solution. The qualitative properties of the solution include:

- (a) *The auditor is underbudgeted.*
- (b) *Either $\tau=1$ or the government spends more than R^* on public goods.*

Again the results cited in Proposition 5 do not depend strongly on the form of the objective function that we chose. In fact, other ways of posing the problem lead to simple conditions that guarantee that the government's optimal tax rate is less than one.

⁴One sufficient condition is to assume that $\phi'(\tau l)D(l, \tau) = \phi'(\tau l)[1 - f(l)c/\tau(1 + \pi)] > 1$ for all $\tau \in (0, 1]$.

We now discuss the impact of the government's incomplete ability to observe incomes. How does the presence of auditing costs influence the government's optimal policy? We discussed one property in the previous section: The optimal auditing policy causes the effective tax rate to be more regressive than the nominal tax rate. Now we discuss the effect of a change in the cost of auditing on the policy of the government.

An increase in the cost of auditing makes it more expensive to collect taxes. Therefore it must lead to a decrease in the amount of auditing, a decrease in the spending for public goods, or an increase in the marginal tax rate. In general, we are unable to say which of these effects occur. We have no guarantee that the first-order conditions characterize the solution to the government's problem and the comparative-static results are ambiguous in any event.

The comparative-statics exercise is simpler when we examine how the government's policies change when the cost of auditing increase from zero. When $c=0$ the government audits all of the returns. Consequently, an increase in c must lead to a reduction of auditing. Computation, using (13) and (14), shows that at an interior solution.⁵

$$(1 + \pi) \frac{\partial \tau}{\partial c} = \frac{\int_i^h \{ W''[u(i)] \phi'(R) [\mu \phi'(R) - i] + W'[u(i)] \mu \phi''(R) \} dF(i)}{\int_i^h \{ W''[u(i)] [\mu \phi'(R) - i]^2 + W'[u(i)] \mu^2 \phi''(R) \} dF(i)}, \quad (16)$$

$$(1 + \pi) \tau \frac{\partial a}{\partial c} = \frac{\phi'(R) \int_i^h W'[u(i)] dF(i)}{W'[u(h)] - \phi'(R) \int_i^h W'[u(i)] dF(i)}, \quad (17)$$

and

$$(1 + \pi) \frac{\partial R}{\partial c} = -1 + (1 + \pi) \mu \frac{\partial \tau}{\partial c} \\ = \frac{\int_i^h W''[u(i)] \phi'(R) [\mu \phi'(R) - i] i dF(i)}{\int_i^h \{ W''[u(i)] [\mu \phi'(R) - i]^2 + W'[u(i)] \mu^2 \phi''(R) \} dF(i)}, \quad (18)$$

where $\mu = \int_i^h i dF(i)$ is the mean income and the utility of an individual with income i is $u(i) = \phi(R) + i(1 - \tau)$. Since

$$\int_i^h W'[u(i)] [\mu \phi'(R) - i] dF(i) = 0 \quad (19)$$

at the optimum by (13) and $W'(\cdot)$ is decreasing, the denominator of (17) is negative. This shows that the amount of auditing decreases as c increases

⁵It follows from (13) that if $\int_i^h i dF(i) > R^*$, then $\tau < 1$.

from zero. The result is not remarkable since everyone is audited when $c=0$. The concavity of $\phi(\cdot)$ and $W(\cdot)$ imply that the denominators of (16) and (18) are negative. To sign $\partial R/\partial c$ note that

$$\begin{aligned}
 & -\int_l^h W''[u(i)]\phi'(R)[\mu\phi'(R)-i]i dF(i) \\
 & = \int_l^h \{-W''[u(i)]i/W'[u(i)]\}W'[u(i)]\phi'(R)[\mu\phi'(R)-i] dF(i). \quad (20)
 \end{aligned}$$

We can write

$$-iW''[u(i)]/W'[u(i)] = [i/u(i)]\rho_w[u(i)], \quad (21)$$

where $\rho_w(i) \equiv -iW''(i)/W'(i)$ is the coefficient of relative risk aversion. Eq. (19) implies that (20) is negative if $-W''[u(i)]i/W'[u(i)]$ is increasing. Therefore, since $i/u(i)$ is positive and increasing and $\rho_w(\cdot)$ is positive, if we assume that $\rho_w(\cdot)$ is nondecreasing, then we can conclude that public spending decreases as the cost of auditing increases from zero. Here is a rough intuition for the result. Imagine that the government is a single decision-maker with uncertain income. Increasing c effectively makes the decision-maker poorer for each realization of income. If the decision-maker has increasing relative risk aversion, then an increase in c causes the individual to be less averse to proportional risks. So she would be less willing to insure against those risks. When taxes are linear, government spending plays a role similar to partial insurance against a proportional risk. If $\rho_w(\cdot)$ is nondecreasing, then demand for government spending decreases as c increases.

The result emphasizes the two roles of tax policy in our model. The government collects taxes to pay for public goods and uses its tax policy to redistribute income. When the cost of auditing increases, the cost of public goods increases. At the same time, effective tax rates become more regressive. Consequently, the government may wish to increase spending on public goods when auditing costs go up in order to redistribute income.

The possibility that an increase in the cost of auditing leads to an increase in spending for public goods depends on our specification of the objective function. If the government chose to maximize $\phi(R) + \int_l^h V[u(i)] dF(i)$, then straightforward computations show that

$$(1 + \pi)\tau \frac{\partial a}{\partial c} = \phi'(R)/[V'(u(h)) - \phi'(R)] < 0,$$

$$(1 + \pi)\tau \frac{\partial \tau}{\partial c} = \frac{\mu \phi''(R)}{\int_i^h i^2 V''[i(1-\tau)] dF(i) + \mu^2 \phi''(R)} > 0,$$

and

$$(1 + \pi)\tau \frac{\partial R}{\partial c} = \frac{-\int_i^h i^2 V''[i(1-\tau)] dF(i)}{\int_i^h i^2 V''[i(1-\tau)] dF(i) + \mu^2 \phi''(R)} < 0,$$

so, while an increase in auditing costs from $c=0$ leads to an unambiguous increase in the marginal tax rate, it also leads to an unambiguous drop in the spending for public goods.

Now let us turn to the case where the government can only select the tax function. The IRS controls the auditing policy by choosing an auditing policy to maximize expected net revenue. Assuming that the government sets affine taxes and the hazard rate is strictly decreasing, Proposition 3 guarantees that the IRS's problem has a unique solution; it audits all of the lowest reports with probability $1/(1+\pi)$ and none of the higher reports.

The difference between the net revenue-maximizing policy and the policy induced when the government selects the auditor's budget is that to maximize net revenue the IRS selects the cutoff income a . When the hazard rate is strictly decreasing, (6) characterizes the IRS's policy. The solution to the hierarchical problem is described by (6), the derivative of the government's objective function with respect to α (12), and

$$\begin{aligned} H(a) \left\{ -\int_i^h W'[u(i)] dF(i) + \int_i^h W'[u(i)] dG(i) \right\} \\ + (a - H(a))(\phi'(R) - 1) \int_i^h W'[u(i)] dF(i) \\ + [\gamma(a)/\gamma'(a)] \int_a^h W'[u(i)] dF(i), \end{aligned} \quad (22)$$

which is the first derivative of the government's objective function with respect to τ . Expression (22) adds to (13) a term that reflects the dependence of a on τ . Parts (c) and (d) of Proposition 4 follow as before. Of course, since the auditor controls its spending on auditing, it will allocate the net revenue-maximizing amount; part (b) of the Proposition is no longer true.

It need not be true that $\tau=1$. Crémer, Marchand and Pestieau (1990) also assume that the IRS seeks to maximize net revenue. They show that the marginal tax rate need not be one in their model, in contrast to (a) in Proposition 4. The reason for the difference lies in the failure of part (b) in

the proposition. The auditor always wishes to spend more on auditing than does the government. In the setting of Proposition 4 this tendency can be controlled by reducing the budget of the auditor. When the IRS maximizes net revenue, this instrument is unavailable. Instead, the level of auditing must satisfy $\tau\gamma(a) = c/(1 + \pi)$ [eq. (6)]. Hence the marginal tax rate is distorted downward to reduce the level of auditing. Since providing a budget that allows the IRS to select its net revenue-maximizing policy is feasible in our problem, the government strictly benefits from controlling the auditor's budget.

We have little to say about the solution to the government's problem when the hazard rate is not monotonic. In principle one could characterize solutions as we have done in this section. A problem that arises is that while the auditor maximizes revenue by using a three-step auditing policy, there may exist more complicated policies that the government prefers.

While we did not talk about optimal nonlinear taxes for the government, arguments that force $\tau = 1$ suggest that the ability to use nonlinear taxes does not aid the government.

5. The effect of delegation on the government's policy

We described the government's behavior when only the IRS had the ability to set auditing policy in the previous section. If the government has direct control over auditing, then it can do at least as well. Since the incentives of the IRS differ from the government's, one would expect that if the government had direct control over the auditing policy, then social welfare would strictly increase. In this section we discuss the effect that the government's limited ability to control the auditing process has on tax policy.

In section 4 we discussed two hierarchical problems. In one problem the government could control the budget of the IRS; in the other it could not. These problems have different solutions. Plainly, the social welfare is strictly higher when the government controls the budget of the IRS. When it does so, we know that less will be spent on auditing.

We would like to show that if the government could control not only the budget of the IRS but the entire auditing policy, then social welfare would strictly increase. Unfortunately, we are unable to find the government's optimal policy when it is directly responsible for conducting audits.⁶ We limit ourselves to a discussion of what could happen.

⁶Mookherjee and Png (1989) analyze the problem of a government that picks taxes and audit policy to maximize social welfare subject to a net revenue constraint. They describe qualitative features of the optimum, but their characterization is not sharp enough to allow a comparison with the hierarchical policy of section 4. Although we restrict attention to linear taxes and a particular form of penalty function we cannot completely characterize the optimum.

Since the auditor maximizes tax revenue subject to a budget constraint, it collects funds efficiently. Consequently, if the government's direct policy differs from the hierarchically optimal one, then it must audit more individuals or raise taxes in order to collect the same net revenue. These changes are desirable only if they lead to an increase in social welfare by inducing more of the tax payments to come from high income individuals.

The optimal hierarchical policy leads to after-tax incomes that are deterministic (*ex ante*). Border and Sobel (1987) and Mookherjee and Png (1989) give conditions under which individuals are uncertain about their after-tax incomes in the solution to the problem in which the government picks both tax and audit policy. While these results suggest that direct control leads to a change in policy, so must be strictly beneficial to the government, they are not relevant to our discussion because they allow the government to use a richer class of tax and penalty functions. There is no reason to believe that the IRS's optimal policies will cause after-tax incomes to be deterministic in the more general case.

6. Limitation and extensions

The most restrictive assumption that we made is that taxpayers are risk neutral. As the auditing policy or taxable income is unlikely to be deterministic [or known by the taxpayers, see Scotchmer and Slemrod (1987) and Scotchmer (1989)], the risk attitudes of individuals should play an important role in determining the qualitative features of the tax system. Adding risk aversion puts the characterization problem out of our reach. Few general statements can be made about the solution to the auditor's problem, let alone the hierarchical design problem that we study, when the taxpayers are risk averse [see Mookherjee and Png (1989) for an analysis of a welfare-maximizing government that is able to set both taxes and an auditing policy and in which taxpayers are risk averse]. We hope that our analysis provides insight into the conflict between the IRS and the government nevertheless.

Our study also neglects the labor supply problem. Some of our results (especially the marginal tax rate of one that we obtain in section 4) would certainly not hold if individuals selected their labor supply in response to the tax structure. Many papers, starting with Mirrlees (1971), investigate the structure of optimal tax policies when tax evasion is not possible but individuals choose their labor supply. Kanodia (1985), Sánchez (1987, ch. 2) and Mookherjee and Png (1989), have examined tax auditing problems in the presence of moral hazard; we are unaware of any work on the hierarchical design and enforcement problems where taxpayers choose labor supply.

Our auditing technology is special. The IRS either knows only the reported income of a taxpayer or, on the basis on an audit, the taxpayer's

income exactly. Evans (1980), Baron and Besanko (1984), Laffont and Tirole (1986), and other presents models in which auditing is feasible but not perfect. For our study of tax auditing, a more realistic framework would allow the IRS to divide the population into income classes at little or no cost (tax returns can be sorted by the zip codes and occupation of the individual, or even sorted on the basis of age or sex). We can interpret the IRS's policy as applying to a particular income class, but our analysis does not discuss how the government can or should discriminate over different groups. Scotchmer (1986) and Crémer, Marchand and Pestieau (1990) discuss inequities that arise if the auditor is able to divide individuals into different income classes. Pestieau and Possen (1991) show that if the cost of auditing depends on the source of income, then the auditing policy influences individuals' choice of occupation.

We assume that the auditor can commit to an auditing policy. Graetz, Reinganum and Wilde (1986) and Reinganum and Wilde (1986) examine the optimal auditing policy when commitment is not possible. Melumad and Mookherjee (1989) explain how the auditor can implement its full commitment auditing policy if delegation is possible.

The relationship between the government and the IRS could be enriched by explicitly modeling the auditor's choice of effort. Demski and Sappington's (1987) hierarchical model studies an organization where the middle level supervisor must be induced to work.

Our model is static. Landsberger and Meilijson (1982) and Greenberg (1984) study models of tax auditing in which the audit function depends on an individual's past behavior. Harris (1986) analyzes an infinite-horizon model of optimal taxation with moral hazard and observable incomes. Without enriching the relationship between the government and the IRS there is nothing to be gained by studying a multiperiod model of their relationship.

On a technical level, we have placed several ad hoc restrictions on the model. First, we assume that the IRS's auditing policy depends only on the individual characteristics of a taxpayer. More general descriptions are possible. To take an extreme case, suppose that the population is finite, and that the IRS knows the empirical distribution of income. The auditor could announce that everyone will be audited with probability one if the realized distribution of reports differs from the empirical distribution. It is an equilibrium for every taxpayer to tell the truth (there are others), so the IRS can obtain total compliance at no cost. While we think the policy is ridiculous, its major flaws are not captured by our model. The procedure works badly if some taxpayers are uncertain about their income, if the auditor is uncertain about the distribution of income, if there is a possibility that a report will be misinterpreted, or if implementing a policy that requires centralized comparison of returns is more costly than audits triggered by

individual reports. If auditing is decentralized so that areas or income classes are treated separately and audited with independent budgets, then the large numbers assumption may be inappropriate. In these situations auditing policies will be based on more than an individual's report. For example, if taxpayers file returns at different times, then the probability of auditing a particular return may depend nontrivially on when reports are filed.

Another technical limitation of our analysis is that we require that the budget constraint of the IRS hold only in expectation. We are implicitly appealing to the law of large numbers. We have not formally supported this assumption.

We have also restricted attention to a special class of tax policies. A complete model would allow the government to set the fines for underreporting income. If individuals are risk neutral and auditing is perfect, then it is well known that cheating can be deterred at a tiny cost by using large penalties [see Becker (1968), Stigler (1970), Kolm (1973) or, for more general results, Mirrlees (1975)]. Even if feasible fines are bounded below, Border and Sobel (1987) show that the auditor's problem may not have a solution if it is possible to give rewards to honest taxpayers. Mookherjee and Png (1989) give conditions on taxpayers' preferences which guarantee that the auditor's problem has a solution.

A more general approach is possible. The government could design a mechanism in which it defines the set of possible reports that taxpayers could make to the IRS and determines taxes and penalties contingent on the reports and results of the audit. Often one can appeal to the revelation principle in order to show that there is no loss in generality in assuming that the set of possible reports are income levels and that the government limits attention to mechanisms in which taxpayers optimize by reporting their true incomes. Border and Sobel (1987) and Mookherjee and Png (1989) use this approach in models where one agency sets both the schedule and the auditing policy. Example show that the revelation principle does not hold in hierarchical settings. It may be optimal for the government to restrict the set of possible reports in order to limit the authority of the auditor.

References

- Allingham, A. and A. Sandmo, 1972, Income tax evasion: A theoretical analysis, *Journal of Public Economics* 1, 323–338.
- Baron, D. and D. Besanko, 1984, Regulation, asymmetric information, and auditing, *The Rand Journal of Economics* 15, 447–470.
- Becker, G., 1968, Crime and punishment: An economic approach, *Journal of Political Economy* 76, 169–217.
- Border, K. and J. Sobel, 1987, Samurai accountant: A theory of auditing and plunder, *Review of Economic Studies* 54, 525–540.
- Crémer, H., M. Marchand and P. Pestieau, 1990, Evading, auditing and taxing: The equity–compliance tradeoff, *Journal of Public Economics* 43, 67–92.

- Christian, V., 1980, Two comments on tax evasion, *Journal of Public Economics* 13, 389–393.
- Demski, J. and D. Sappington, 1987, Hierarchical regulatory control, *The Rand Journal of Economics* 18, 369–383.
- Evans, J., 1980, Optimal contracts with costly conditional auditing, *Journal of Accounting Research* 18, 108–128.
- Graetz, M., J. Reinganum and L. Wilde, 1986, The tax compliance game: Toward an interactive theory of law enforcement, *Journal of Law Economics, and Organization* 2, 1–32.
- Greenberg, J., 1984, Avoiding tax avoidance: A (repeated) game-theoretic approach, *Journal of Economic Theory* 32, 1–13.
- Harris, C., 1986, On strategic taxation, Mimeo. (Oxford University).
- Kanodia, C., 1985, Stochastic monitoring and moral hazard, *Journal of Accounting Research* 23, 175–193.
- Kolm, S., 1973, A note on optimum tax evasion, *Journal of Public Economics* 2, 265–270.
- Laffont, J.J. and J. Tirole, 1986, Using cost observation to regulate firms, *Journal of Political Economy* 94, 614–641.
- Landsberger, M. and I. Meilijson, 1982, Incentive generating state dependent penalty system, *Journal of Public Economics* 19, 333–352.
- Maskin, E. and J. Riley, 1984, Monopoly with incomplete information, *The Rand Journal of Economics* 15, 171–196.
- Melumad, N. and D. Mookherjee, 1989, Delegation as commitment: The case of income-tax audits, *The Rand Journal of Economics* 20, 139–163.
- Mirrlees, J., 1971, An exploration in the theory of optimum income taxation, *Review of Economic Studies* 48, 175–208.
- Mirrlees, J., 1975, The theory of moral hazard and unobservable information, Part 1, Mimeo.
- Mookherjee, D. and I. Png, 1989, Optimal auditing, insurance, and redistribution, *Quarterly Journal of Economics* 104, 399–415.
- Morton, S., 1988, Strategic auditing for fraud, Mimeo. (Pennsylvania State University).
- Myerson, R., 1981, Optimal auction design, *Mathematics of Operations Research* 6, 58–73.
- Pestieau, P. and U. Possen, 1991, A simple model of income taxation with occupational choice and tax evasion, *Journal of Public Economics* 45, 107–125.
- Reinganum, J. and L. Wilde, 1985, Income tax compliance in a principal-agent framework, *Journal of Public Economics* 26, 1–18.
- Reinganum, J. and L. Wilde, 1986, Equilibrium verification and reporting policies in a model of tax compliance, *International Economic Review* 27, 739–760.
- Samuelson, W., 1984, Bargaining under asymmetric information, *Econometrica* 52, 995–1006.
- Sánchez, I., 1987, Three essays on the theory of hierarchical design, UCSD Ph.D. Dissertation.
- Sánchez, I. and J. Sobel, 1991, Hierarchical design and enforcement of income tax policies, UCSD Discussion paper 91-02.
- Scotchmer, S., 1986, Equity and tax enforcement, Harvard Discussion paper.
- Scotchmer, S., 1989, Who profits from taxpayer confusion?, *Economics Letters* 29, 49–55.
- Scotchmer, S., 1991, The regressive bias in tax enforcement, Sept.
- Scotchmer, S. and J. Slemrod, 1987, Optimal obfuscation in tax enforcement, Berkeley Discussion paper.
- Slemrod, J. and S. Yitzhaki, 1987, The optimal size of a tax collection agency, *The Scandinavian Journal of Economics* 89, 183–192.
- Srinivasan, T.N., 1973, Tax evasion: A model, *Journal of Public Economics* 2, 339–346.
- Stigler, G., 1970, The optimum enforcement of laws, *Journal of Political Economy* 78, 526–536.
- Tirole, J., 1986, Hierarchies and bureaucracies: On the role of collusion in organizations, *Journal of Law, Economics and Organization* 2, 181–214.
- Yitzhaki, S., 1974, A note on income tax evasion: A theoretical analysis, *Journal of Public Economics* 3, 201–202.