A Stochastic Game Model of Tax Evasion

T.E.S. Raghavan
Department of Mathematics
Statistics and Computer Science
University of Illinois at Chicago, USA
ter@uic.edu

Abstract

The problem of tax evasion is modelled as a zero-sum two-person generalized stochastic game with incomplete information. This model incorporates the classical statistical classification procedures used in classifying a random observation from a mixed population. The model incorporates the secrecy of the tax office and lack of information about the past history of the taxpayer. With full information, the model is closer to certain structured classes of stochastic games that admit efficient algorithms for optimal solutions.

1 Historical Introduction

Tax evasion as a topic for theoretical investigation was first suggested by J.A. Mirrlees in a paper prepared for the International Economic Associations Workshop in Economic Theory, in Bergen, Norway in 1971 (Allingham and Sandmo [2]). Independently Allingham and Sandmo [2],[1] and Srinivasan [45] considered static models which were almost identical but still different in terms of tax function and taxpayer's aim. The following is their model: Suppose the income y of a taxpayer when reported results in a tax $T(y)$. Let λ be a proportion by which y is understated. Since the government can find it only when the return is audited, the government does not know y, but only the reported $(1-\lambda)\,y$. Let π be the chance for being audited (of course π can depend on y). Let $P(\lambda)$ be the penalty multiplier, i.e., $P(\lambda)\lambda y$ is the penalty on the undeclared income λy. Let us assume that the individual chooses λ that minimizes his expected income (or suitable expected utility function of income). If the taxpayer is risk averse and if his expected payment on unreported income is less than what he/she has to pay otherwise, he/she will declare less when $\pi(y) = \pi$. When the probability of audit increases, the optimal proportion λ^* by which income is understated decreases. If the audit chances are independent of the income level, then, richer income is underreported at greater λ's. However, this will not be the case when π depends on income levels. When the actual income varies, the fraction declared increases or decreases according to the relative risk aversion in the sense of Arrow [4] is an increasing or decreasing function of income. Allingham and Sandmo [2] also considered a dynamic model for a constant tax rate with t periods. They assumed that by audit, the government could recover all the dues up to that day from the remote past. They showed that the optimal strategy for the taxpayer would be to choose a period T and evade until period T and report fully after period T. This line of research presupposes that *while the government is ignorant of a taxpayer, the taxpayer is fully informed of their audit chances!* Based on a survey conducted in Belgium, Frank and Dekeyser-Meulders [17] calculated certain tax discrepancy coefficients. They found that wage earners and salaried persons gained the least by tax evasion and that certain types of evasions could not be caught even after an audit. Balbir Singh [6] observed in Srinivasan's model that with fixed chance π for audit, if $\pi < 1/3$, taxpayers could even evade income tax completely. Kolm [21] pointed out that their models never involved any auditing costs. Based on a survey, Monk [24] suggested that greater resources should be allocated to auditing higher income groups. Spicer and Lundstedt [44] pointed out that tax evasion was more than just gambling. A psychological survey conducted by Spicer and Lundstedt [44] (also see Spicer [42]) revealed the following phenomena.

(1) Evasion is less likely when sanctions against evasion are perceived to be severe.
(2) Evasion is less likely when probability of detection is perceived to be high.
(3) Evasion is more likely when a taxpayer perceives that his terms of trade with the government are inequitable compared to others.

Vogel, in a survey conducted in Sweden [47], observed that taxpayers were vulnerable for tax evasion when their aspirations were not matched by the government's services. They also observed that direct cash flow resulted in greater tax evasion. Cross and Shaw [9] corroborated the same view on many professionals who were self-employed. Allingham, by a simple model, pointed out [1] that progressive taxation need not be a solution for removing inequities.

The models by Reinganum and Wilde [34], [35] and Erard and Feinstein [13] were clearly game theoretic and allowed strategic behavior by the Internal Revenue Service (IRS) against taxpayers. Reinganum and Wilde [34] through a simple model showed that an audit cut-off policy would be more desirable as it would dominate any random audit policy. Erard and Feinstein [13] expanded on the model of Reinganum and Wilde [34] and showed that unlike the model of Allingham and Sandmo [2] where honest taxpayers had no influence on the rest of the population indulging in tax evasion strategies, in their extended [13] model, in equilibrium, honest taxpayers had indirect peer pressure on tax evaders. Mookherjee and Png [25] develop a model and find sufficient conditions for random audits to be optimal.

There is a small amount of recent empirical work on what determines tax compliance (see [8], [11]). By partitioning the set of all taxpayers into three distinct classes, called 1. *honest*, 2. *susceptible*, and 3. *evading* types, Davis, Hecht and Perkins [10] study the problem via an explicit law of motion and its solution. For example, they assume that the rate of change of the population of honest taxpayers with respect to time is a negative proportion of the product of honest and evasive taxpayers. Justification for this assumption of their model is based on the empirical observation by Vogel [47] and Spicer and Lundstedt [44] that even honest people can become evaders in the future when their colleagues are evaders. One can contrast this with the assertions of Erard and Feinstein [13].

While many of these models are static, tax evasion and tax compliance are dynamic phenomena. One of the earliest dynamic game models of tax evasion was initiated by Greenberg [19]. See also Landsberger and Meiljison [22]. Greenberg, who formulated the problem as a repeated game with absorbing states, imposed some strong assumptions on the law of motion in order to achieve an elegant characterization of the optimal strategies. These were all zero-sum models.

The model proposed here is a generalized zero-sum stochastic game but with incomplete information. For the case when past history and immediate payoffs and transitions are common knowledge, these games reduce to tractable classes admitting efficient algorithms for computing good strategies (see Parthasarathy and Raghavan [28], Filar and Vrieze [15], Raghavan and Syed [30],[31]). The model is capable of incorporating empirical evidences via the immediate payoffs and transition probabilities.

Tax agencies like the IRS will show greater interest in the game theoretic approach only when the suggested solutions are further refinements that are closer to their current audit procedures developed in cooperation with their electronic data processing (EDP) units. Even popular books by IRS agents and supervisors (see Murphy [27], Monk [24] and informative articles by tax agency directors (Pond [29], Smith [39]) agree on the power and usefulness of discriminant analysis. Our models here complement and refine the discriminant function approach. We will still need the valuable and ingenious techniques of conducting sample surveys as in Frank and Dekeyser [17], Monk [24], Strumpel [46] to gather information about psychological behavior patterns of taxpayers in forming immediate payoffs. In this context the psychological studies in simulating income tax evasion by Friedland, Maital, and Rutenberg [18] and Spicer and Becker [43] will be very useful.

2 Secrecy and Lack of Information

An essential feature of taxation is the secrecy behind auditing procedures implemented by the tax office and the lack of full information about any taxpayer and his possible tax evasion strategies. These aspects have not been effectively incorporated into the models of tax evasion considered thus far in public economics literature. Often, in order to characterize equilibrium strategies and optimal strategies in closed form, model builders tend to make drastic assumptions. Our approach to modelling tax evasion is certainly not to look for closed form solutions but to look for models that retain the notions of *incomplete information* and *secrecy of actions* intrinsic to the tax evasion problem. At the same time, these models are quite close to existing stochastic game models where efficient solution techniques have already been developed. The notions of capturing secrecy and asymmetry in information have both been part of substantial research in the area of game theory known as *games with incomplete information.* Existence theorems are much harder to obtain in many such games with incomplete information. Here we propose a dynamic game theoretic approach to the study of the tax compliance problem that incorporates the dual secrecy inherent in the problems of tax evasion and auditing. The problem is viewed as a multistage game between the IRS (player I) and a taxpayer in a socioeconomic group (player II). The taxpayer adopts, either by choice or by ignorance of tax laws, a strategy to evade taxes on certain selected taxable items. Based on the particular socioeconomic group of the person, the IRS has a prior perception about the taxpayer with respect to his methods and modes. This perception is modified from year to year based on the tax returns and the dictates of the discriminant function and the norms of the IRS. This is modelled as a stochastic game with transition laws and states unknown to the taxpayer (player II) but known to the IRS (player I).

3 Detecting Tax Evasion via Discriminant Analysis

Fisher, in his seminal paper on taxonomic problems [16], suggested an ingenious procedure to classify any observation drawn randomly from a mixture of populations into one of them, based on the densities of the sub-populations. For many practical applications see [26], [3], [32], [33]. The procedure is easily adaptable to problems involving classifications in many other areas including bankers lending credit facilities for small businesses, taxation, and credit card approvals. Intuitively, we can describe this procedure for tax evasion problems as follows.

Although the tax paying population is quite heterogeneous, people in each professional group are relatively homogeneous. They tend to associate with people in the same professional group and inherit similar socioeconomic patterns of life. Thus, the population can be made more homogeneous by stratifying according to profession. Having stratified the population into sub-populations, such as executives, doctors, lawyers, salesmen etc., the next problem is to further divide each sub-population into two types, namely those filing legally correct and honest returns and those filing legally incorrect or manipulated tax returns. Apparently, in the early 1940s nearly 25% of the tax returns belonged to the second type [27]. While deductions in income tax returns accounted for less than 12% in 1947, a decade later the same deductions were almost 15% of the reported gross income [27]. Apparently, tax loopholes and manipulations were used in the process.

The classification of all members of a professionally homogeneous group into the above two distinct types is much more complex. This could only be achieved when persons in the profession were targeted earlier with a foolproof audit. As a first step one needs norms for auditing, so that tax items violating these norms conspicuously can be considered as candidates for auditing. Only expert tax inspectors can be relied upon to come to grips with this initial data classification problem.

Assume that there is data available from the past for this classification. Our hunch is that the IRS will know from past data the chance that a random tax return from a specified professional group is legally correct and honest. Needless to say, this chance will vary from profession to profession. It is known that many self-employed professionals and especially those who deal exclusively with cash transactions are often involved in tax evasions. Given all tax returns, the main statistical approach is to partition data into two disjoint sets where data in one set is classified as honest and requires no prima facie reason for auditing and the data from the complement is classified as incorrect or dishonest reports that need auditing. There are two costs associated with any such classification. If an honest return is audited, the cost of auditing time is wasted on the return. If a manipulated return is not audited, then the cost is the loss in taxes properly due. We have to convert all costs into money for proper comparisons. Now any tax return x is simply a vector whose coordinates correspond to tax items such as 1) married or single, 2) gross income, 3) dividend or interest income, 4) employee business expense, 5) real estate taxes paid, etc. Thus, in general $x = (x_1, x_2, \dots , x_p)$ is a vector of observations. Some are qualitative and some quantitative. For simplicity, we will assume all are quantitative. Then the past data collected from the two sub-populations adjusted for inflation will give mean values and variances and covariances for each sub-population. Let $f_1(x)$ and $f_2(x)$ be the densities that represent the populations. If the populations are normal, f_1 and f_2 are uniquely determined by the mean vectors μ_1 and μ_2 and variance covariance matrices Σ_1 and Σ_2. Thus, an optimal procedure is one that minimizes expected costs of misclassification. For example, if population 1 corresponds to the honest and correct tax returns, then c_{12} = cost of classifying 1 into 2 for auditing = audit costs. Now a tax audit procedure has to decide which observations x have to be audited. Let $(\mathbb{R}, \mathbb{R}^c)$ be a partition of all observations into don't audit, audit classifications. Then given the prior $\xi = (\xi_1, \xi_2)$ and $\mathbb{R}$ the expected cost is simply

$$c_{12}\xi_1 \int_{\mathbb{R}^c} f_1(x)\, dx + c_{21}\xi_2 \int_{\mathbb{R}} f_2(x)\, dx.$$

We could rewrite the same as

$$c_{12}\xi_1 + \int_{\mathbb{R}} (c_{21}\xi_2 f_2 - c_{12}\xi_1 f_1)\, dx.$$

Thus, when the integrand is less than zero on $\mathbb{R}$ the expected cost is minimized, which is the same as saying that the optimal classification procedure $\mathbb{R}^*$ satisfies

$$\mathbb{R}^* = \left\{ x : \frac{c_{21}\xi_2}{c_{12}\xi_1} \le \frac{f_1(x)}{f_2(x)} \right\}.$$

Equivalently, by taking logs, the procedure $\mathbb{R}^*$ reduces to

$$\mathbb{R}^* = \left\{ x : \log \frac{f_1(x)}{f_2(x)} \ge c \right\},$$

where c is known since $c_{21}, c_{12}, \xi_1, \xi_2$ are known.

If $\Sigma_1 = \Sigma_2$, $\log(f_1/f_2)$ to within some constant factor reduces to $\varrho(x) = (\mu_1 - \mu_2)^T \Sigma^{-1} x$. This is the famous *linear discriminant function* of Fisher ([3], [33]). The function $\varrho(x)$ is simply a linear combination of the x_i's for some suitable weights w_i's. We have the following intuitive interpretation of the discriminant function.

> *Each tax item i with reported x_i is given a weight w_i. The return is not audited if $\sum w_i x_i > c$, otherwise an audit is suggested.*

Of course, what is mathematically easily said is quite hard to implement. Even statistical problems with cost coefficients, prior distributions, etc., are quite difficult to compute exactly.

For example, when certain professions are hard hit by federal regulations, the changes in the pattern of expenditures may not come through immediately. Say that doctors and hospitals are being pressured to charge only a fixed amount for a certain diagnostic treatment, then clearly the income of the profession is much affected. The life style cannot be changed and the temptation to get away from tax payments increases. One needs to study such complex phenomena with suitable models. As the priors (ξ_1, ξ_2) will also change, we need to find suitable models to analyze them.

4 A Need for Further Game Theoretic Refinement of the Discriminant Function Approach

In the discriminant function approach, though the individual returns are classified into one of two sub-populations within each professional category, *the dynamics of tax evasion from tax year to tax year and the strategic audit manipulations to curb the evasions are not at all captured by such a purely statistical model. Straightforward discriminant analysis ignores the strategic manipulations of individual taxpayers, a key element in tax returns.*

As a further refinement of the statistical discriminant function approach we propose to formulate various game theoretic models of multistage games that conceptually capture the essence of tax games between the IRS and individual returns.

Before modelling in full generality, we will introduce the notion of a zero-sum two-person stochastic game with two states and two actions for each player.

A Stochastic Game with Two States and Two Actions Consider two players playing one of the games A or B. In both games players secretly select one of the numbers 1 or 2. Depending on their choices an immediate reward is received by player I from player II. Their choices and the current game they play determine which game will be played next time. The following is a simple example of such a game:

$$\begin{array}{cc} \text{A} & \text{B} \\ \begin{array}{cc} 5/A & 0/B \\ 0/B & 3/A \end{array} & \begin{array}{cc} 4/B & 0/A \\ 2/A & 1/B \end{array} \end{array}$$

Let player I secretly select one of the rows and II secretly select one of the columns. If in A row 1 and column 2 are chosen, player I receives nothing and the game moves to playing B next round. If in B row 2, column 1 are their choices the game moves to A after a reward of 2 to player I from player II. The payoff accrues and future payoffs are discounted at a fixed discount rate β. The aim of player I is to maximize the total discounted payoff. The aim of player II is to minimize the same. If x_n is the payoff on the nth day, $\sum_{n=0}^{\infty} \beta^n x_n$ is the total payoff where $0 < \beta < 1$.

Shapley [38] proved the remarkable theorem that these games can be intelligently played by locally randomizing the selection of rows in each matrix independant of the history of the play leading to the given game. For example if $\beta = .8$, the game value starting in A is approximately 6.79; in B it is approximately 5.43. A good strategy for I is to choose row 1 in matrix A with a chance 0.223 and to choose row 1 all the time in matrix B. Player II should choose column 1 in A with chance 0.223 and column 2 all the time in matrix B.

5 A Simple Model of a Tax Return-Audit Game

Consider the population of professional engineers employed by engineering firms. Suppose that from past auditing the IRS has a hunch that 10% of them manipulate returns, while 90% are honest. Given a tax return x, the IRS can compute the discriminant function which could decide whether to audit or not. *However, the IRS may have an initial perception on a return, which may cause the agency to audit, even though the discriminant function may indicate the opposite.* Namely, besides the two actions available to the IRS, the perception of the IRS is a variable which could vary from year to year depending on the years past. *This year's data may not reveal it. Last year's perception alone could give some clue.* Thus, the perception of the IRS can be thought of as states of the game which, for example, can also vary between the two states: honest and manipulating. Only an audit can make perceptional changes. Even if the discriminant function favors auditing, it cannot be immediately implemented for want of staff. One may have to manage with existing staff, which means limiting thee auditing facility. In such a case, strategic selection of auditing may be the only alternative. Thus, we can think of a tax return as a game with the following interpretation.

- **Players:** I - IRS, II - individual or firm filing tax return
- **Pure strategies:**
  For player I: 1. audit 2. don't audit
  For player II: 1. honest return 2. cheat
- **States:**
  A. IRS perceives a return as honest
  B. IRS perceives a return as manipulating.
- **Law of motion or transition probabilities:**
  If a return is not audited, then the perception of the IRS is the same as it was the previous year. If an audit finds someone guilty of manipulation, the perception changes from honest to manipulating. This is our stochastic game. The other situations are given below as in our mathematical example of a stochastic game. The perception of the IRS in states A and B is given below:

$$\text{state A} = (\text{honest})$$
$$\begin{array}{ll} & \quad\text{Honest} \quad \text{Manipulate} \\ \text{Audit} & \left[\begin{array}{ll} a_1/A & b_1/B \\ c_1/A & d_1/A \end{array}\right] \\ \text{Don't audit} & \end{array}$$

$$\text{State B} = (\text{manipulating})$$
$$\begin{array}{ll} & \quad\text{Honest} \quad \text{Manipulate} \\ \text{Audit} & \left[\begin{array}{ll} a_2/.5A & b_2/B \\ c_2/B & d_2/B \end{array}\right] \\ \text{Don't audit} & \end{array}$$

  For example, in state B, the IRS could perceive a taxpayer as being susceptible for manipulations even if the current audit finds no tax evasion on the items audited. As a measure of deterrence, the IRS continues to view any past tax violators with a 50:50 suspicion even after a current audit finds them honest.
- **Rewards:**
  In parlor games the immediate rewards are well defined simply by the rules of the game. In modelling real problems as games the most thorny issue is to define meaningful payoffs. In the case of tax returns, the actual tax collected with or without audit can be taken to be the immediate payoff corresponding to independent choices by the tax office and taxpayer. This immediate payoff can be defined as the expected tax paid when not audited and the expected tax collected with suitable fines imposed when audited finds someone guilty or not guilty less audit costs.

First one needs to estimate the prior perception probabilities. The IRS will have on k random persons, data $x_1, x_2, \ldots, x_k$ on tax item i and $y_1, y_2, \ldots, y_k$

on tax item j where i and j are independent deduction items and where the k persons were audited for the first time. Let $\xi_1, \xi_2, \ldots, \xi_k$ and $\eta_1, \eta_2, \ldots, \eta_k$ be the revised amount for the same two items after audit. Let $\bar{x}, \bar{y}, \bar{\xi}$, and $\bar{\eta}$ be the averages.

$$\frac{1}{k}\left[\left|i : x_i - \bar{x} > 0, y_i - \bar{y} > 0\right| - \left|i : \xi_i - \bar{\xi} < 0 \text{ or } \eta_i - \bar{\eta} < 0\right|\right]$$

is a rough overestimate of the proportion of people who would have manipulated.

The credibility is maintained until an audit proves otherwise. The threat of audit should always be on any person to discourage any future manipulation. This is incorporated in the first row first column entry in matrix B. The game is played as an ordinary stochastic game with discounted payoff. The above model, though completely in line with a model of an ordinary stochastic game, misses an important ingredient of our tax return problem.

Suppose a tax officer has two file cabinets to store all tax returns. Depending on the current perception of the tax officer that a taxpayer is honest or cheating, he stores honest ones in cabinet A and the rest in cabinet B. Thus the actual cabinet in which one's current tax return is saved will be known only to the tax officer. A taxpayer can assume that his file is in file cabinet A when he has never been audited. If a taxpayer was, after an audit, found cheating some time in the past, even if he is found honest by later audits, the taxpayer cannot be sure where his file will be stored by the officer. Thus, the taxpayer is often ignorant of the current state (*perception of the officer*) of the stochastic game. Similarly, if the taxpayer cheats, the tax office will not know this without auditing. Thus the tax office is in general not fully informed about the past actions of the taxpayer. Full information about the current state and past actions, namely the partial history of the game is not fully known to both players. Currently, all the standard existence theorems for zero-sum stochastic games assume full information about past history of actions for both players. See [41].

6 Generalized Stochastic Game

A population Π is partitioned into n sub-populations $\pi_1, \pi_2, \ldots, \pi_n$. Player II selects secretly a sub-population π_j and chooses a random observation x from π_j. Only the observation x is revealed to player I. Independent of the observation revealed, player I has a fixed prior distribution $\xi = (\xi_1, \xi_2, \ldots, \xi_n)$ on the sub-population selected by player II. Initially player I selects secretly a P_i according to ξ. Given the data x from π_j unknown to player I, based on the observation x revealed he computes a set $A_i = i(x)$, a finite set of actions available in P_i. Now he chooses secretly an action $a \in i(x)$ and receives from player II an immediate reward $r(x, i, a)$ and the game moves to P_k from P_i with chance $q(k/i, x, a)$. Player II secretly chooses a new j and a random x' from π'_j and the x' is revealed to player I. He computes possible actions $A_k = k(x')$ and selects an action $a' \in k\,(x')$. Again he receives a reward $r\,(x', k, a')$ and so on. The payoff accrues each time and the future payoffs are discounted at a fixed discount rate $\beta, 0 < \beta < 1$. The aim of player I is to maximize the expected discount reward. The aim of player II is to minimize the same.

It will be convenient to motivate our above generalized stochastic game both from the point of view of statistical decision theory (Wald [49], Blackwell and Girshick [7], Ferguson [14] and multistage game theory (Filar and Vrieze [15]). First we will set up the correspondence between our generalized stochastic game and tax return-audit game. This is shown in Table 1.

Let A be the maximal finite set of all possible actions for all possible tax returns. Suppose that the set A has ϱ elements $\{a_1, a_2, \dots, a_\varrho\}$. For each x one can associate a probability distribution $\left\{\phi_1^s(x), \phi_2^s(x), \dots, \phi_\varrho^s(x)\right\}$ on the exhaustive action space $A = \{a_1, a_2, \dots, a_\varrho\}$, where $\sum_{i=1}^{\varrho} \phi_i^s(x) = 1$. Here s is the current perception of the IRS. Thus, we can associate a stationary strategy $\left\{\phi_1^s, \phi_2^s, \dots, \phi_\varrho^s\right\}$ on the action space A for each x and current perception s. Let $\psi_1, \psi_2, \dots, \psi_N$ be a probability distribution on $\{\pi_1, \pi_2, \dots, \pi_N\}$. We are now ready to state some open problems.

Problem 6.1. Does the generalized β-discounted stochastic game admit a stationary optimal strategy for player I *assuming the following* conditions (1)–(4)?

(1) The partition $\pi = (\pi_1, \pi_2, \dots, \pi_N)$ is the same for both players.
(2) $q\,(k/x, i, a_i)$ is known to both players.
(3) The discriminant function $i(x)$ and the associated norm violation resulting in possible audit action set $A_i = i(x)$ is known to player II, for each data x.
(4) The perception s of player I about player II is also known to player II.

Problem 6.2. When the data x comes from continuous densities corresponding to $\pi_1, \pi_2, \dots, \pi_n$ and when the conditions of Problem 6.1 are satisfied, can one replace the stationary strategies $\{\phi_1^s(x), \phi_2^s(x), \dots, \phi_r^s(x)\}$ by a pure strategy? That is, given data x do we have a single action for each perception P_s which is equivalent to $\{\phi_1^s(x), \phi_2^s(x), \dots, \phi_r^s(x)\}$ in the sense of equivalent rewards?

In this context we want to recall the theorem of Dvoretsky, Wald, and Wolfowitz [12] in statistical decision theory.

Theorem 6.1. *Let Ω be a finite set of parameters representing states of nature. Let A be a finite subset of $\mathbb{R}^n$ representing actions of a statistician. Let f_ϖ for each ϖ be a continuous density function. Let D be the space of decisions where each $d \in D$ is a map $d : X \to A$ where X is the sample space. Let $L\,(\varpi, a)$ be a bounded measurable loss function. Then any randomized decision $\phi : x \to \{\phi_1(x), \phi_2(x), \dots, \phi_k(x)\}$, where $\phi_i(x) =$ the chance action i is taken*

Table 1: Correspondence between the two games.

Player I	IRS
Player II	Individual taxpayer
sub-populations $\pi_1, \pi_2, \ldots, \pi_n$	Those within the professional group who manipulate a specific set of tax items in the tax return.
$P_1, P_2, \ldots, P_n$	In the eyes of IRS possible sets of items that are being manipulated by various types of persons in the profession. (Partition according to perception based on past data.)
A_i	For perception P_i, the set of actions available to IRS. (For example, if IRS suspects on, say, items 1) moving expenses and 2) charitable contributions. They may choose to audit on item 1, 2, both or none. These are 4 possible actions.)
$x = (x_1, x_2, \ldots, x_p)$	A vector of items filled in the tax return with some of the x_i's as qualitative variables such as marital status, filing status, etc.
$i(x)$	With each tax data vector x, a set A_i of audit actions that are needed to correct the *norm violations, as found in the data* via discriminant analysis. Whereas the perception of the IRS is based on *past tax returns* by the taxpayer, the actions at the current time are dictated by the current discriminant function. Thus the importance of the discriminant function lies in not just classifying the observation, but also suggesting possible audit actions, based on the *current data*.
$q\,(k/x, i, a)$	The chance that the tax return x, with an initial perception P_i and an action $a \in A_i = i(x)$ by IRS, changes the perception of the IRS from P_i to the new perception P_k.
$r\,(x, i, a)$	The actual tax collected when the return reports data x, when the perception by IRS is P_i and when action a is taken by the IRS.
β	Discount factor accounting for inflation rate, etc.

when x is observed, can be equivalently replaced by a pure decision $d : x \to A$ in the sense that the expected risk $r(\varpi, d) = r(\varpi, \phi)$ for all $\varpi \in \Omega$.

Our Problem 6.2 is to extend this theorem in the context of our generalized stochastic games.

7 Stochastic Games with Incomplete Information

From the point of view of our actual tax return-audit problem we need to handle the more difficult problem of lack of information from either side.[1] For example, the perception P_i of the IRS is rarely known to player II, the taxpayer. Also the law of motion $q(k/s, x, a_i)$ is unknown to the taxpayer. Actually, the theory of stochastic games with incomplete information has few computable solutions. The theory of structured stochastic games has many existence theorems and efficient algorithms to compute value and optimal or equilibrium strategies ([28],[48],[15],[30],[31]), and for structured repeated games (a very special class of stochastic games) with incomplete information of a special type, one has some existence theorems. However, there are very few computational tools. See [23], [40], [41], [36], and [37] (this volume) in recent years. The researches in the area of stochastic games with incomplete information that are close to our model are the ones by Melolidakis [23] and Rosenberg, Solon, and Vieille [36]. We could call our tax return-audit problem a statistical extension of stochastic games with incomplete information. In the next section, we will briefly discuss the notation of stochastic games with lack of information and show what our generalized stochastic game is with reference to this setup.

8 Games with Lack of Information on One Side

Games with incomplete information were pioneered by Harsanyi [20], and later formulated in some precise mathematical models for certain special kinds of information lags by Aumann and Maschler [5]. For more recent developments on stochastic games with incomplete information, see Sorin [41]. For our tax model what we will need is a certain subclass of games called stochastic games with lack of information on one side (SGLIOS) in the sense of [23], an adaptation of the Aumann–Maschler model for discounted and undiscounted stochastic games.

SGLIOS Model: A stochastic game with lack of information on one side consists of:

(i) A set of $m \times n$ matrices $S = \left\{A^1, A^2, \ldots, A^N\right\}$ called the “states” of the game. We identify A^s with state s.

---

[1] This critical aspect of the problem was first pointed out to the author by Professor Ritzburger, of the Institute of Advanced Study, Vienna, Austria.

(ii) A prior distribution $\xi^0 = \left(\xi_1^0, \ldots, \xi_N^0\right)$ on S.
(iii) A law of motion $q\,(t/s, i, j)$ where the game moves to state t from state s, when row i and column j are chosen secretly by the players in state s resulting in an immediate payoff $(A^s)_{ij} = a_{ij}^{(s)}$.
(iv) Player I alone knows the true state. Each time the choices i, j are revealed to both players after the choices are made.
(v) The immediate payoff $a_{ij}^{(s)}$ is kept secret from player II, though player I knows the same.
(vi) The payoff is evaluated by discounting each time with a discount factor β $(0 < \beta < 1)$.

The following is the main theorem.

Theorem 8.1 (Melolidakis)**.** *Let* Γ *be a* β*-discounted stochastic game of the above type SGLIOS.* We can associate an ordinary stochastic game, Γ^* *where player I has a pure optimal stationary strategy* $f^*(\xi)$ *and player II has a stationary optimal strategy* $g^*(\xi)$*. Here the game* Γ^* *is played as follows. Let player I, as in SGLIOS, use his usual information in selecting his behavioral strategy. Unlike in* Γ, *here player II is informed of the posterior distribution at each stage based on the state of the game, the actions of the player, and the law of motion. One of the main observations of Melolidakis is that the value* $v(\Gamma) = v(\Gamma^*)$, *for player I loses nothing by revealing the posterior.*

As player I knows all about the law of motion, prior, state of the game, etc., one can consider the following stochastic game.
Game Γ^{**}*:* Let the action space of I be $\{f(s) : s \in S\}$ where $f(s)$ is a mixed strategy on the rows of A^s. Let the action space of II be the set $\{1, 2, \ldots, N\}$. Let the state space be all probability vectors in $\mathbb{R}^N$. Let the law of motion Q be $q\,(./\xi, f, j)$ where the new prior at t is the posterior given f, j, To clearly understand the new prior, as the posterior given the actions of the players in the original game, we evaluate the posterior probability of the game to be in state s, given the actions i, j of the two players. This is $\eta\,(s/f, j, \xi)$ where for i fixed the entry is

$$\begin{aligned}\eta\,(t/i, j, \xi) &= \sum_s q\,(t/s, i, j)\, f_i(s)\xi(s) / \sum_s f_i(s)\xi(s).\\ r\,(\xi, f, j) &= \sum_s \sum_i \xi(s) f_i(s) r\,(s, i, j)\,.\end{aligned}$$

Here $(A^s)_{ij} = r\,(s, i, j)$. This is the stochastic game induced by SGLIOS. An important observation is that the value of this stochastic game coincides with the value of the original stochastic game.

9 Similarities and Differences between the SGLIOS and Our Tax Return-Audit Game

In our tax return game player I (the IRS) knows about the actual state (the perception). The prior is also known to the IRS (based on past data). Instead of a finite set of actions for player II, the selection involves a partition and the selection of a random observation from a partitioned sub-population. While the actual actions are revealed in SGLIOS, here only the data x and action i, corresponding to the audit decision by the IRS, are known to each other. Whereas the actual payoff is unknown to player I but the law of motion $q\,(t/s, i, j)$ is known in SGLIOS, here the actual payoff is known but the law of motion $q\,(t/s, x, j)$ corresponding to data x, action j and change of perception to t from s is kept secret by the IRS. Intuitively the change of perception of the IRS about a taxpayer is kept secret even though the audit resulting in tax payments is known to both sides.

10 Some Thoughts on Norms for Audits and Some Questions on Revealing Audit Policies

In a penetrating paper on deriving norms for income tax audits, Pond [29] makes the following remarks: "There are always some, who through inadvertence or design, minimize their tax liability. Deductions offer one of the greatest avenues for minimizing tax liability and it is evident that to be most successful with the available staff, the audit program should concentrate on taxpayers whose deductions are excessive in relation to others in a comparable income classification. The first step in deriving the norms is to determine a frequency distribution for each deduction "and" calculate the ratio of net income to gross income! The carefully chosen tax returns for closer audit saves audit time without losing tax arrears. Such carefully chosen returns represent 5/6 of the total on all the cases. Since the audit agency was handling only 2/3 of all cases on a nonselected basis and therefore only was producing 4/6 of the potential, the application of norm method has an imputed gain of 25%."

In a sense many statistically heuristic procedures are already perhaps adopted by the IRS! As part of the deduction of Pond's paper the first question raised was whether it was possible for a taxpayer to become familiar with the selection criteria and thus become able to evade taxes and be sure of escaping detection. Two factors were seen mitigating against this: 1) The norms are kept secret, and 2) they are constantly being reevaluated. There was no general agreement on how the norms were to be evaluated. Fault was found with putting emphasis on assessment/cost ratio; the failure to audit in such cases might reduce voluntary compliance within those groups.

In our opinion, the verbal language above and the discussions pertaining to the problem of tax are in the spirit of multistage games. Thus a proper analysis of the problem via models of stochastic games is only desirable from both the theoretical and practical points of view. Perhaps solutions to such models might reveal answers to unanswered questions like the following.

Problem 10.1. Can the norms for auditing tax returns be made public?

An answer to Problem 10.1, though volatile, might still be valuable to look into for suitable models. With reference to our formulated model, this is the same as the following mathematical interpretation.

Problem 10.2. In our game with incomplete information if the state (IRS's perception) alone is kept secret, but not the actual law of motion of the game and the discriminant function and the partition $P_1, P_2, \ldots, P_M$ by the IRS, will the value of the stochastic game change?

We could conceptually understand Problem 10.1 by modelling the game as the following single controller stochastic game with incomplete information.

11 A Single Controller Game with Incomplete Information

Players I and II know that a population Π is a mixture of sub-populations $\Pi_1, \Pi_2, \ldots, \Pi_n$ with densities $f_1(x), f_2(x), \ldots, f_n(x)$, respectively. Player II chooses secretly a $j \in \{1, 2, \ldots, n\}$ and then selects a random observation x from Π_j. The choice j is not revealed to player I. However, the randomly selected observation x from density f_j is revealed to player I. Player I has prior $\xi_1, \xi_2, \ldots, \xi_n$ on his perceptions about the current choice of player II. We call player I's current perception the state of the game. The perception of player I remains unchanged and stays at state s with probability $1-\phi(s)$, unknown to player II (here $0 \leq \phi(s) \leq 1$). With probability $\phi(s)$ the perception of player I changes to a new state k taking into account the posterior dictated by the data x. Based on x, he selects an action $i \in \{1, 2, \ldots, n\}$ with probability $\psi_i(x)$. In the case $i \neq j$, player I receives a reward $c_i(s)$ from player II. In the case $i = j$ he receives an amount $u_j(x)$ from player II. The play continues with the posterior as the new prior. The payoff accrues at a fixed discount rate β. The aim of player I is to maximize the total discounted reward. The aim of player II is to minimize the same.

We could convert the above model into the following single controller stochastic game with incomplete information. We need to define immediate rewards and transition probabilities.

Let $q\,(k/s, \psi, j)$ = expected transition probability of the perception of player I to move from perception s to perception k given the strategies ψ and j by players I and II, respectively. Since the transition depends only on the posterior and the preassigned norms for remaining in status quo or following the decision based on the data, we get

$$q\left(k/s,\psi,j\right)=\begin{cases}\phi(s)\cdot\int\left[\frac{\xi_k f_k(x)}{\sum_t \xi_t f_t(x)}\right]\cdot f_j(x)\,dx, & k\neq s\\ [1-\phi(s)]+\phi(s)\cdot\int\left[\frac{\xi_s f_s(x)}{\sum_t \xi_t f_t(x)}\right]\cdot f_j(x)\,dx, & k=s.\end{cases}$$

Notice that $\xi_k f_k(x)/[\sum_t \xi_t f_t(x)]$ is the posterior probability given the data and q gives the expected transition probability.

An important observation is that this transition probability $q\left(k/s,\psi,j\right)$ depends only on the action of player II. Since $\phi(s)$ is unknown to player II, the actual law of motion is unknown to player II, although he controls the law of motion! The expected immediate reward $r\left(s,i,j\right)$ to player I can be written as

$$r\left(s,\psi,j\right)=\int_\Omega\left(\sum_{i\neq j}\psi_i(x)c_i(s)\right)f_j(x)dx+\int_\Omega\psi_j(x)u_j(x)f_j(x)dx.$$

Thus, our problem reduces to an ordinary one player control game if the perceptions of player I, u_j, c_i, ξ_i, etc., are common knowledge. Even if c_i's and ϕ's are known to player II, as long as the actual perception of player I about player II is kept secret, the game will still be a single controller stochastic game with lack of information on the law of motion for player II.

Thus we are led to the following problem.

Problem 11.1. Let Γ be a single controller stochastic game with reward $r\left(s,\psi,j\right)$, transition probabilities $q\left(k/s,j\right)$, and discount factor β, $0<\beta<1$. A prior distribution ξ on the states is chosen. Though the prior is known to both players, the actual state is known only to player I.

The law of motion $q\left(k/s,j\right)$ will be known to player I if action j of II is known. Even if $\phi(s)$ is revealed to player II, only the law of motion will be known to player II, but the true state s of the game will still be unknown to player II. The reward $r\left(s,\psi,j\right)$ is unknown to player I as he does not know j, the choice of player II. If $c_i(s)$ is independent of s, the reward in each state is $r\left(\psi,j\right)$, which depends only on the actions of players I and II. Even in this case the immediate payoff is unknown to either player as each one's choice remains secret in each round. A final case is when player II chooses a fixed j once and for all and all he does from one round to the next is choose an independent observation from the same density f_j. This is the closest to the single control games considered by Rosenberg, Solon, and Vieille [36]. In this case we are led to the one player control game with known reward, but unknown law of motion for the controlling player. The main problem is to find whether such games admit value and, given the data, to solve for the value and good strategies if any.

This game captures the spirit of tax return-audit in the following sense. The perceptions about a taxpayer by the IRS as honest, moderate, cheats on moving expenses, etc., have to be solely based on the data x the taxpayer submits and the action he chooses to get x. Thus, his own actions essentially contribute towards any changes in the perception of the IRS. The threat to keep him obedient to tax laws needs the secrecy of the perception (state). The norms governing the status of an audited taxpayer are captured by the function $\phi(s)$. We will briefly summarize the research findings of some earlier models that are found in the literature on public finance.

Generalized stochastic games: We have already described these games earlier. To focus just on the mathematical formulation of these games, we need only to modify our preceding formulation on stochastic games.

Generalized stochastic game: Players I and II play the following game. The game has a finite number of states $1, 2, \ldots, S$. In each state player I has m actions and player II has n actions. Player II selects an action j in state s. This is revealed to a referee. The referee picks a random observation x according to the density function $f_j(x)$ and reveals x to player I. Not knowing j, but knowing x, player I selects an action i among his m actions. Then he receives an amount $r(s, i, x)$, and the game moves to a state k with chance $q(k/s, i, j)$, and so on. The payoff as before is the total discounted payoff. Here a stationary strategy for player II is the same as before. However, a stationary strategy for player I is of the type $\phi_i^s(x)$ where $\phi_i^s(x) =$ the chance action i is selected in state s when observation x is given. It is a pure stationary strategy if $\phi_i^s(x) = 0$ or 1 for each x, s. The law of motion is common knowledge and the reward is known to both players.

Our problem is to check whether the game has optimal stationary strategies: Also, we are interested in the situation where the optimal stationary strategies are replaceable by pure stationary optimals. We have already discussed the special subclasses of such single controller games with incomplete information as a model of our tax problem.

12 A Model of Tax Evasion as a Stochastic Game with Incomplete Information on the States

Consider a population of taxpayers all belonging to a single professional category. A tax return is simply a p-vector $X = (X_1, X_2, \ldots, X_p)$ with X_1 as the adjusted gross income and X_p as the tax due, *as reported by the taxpayers* in their tax returns. Based on their past tax returns and past audit actions by the IRS, the current tax returns are stratified and stored in S distinct file cabinets $1, 2, \ldots, S$. In the *perception of the IRS*, based on most recent audits, the returns stored in file cabinet $j > i$ are viewed as higher-order tax violations than those in file cabinet i. We will assume that, by an audit, the tax office can always find out the true values $(Y_1, Y_2, \ldots, Y_p)$ for a taxpayer's tax return. If a taxpayer resorted to, say, the kth level of tax violation, then his reported tax return will be taken to be a vector function of the true values defined by

$$X_t = \phi_t^k(Y_1, Y_2, \dots, Y_p), \quad t = 1 \dots, p.$$

Let $f^1(x), f^2(x), \dots, f^S(x)$ be the joint density functions corresponding to the reported tax dues of the population of tax returns in file cabinets $1, \dots, S$. We will often use the random variables (X_1, X_p) with marginal joint densities $g^j(x_1, x_p)$, $j = 1, \dots, S$. In general the tax office maintains secrecy of the locations of individual files, and information about the marginal densities and joint densities of the returns stored in various file cabinets. When the IRS decides to audit a tax return from file cabinet k, the taxpayer will be notified about the current location k from where the return was chosen for audit.

Since the number of auditors is fixed, the tax office has to allocate the available auditing time A efficiently. An intuitive policy would be to rearrange the files in each file cabinet i from the smallest to the largest values of R where

$$R = \frac{\text{Adjusted Gross income}}{\text{Tax due}} = \frac{X_1}{X_p}$$

and target the upper end among them, namely those for whom $R > \rho$ for some ρ chosen secretly by the tax office. Obviously, there should be many deductions of various kinds to arrive at a relatively small tax, and the auditing hours will be longer on such tax returns. Just because R is large one cannot immediately conclude that the person is a cheater. The deductions could be genuine and the person could be honest. It could have been a bad year for the taxpayer with large hospital bills beyond insurance coverage. However, this ratio R is more likely to exceed the given value ρ in a population of higher-order tax violators than in a population of lower-order tax violators and in particular for honest taxpayers. Therefore, for any random tax return X^j from file cabinet j, let R^j denote the above-mentioned R value. Then for any random tax return X^i from file cabinet $i, P(R^j > \rho) \geq P(R^i > \rho)$ if $j > i$.

A strategy for the tax office is to choose a threshold value ρ and to target for audit all tax returns with $R > \rho$. Thus by the stochastic ordering assumption, a greater proportion of tax returns from file cabinet j will be targeted than from file cabinet $i < j$. For simplicity let us suppose the audit time to audit a tax return with value R is cR. Let $h_i(r)$ be the density of the reported value of R for the tax returns in the file cabinet i. Thus the expected audit time for file cabinet i is given by

$$c\int_\rho^\infty r h_i(r)dr = q_i.$$

This will also immediately fix the total audit time for all file cabinets as

$$\sum_i c\int_\rho^\infty r h_i(r)dr = \sum_i q_i.$$

Let $u^i(x_p)$ denote the density of the random variable X_p representing the tax amount on any random tax return from file cabinet i.

If by an audit the IRS comes to know that a taxpayer has indulged in a tax violation of order $k > i$, then the IRS classifies the current and future returns of the taxpayer in file cabinet k. If the audit reveals that the tax violation is of order $k < i$, the IRS classifies the current return and future returns of this taxpayer in file cabinet k. The following is the intuition for such a transition. If a person's return, either based on prior allocation or on recent audit was found to be a tax violation of a certain order, when audited currently is found to be one of a higher order, he/she deserves to be watched with immediate reclassification with necessary caution. The persons who are found from current audit to be improved with lower levels of violation are recognized for their acceptance of law and order with a slight bit of reservation. When a tax return is not audited the IRS loses tax on a taxpayer when he/she becomes a tax evader of a higher order. However when a tax return from a taxpayer who has considerably toned down from his/her original level of tax evasion is audited, the tax office incurs higher cost due to unnecessarily prolonged auditing. The transition probability based on the preceding intuitive principles can be defined as follows.

Let the transition probability be $q(j/s, k, \rho)$ where s is the current location (file cabinet) of the return from where the return with data X was picked for audit using ρ strategy and found to be of violation level k. In case $k \geq s$ the file is immediately transferred to file cabinet k with probability one. Suppose that the audit reveals a violation level $k < s$; then with a small probability α it is kept in the same file cabinet and with probability $1 - \alpha$ it is moved to file cabinet k. Thus if a tax return X from file cabinet s is audited and if the taxpayer has chosen a tax violation level k currently, then the tax return moves to state j with transition probabilities given by

$$\begin{aligned} q(j/s, k, \rho) &= 1 \quad \text{if} \quad j = k \text{ and } k \geq s \\ &= \alpha \quad \text{if} \quad j = s \text{ and } k < s \\ &= 1 - \alpha \quad \text{if} \quad j = k \text{ and } k < s. \end{aligned}$$

Suppose that the audit strategy ρ is chosen by the IRS. If a taxpayer has never been audited, then he can assume that his tax return is located in file cabinets $1, 2, \ldots, S$ with respective priors ξ_1, ξ_2, ξ_S. The priors are known to all taxpayers. If a taxpayer was audited in the past, based on the most recent audit he can evaluate the posterior probabilities for the current location of his tax return.

The revenue for the IRS from a taxpayer will depend on the following:

- Was he ever audited and if so what was the violation level of the most recent audit?
- Is he currently being audited?
- What is the current level of violation of the taxpayer?

Suppose that the taxpayer chooses currently a level $j = j(X)$ for tax violation. Including the current year suppose he has never been audited. Since his return could be from file cabinet i with stationary prior probability ξ_i it would have escaped the current audit if the calculated R value $X_1/X_p < \rho$. When it is not audited, he pays only X_p. Since he has chosen level j, the tax return can be thought of as a random observation from file cabinet j with density $f^j(x)$ and with $R < \rho$. Thus the conditional expected payoff to the tax office given that the tax return was in file cabinet i, and the current choice was j by the taxpayer, and it escaped audit currently is given by

$$\int\int_{\{(x_1,x_p):x_1<x_p\rho\}} x_p g^j(x_1, x_p)dx_1 dx_p.$$

Thus the expected income to the IRS from such a never-audited tax return is given by

$$\sum_j \xi_j \int\int_{\{(x_1,x_p):x_1<x_p\rho\}} x_p g^j(x_1, x_p)dx_1 dx_p.$$

Suppose that audit costs are w dollars per hour. The IRS charges a suitable penalty for tax violations depending on the level of tax violation when audited. Let each dollar due be multiplied by a penalty factor θ_k for tax returns audited from file cabinet i found to be a tax violation of level k. If the IRS charges a penalty proportional to the difference between the true tax due and reported tax amount X_p specified in the tax return, then the net expected income to the IRS from an audited tax return from file cabinet i with violation level $k > i$ is $r(i,k) = \theta_k(\mu_1 - \mu_k) + \mu_k - cw\int_\rho^\infty rg_i(r)dr$. Here $g_i(r)$ is the density of the statistic R from file cabinet i and μ_i = expected tax from file cabinet i for all files that escaped audit. Also $\theta_k > \theta_{k-1} \cdots > \theta_1 = 1$. In the case $k < i$, and if the taxpayer is audited then the tax office finds that the taxpayer is relatively reformed and the expected income to the tax office is $\theta_i(\mu_1 - \mu_i) + \mu_i - cw\int_\rho^\infty rg_i(r)dr$. The expected income to the tax office with the ρ strategy when the taxpayer in file cabinet i wants to choose tax violation level k is given by

$$P_i(R < \rho)\theta_i(\mu_1 - \mu_i) + \mu_i - cw\int_\rho^\infty rg_i(r)dr$$
$$+ P_i(R > \rho)\theta_k(\mu_1 - \mu_k) + \mu_k - cw\int_\rho^\infty rg_i(r)dr.$$

If the tax form for a taxpayer has p items to fill in with numerical values, any subset S among those items can be misrepresented by the taxpayer by deviating from the true value. Suppose that the tax office can identify the deviated items by audit; then the set of such deviators will constitute a sub-population with a density function f_S. In our model we assume that their tax returns are to be stored in a file cabinet labelled S. When a tax office calls a taxpayer for audit, and spells out where they have doubts on the tax return, they essentially reveal the label of the file cabinet from which this return is chosen for audit on the labelled items. Given the information that a taxpayer's file was stored in file cabinet S, after the most recent audit, a simple class of pure strategies for the taxpayer who wants to act like a random person from file cabinet A can be generated by any scale vector $a = (a_1, a_2, \ldots, a_p)$ that is used to fudge the true data X and report it as $Y = (Y_1, \ldots, Y_p)$ where $Y_i = a_i X_i, i = 1, \ldots, p$ and $A = \{i : a_i \neq 1\}$. Similarly, a simple pure strategy for the tax office is a choice of ρ that selects in the first round all tax data whose R value exceed ρ. If there are too many selected this way, a suitable stratified random sampling scheme can be used to select the size that is manageable with existing audit resources.

REFERENCES

[1] Allingham, M., Inequality and progressive taxation, *J. Public Economics,* 11, 273–274 (1979).

[2] Allingham, M. and A. Sandmo, Income tax evasion: A theoretical analysis, *J. Public Economics,* 1, 323–338 (1972).

[3] Anderson, T.W., *An Introduction to Multivariate Statistical Analysis*, Wiley, New York (1972).

[4] Arrow, K.J., *Essays in the Theory of Risk Bearing*, Markham, Chicago (1971).

[5] Aumann, R.J. and M. Maschler, *Repeated Games with Incomplete Information*, M.I.T Press, Cambridge, MA (1995).

[6] Balbir, S., Making honesty the best policy, *J. Public Economics,* 2, 257–263 (1973).

[7] Blackwell, D. and G.A. Girschick, *Theory of Games and Statistical Decisions*, Wiley, New York (1955).

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

[9] Cross, R.B. and G.K. Shaw, The evasion-avoidance choice: A suggested approach, *Nat'l Tax Journal,* 34, 487–491 (1981).

[10] Davis, J.S., Hecht, G., and J.D. Perkins, Social behaviors, enforcement and tax compliance, *The Accounting Review,* 78, 39–69 (2003).

[11] Dubin, J.A., Graetz, M.J., and L.L. Wilde, Are we a nation of tax cheaters? New econometric evidence on tax compliance, *The American Economic Review,* 77, 240–245 (1987).

[12] Dvoretzky, A., Wald, A., and J. Wolfowitz, Elimination of randomization in certain statistical decision procedures and zero-sum two person games, *Ann. Math. Stat.,* 22, 1–21 (1951).

[13] Erard, B. and J. S. Feinstein, Honesty and evasion in the tax compliance game, *Rand J. Economics,* 25, 1–19 (1994).

[14] Ferguson, T.S., *Mathematical Statistics: A Decision Theoretic Approach*, Academic Press, New York (1967).

[15] Filar, J.A. and O.J. Vrieze, *Competitive Markov Decision Processes*, Springer-Verlag, Berlin (1996).

[16] Fisher, R.A., The use of multiple measurements in taxonomic problems, *Ann. Eugenics,* 7, 179–188 (1936).

[17] Frank, M. and M. Dekeyser, A tax discrepancy coefficient resulting from tax evasion or tax expenditures, *J. Public Economics,* 8, 67–78 (1977).

[18] Friedland, N., Maital, S., and A. Rutenberg, A simulation study of income tax evasion, *J. Public Economics,* 10, 107–116 (1978).

[19] Greenberg, J., Avoiding tax avoidance: A (repeated) game theoretic approach, *J. Economic Theory*, 32, 1–13 (1984).

[20] Harsanyi, J., Games with incomplete information played by Bayesian players, I–III, *Management Science,* 14, 159–182; 320–334; 486–502 (1967–68).

[21] Kolm, S.C., A note on tax evasion, *J. Public Economics,* 2, 265–270 (1973).

[22] Landsberger, M. and I. Meiljison, Incentive generating state dependent penalty system: The case on income tax evasion, *J. Public Economics,* 19, 333–352 (1982).

[23] Melolidakis, C.A., Stochastic games with lack of information on one side and positive stop probabilities, pp. 113–126, in: *Stochastic Games and Related Topics*, Raghavan, T.E.S., Ferguson, T.S., Parthasarathy, T., and O.J. Vrieze (eds.), Kluwer Publishers, Dordrecht, The Netherlands (1991).

[24] Monk, K.A., Income tax evasion: Some empirical evidence, *J. Public Finance*, 30, 70–76 (1975).

[25] Mookherjee, D. and I. Png, Optimal auditing, insurance and redistribution, *The Quarterly J. Economics,* 104, 399–415 (1989).

[26] Morrison, D.F., *Mutivariate Statistical Methods*, Fourth Edition, Duxbury Press, Pacific Grove, CA (2002).

[27] Murphy, J.H., EDP and tax administration in New York, *Nat'l Tax Journal*, 14, 223–227 (1961).

[28] Parthasarathy, T. and T.E.S. Raghavan, An order field property for stochastic games when one player controls transition probabilities, *J. Optimization Theory and Appl.*, 33, 375–392 (1981).

[29] Pond, C.B., Derivation and application of norms in selecting audit, *Nat'l Tax Journal,* 14, 227–231 (1984).

[30] Raghavan, T.E.S. and Z. Syed, A policy-improvement type algorithm for solving zero-sum two-person stochastic games of perfect information, *Math. Program.,* 95, 513–532 (2003).

[31] Raghavan, T.E.S. and Z. Syed, Computing stationary Nash equilibria of undiscounted single-controller stochastic games, *Math. Oper. Res.,* 27, 384–400 (2002).

[32] Rao, C.R., *Advanced Statistical Methods in Biometric Research*, Wiley, New York (1952).

[33] Rao, C.R., Discrimination among groups and assigning new individuals: The role of methodology of classification in psychiatry and psychopathology, 229–240, U.S. Dept. of Health Education and Welfare (1966).

[34] Reinganum, J. and L. Wilde, Income tax compliance in a principal agent framework, *J. Public Economics,* 26, 1–18 (1985).

[35] Reinganum, J. and L. Wilde, Equilibrium verification and reporting policies in a model of tax compliance, *International Economic Review,* 27, 739–760 (1986).

[36] Rosenberg, D., Solon, E., and N. Vieille, Stochastic games with a single controller and incomplete information, *SIAM Journal on Control and Optimization,* 43, 86–110 (2004).

[37] Rosenberg, D., Solon, E., and N. Vieille, Stochastic games with imperfect monitoring (this volume).

[38] Shapley, L.S., Stochastic games, *Proc. Nat'l Acad. Sci. U.S.A.,* 39, 1095–1100 (1953).

[39] Smith, W.H., Developing a new technique in selecting returns for audit, *J. Accountancy,* 123, 22–23 (1967).

[40] Sorin, S., Big Match with lack of information on one side, *Internat. J. Game Theory,* 14, 173–204 (1985).

[41] Sorin, S., *A First Course on Zero-Sum Repeated Games*, Springer-Verlag, New York (2002).

[42] Spicer, M.W., A behavioral model of tax evasion, Ph.D. Thesis, Ohio State University, Columbus, OH (1974).

[43] Spicer, M.W. and L.A. Becker, Fiscal inequity and tax evasion: An experimental approach, *Nat'l Tax Journal*, 33, 171–175 (1980).

[44] Spicer, M.W. and S.B. Lundstedt, *Public Finance*, 31, 295–305 (1976).

[45] Srinivasan, T.N., Tax evasion: A model, *J. Public Economics,* 2, 339–346 (1973).

[46] Strumpel, B., Contributions of survey research, in: A. Peacock (ed.), *Quantitative Analysis in Public Finance*, 29–32, Praeger Publ., New York (1969).

[47] Vogel, J., Taxation and public opinion in Sweden: An interpretation of recent data, *Nat'l Tax Journal,* 27, 499–514 (1974).

[48] Vrieze, O.J., Stochastic games with finite state and action spaces, Ph.D. Thesis, Free University of Amsterdam, The Netherlands (1983).

[49] Wald, A., *Statistical Decision Functions*, Wiley, New York (1950).