Tax Evasion as an Optimal Solution to a Partially Observable Markov Decision Process

Paraskevi Papadopoulou and Dimitrios Hristu-Varsakelis

Abstract Motivated by the persistent phenomenon of tax evasion and the challenge of tax collection during economic crises, we explore the behavior of a risk-neutral self-interested firm that may engage in tax evasion to maximize its profits. The firm evolves in a tax system which includes many of "standard" features such as audits, penalties and occasional tax amnesties, and may be uncertain as to its tax status (not knowing, for example, whether a tax amnesty may be imminent). We show that the firm's dynamics can be expressed via a partially observable Markov decision process and use that model to compute the firm's optimal behavior and expected long-term discounted rewards in a variety of scenarios of practical interest. Going beyond previous work, we are able to investigate the effect of "leaks" or "pre-announcements" of any tax amnesties on the firm's behavior (and thus on tax revenues). We also compute the effect on firm behavior of any extensions of the statute of limitations within which the firm's tax filings can be audited, and show that such extensions can be a significant deterrent against tax evasion.

1 Introduction

In recent years, there has been increasing interest in the study of optimization and optimal control problems in the area of taxation and tax policy [5, 6, 11, 19, 20]. This activity is motivated in part by the recent global financial crisis which brought the problem of tax revenue collection squarely into the fore, but also by the availability of computational power which allows researchers to delve in areas where analytical results are scarce. This work explores an optimization problem related to tax evasion, a persistent phenomenon with which most countries grapple to some extent. In particular, we seek to determine the actions of a self-interested, risk-neutral tax

P. Papadopoulou · D. Hristu-Varsakelis (✉)
Computational Methods and Operations Research Laboratory, Department of Applied Informatics, University of Macedonia, Thessaloniki, Greece
e-mail: mai16036@uom.edu.gr; dcv@uom.gr

I. C. Demetriou, P. M. Pardalos (eds.), *Approximation and Optimization*,
Springer Optimization and Its Applications 145,
https://doi.org/10.1007/978-3-030-12767-1_11

entity (we will use the term “firm”) which may engage in tax evasion in order to maximize its long-term profits. The firm will be allowed to evolve dynamically in a tax system which includes many of the features commonly used, including a tax rate, penalties for concealing profits, random audits (where the firm’s past tax statements can be checked for a number of years into the past), and occasional optional tax amnesties which the government may offer but which the firm has no advance knowledge of. In this context, which we will make precise shortly, we are interested in (i) determining the firm’s optimal behavior and expected long-term discounted rewards, (ii) finding out whether the firm can profit by reducing its uncertainty with respect to upcoming amnesties (e.g., from government announcements or “leaks” to the press), and (iii) quantifying the effect on maximal firm revenues (and thus on government revenues) of possible increases to the statute of limitations within which the firm’s past actions can be audited.

Prior work related to optimal taxation and tax-evasion modeling includes early approaches, such as [1] which examined tax evasion as a portfolio allocation problem, and subsequent improvements (e.g., [4, 18, 33]). One disadvantage of these and later analytical approaches was that, in order to remain tractable, they often took on a macroscopic viewpoint, and could not express taxpayer heterogeneity nor could they fully capture the dynamics of tax evasion [21]. In recent years, there has been interest in modeling taxpayer entities at a finer-grained level in order to study their year-to-year evolution through the tax system [11]. These dynamics may include the random transitions in a firm’s tax status (e.g., being subject to a surprise audit, or being included in an amnesty program) or the changing preferences of firms, viewed as interacting agents. In some cases, however, the richness of these models comes at a price in terms of complexity and requires a computational, rather than analytic, approach [9, 10].

The work in [11], on which this paper builds, examined a Markov-based model of the firm’s evolution whose state corresponded to its yearly tax status (i.e., being audited or not, being able to expunge previous tax records through amnesty, etc.). That work showed that if the firm is risk-neutral, then its optimal policy can be computed via dynamic programming, and produced maps of the space of tax parameters identifying those that would remove the incentive for the firm to evade taxation. One limitation of [11] was the fact that the firm knew each year whether the government intended to activate a tax amnesty, and therefore could take advantage of that knowledge as it decided its future course. In practical settings (e.g., Greece, which [11] and the present work use as a case study) this may occur when the government creates expectations either through official announcements or press leaks. Under normal circumstances, however, the firm cannot be assumed to have information on upcoming amnesties. It is thus crucial to develop models that take into account the resulting uncertainty from the point of view of the firm as to its true tax status. Doing so will allow us to explore the effects of that uncertainty on the firm’s actions with respect to tax evasion. At the same time, it is important for policy makers to know what are the consequences, if any (in terms of revenue), of information leaks.

This paper's contribution is twofold. With respect to describing the firm's behavior vis-a-vis tax evasion, we propose a model which is structurally more parsimonious and yet more realistic than its predecessors, taking into account the fact that the firm has incomplete knowledge of its tax status. Our model, in the form of a POMDP, will allow us to approximate the firm's optimal policy given the parameters of the tax environment, to investigate whether it is important for the government to be careful about the information it releases on possible tax amnesties, and to quantify the effect (in terms of revenue) of keeping taxpayers "in the dark" regarding upcoming amnesties. The proposed model is used to identify the combinations of tax penalties and audit probabilities that lead to honest behavior, and to quantify the impact that an increase in the statute of limitations on tax audits would have on tax evasion, an aspect which—to our knowledge—has not been sufficiently explored in the context of relevant dynamical models.

The remainder of this paper is structured as follows: In Section 2 we propose a POMDP model that captures the firm's evolution in the tax system. The firm seeks to maximize its discounted long-term expected profit, taking into account the rules imposed by the tax system and its awareness (or lack thereof) of any imminent opportunity for amnesty that the government may provide. Section 3 discusses the solutions obtained from our model and examines the impact of uncertainty on the firm's decisions and the circumstances under which it is useful for the firm to know the government's intentions in advance. In the same section we explore the effect of extending the statute of limitations on auditing and how it affects tax evasion.

2 A POMDP Model of Tax Evasion

We proceed to construct a mathematical description of the firm's time evolution in the tax system. We will consider a generic tax system similar to that in [11] (see that work for a fuller description), which includes a fixed tax rate on profits, random audits by the tax authorities, and tax penalties for underreporting income. For the sake of concreteness, we will use Greece as a case study [30] when it comes to selecting specific values for the various tax parameters, although the discussion applies to a much broader setting. During an audit, a firm's tax statements can be scrutinized for up to 5 years in the past, meaning that any tax evasion beyond that horizon goes unpunished. Audits tend to focus on firms that have not been audited for three or 4 years running, and thus have tax records which are about to pass beyond the statute of limitations. Finally, the Greek tax system occasionally offers a kind of optional tax amnesty [2], termed "closure," in which a firm may pay a one-time fee for excluding past tax statements from a possible audit. The use of tax amnesties is not unusual, with hundreds of cases documented across many countries, including the USA [25], India [7], and Russia [3]. In Greece, the closure option mentioned above was used during 1998–2006 [13, 14] and was again considered more recently [15]. The availability of the closure option gives an incentive for firms who evade taxation to pay *some* amount in taxes where they might otherwise pay none, as it increases the chance of an audit for those who do not participate. However, it has been shown [11] to encourage tax evasion.

Operationally, at the end of each fiscal year the firm declares its net profit to the government, and also its intent to use the closure option (and pay the associated fees) if it becomes available. In an important—and realistic—departure from previous models [11] the firm does not know in advance whether the option is to be made available or not for the current year; that information is released only after the tax filing deadline. Thus the firm must decide on its tax-evasion policy without knowing, for example, whether it will be able to expunge any imminent tax misdeeds by availing itself of the option. This introduces uncertainty as to the firm's tax status and gives rise to a POMDP [8, 17] which we describe next.

2.1 The Firm's State and Action Sets

Using the notation from [11], we will let $s_k \in \mathcal{S}$ be the tax status of representative firm in year $k = 0, 1, 2, \ldots$ with

$$\mathcal{S} = \{V_1, \ldots, V_5, O_1, \ldots, O_5, N_1, \ldots, N_5\}, \tag{1}$$

where

- V_i: the firm is being audited for the last $i = 1, \ldots, 5$ tax filings,
- O_i: the government is making the closure option available to the firm, whose last audit or closure option usage occurred $i = 1, \ldots, 5$ years ago, and
- N_i: the firm is not being audited nor has a closure option available, and its last audit or closure option usage happened $i = 1, \ldots, 5$ years ago.

For the sake of notational convenience, we will sometimes refer to the elements of $\mathcal{S}$ by integer, in their order of their appearance, i.e., $V_1 \to 1$, $V_2 \to 2$, $\ldots$, $N_5 \to 15$. We note that $\mathcal{S}$ contains five "copies" of each type of tax status (V, O, N) corresponding to the 5-year statute of limitations on tax evasion. Of course, $\mathcal{S}$ (and the discussion that follows) could be appropriately generalized to model a longer, length L time-window on which the government has the chance to detect tax evasion:

$$\mathcal{S} = \{V_1, \ldots, V_L, O_1, \ldots, O_L, N_1, \ldots, N_L\}.$$

We will have more to say about this in Section 3.

We will let $\mathcal{A} = \{1, 2\} \times [0, 1]$ be the firm's action set, where at year k the firm selects $a_k \in \mathcal{A}$, $a_k = [v_k, u_k]^T$, with $v \in \{1, 2\}$ denoting the firm's decision to apply for usage of the closure option ($v_k = 1$) or to forgo the closure option ($v_k = 2$), and $u_k \in [0, 1]$ being the fraction of profits that the firm decides to conceal. Based on the above, the firm's state at time k will be the vector

$$x_k = [s_k, h_k^T]^T, \tag{2}$$

where $s_k \in \mathcal{S}$, and $h_k \in [0, 1]^5$ will contain a history of the firm's latest five decisions with respect to tax evasion.

We note that our model differs from that of [11] in one important and practical point. In [11], the state vector included one additional element, corresponding to whether or not the closure option is available to the firm or not. In our case, the firm has no such knowledge; it can declare its wish (v_k) to avail itself of the option (should the government make it available after the tax filing deadline) but is forced to commit to a decision on tax evasion (u_k) in advance. Put in other words, without any information as to the government's intent to offer the closure option, the states O_i are indistinguishable to the firm from their N_i counterparts at the time the firm makes its decisions.

2.2 State Evolution

Based on the above discussion (see also [11]), the firm will evolve in $\mathcal{S} \times [0, 1]^5$ according to

$$x_{k+1} = Ax_k + Ba_k + n_k, \quad x(0) \text{ given}, \tag{3}$$

where

$$A = \begin{bmatrix} 0 & \\ & H \end{bmatrix} \quad H = \begin{bmatrix} 0&1&0&0&0 \\ 0&0&1&0&0 \\ 0&0&0&1&0 \\ 0&0&0&0&1 \\ 0&0&0&0&0 \end{bmatrix}, \quad B = \begin{bmatrix} 0&0 \\ \vdots&\vdots \\ 0&0 \\ 0&1 \end{bmatrix}, \quad n_k = \begin{bmatrix} w_k \\ 0_{5\times 1} \end{bmatrix}, \tag{4}$$

and the term $w_k \in \mathcal{S}$ is a random process whose transition probabilities depend on whether the firm applies for use the closure option or not, and whether the government ultimately decides to grant it at that particular year. Based on our earlier description of the tax system, we can represent the process driving w_k in graphical form using two transition diagrams, one for the case where $v_k = 1$ (the firm decides to use the closure option, Figure 1) and another when $v_k = 2$ when the firm declines the use of the option (Figure 2).

For example, let us assume that in Figure 1 the firm has the tax status O_2. This means that the firm will pay to exclude its last two tax statements from any audits, thus "cleaning its slate," and will now file its next tax statement—the only one subject to a possible audit next year. The firm's possible transitions are to O_1 (if the government does grant the option again) where the just-filed tax statement will be expunged; to V_1 where it will be audited; or to N_1, where the firm will avoid an audit but its tax statement will be kept in waiting, for possible future audits or closures. The transition diagram in Figure 2 (where the firm has declined the closure option) operates in a similar manner, with the firm never transitioning to an O_i state.

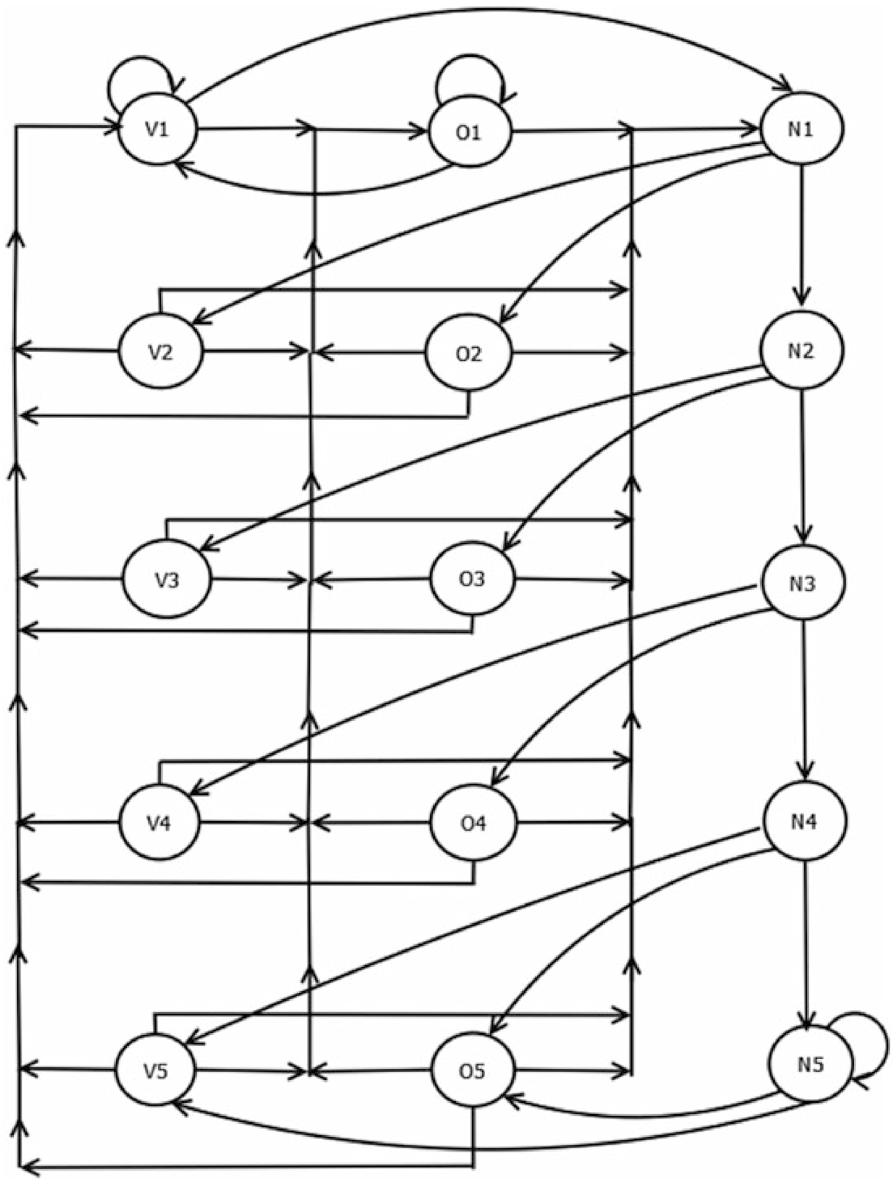

Fig. 1 Transition diagram when the firm asks to use the closure option. The values of the various transition probabilities are discussed in the text and in the Appendix

However, the transition probabilities from O_i states to audit states V_i will be higher compared to the previous case (see below and the Appendix for numerical values) to reflect the fact that once the government offers the closure option, any firm that opted out has a higher chance of being audited because its peers that opted in have now removed themselves from the audit pool.

2.2.1 Transition Probabilities

Based on the transition diagrams of Figures 1and 2, the transition probabilities of the process w_k (and thus the firm's evolution from one tax status to another) are

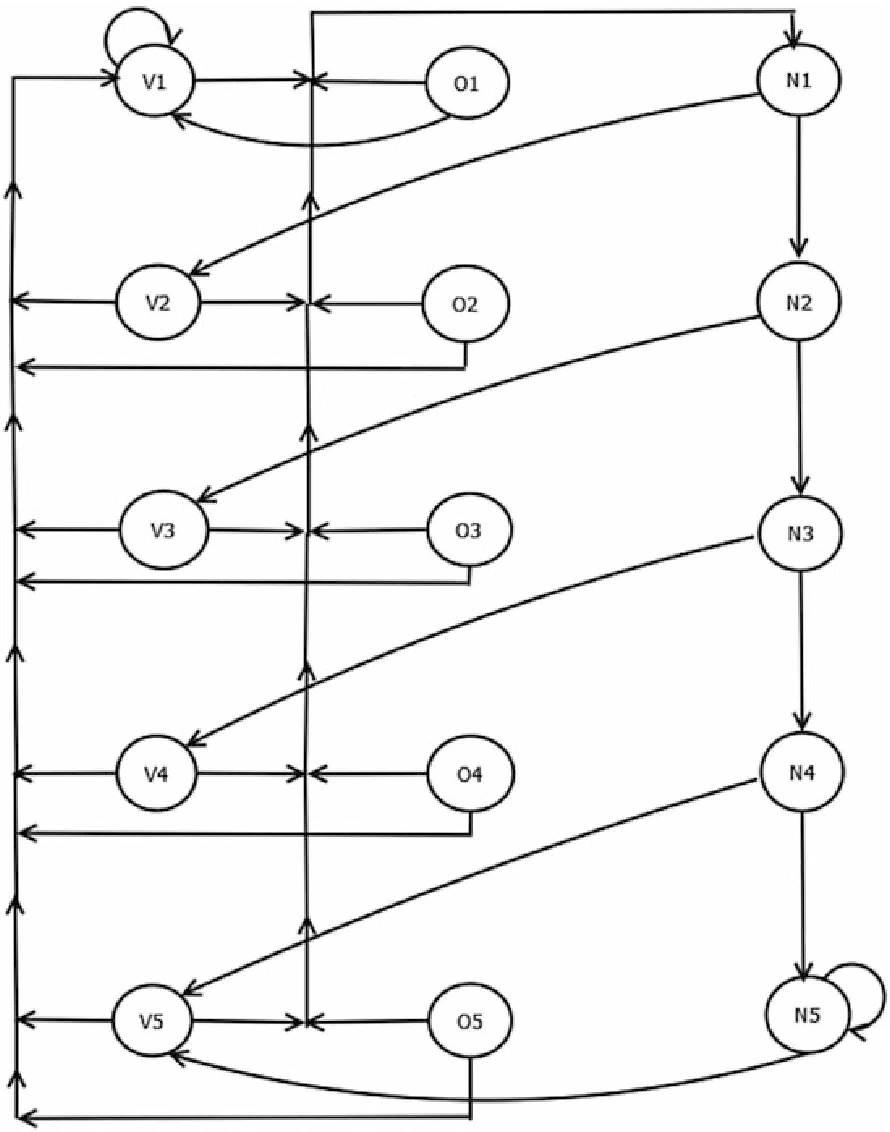

Fig. 2 Transition diagram when the firm decides not to use the closure option. The values of the various transition probabilities are discussed in the text and in the Appendix

$$Pr\left(s_{k+1} = i | s_k = j, a_k = (v_k, u_k)\right) = T_{ij}\left(v_k\right), \quad i, j \in \{1, \ldots 15\}, \tag{5}$$

where T_{ij} is the (i, j)-th element of the transition matrix T given by

$$T\left(v_k\right) = \begin{cases} T_1 \text{ if } v_k = 1 \text{ (apply for closure)} \\ T_2 \text{ if } v_k = 2 \text{ (decline closure)} \end{cases} \tag{6}$$

and the T_1 and T_2 are shown in the Appendix.

2.3 Firm Rewards

Let Π denote the firm's annual profit (assumed to be constant), r the tax rate, β the annual penalty rate for any unpaid taxes applied in the event of an audit, and ℓ the cost of closure as a fraction of the firm's profits. Then the firm's reward is as per [11] (given here again for completeness):

$$g(x,a) = g\left([s, h^T]^T, a\right) = \Pi \cdot \begin{cases} 1 - r + ru & s \in \{11, \dots, 15\} \\ 1 - r + ru - \ell(s-5) & s \in \{6, \dots, 10\} \\ 1 - r + ru - r\sum_{i=1}^{s}[h]_{6-i} & \\ -\frac{3}{5}\beta r \sum_{i=1}^{s} i[h]_{6-i} & s \in \{1, \dots, 5\} \end{cases}, \tag{7}$$

where we have labeled elements of $\mathcal{S}$ by integer, and $[h]_i$ denotes the i-th element of the vector h. The top term on the right-hand side of (7) corresponds to the reward obtained if the firm is neither audited nor using the closure option and conceals an amount of Πu. The second term applies when the firm uses the option and thus pays ℓ per year since its last audit or closure. The last term is the firm's reward in the event of an audit, where its past history of tax evasion is used to calculate the back taxes and penalty owed.

2.4 Firm Observations, Belief, and Value Function

As we have already stated, the firm does not know if the government is to offer the closure option until *after* it has filed its taxes for the year. This means that—unless it is being audited—the firm is uncertain of its tax status (the first element of its state vector) which may be N_i (no option available) or O_i (option available if the firm is willing to pay for it). In practice, the firm may have some information from the press, government, or market sources, as to whether a new round of closure may be imminent. Of course, the information may not always be correct. Also, it is in the government's interest to know what effect any such "leaks" would have on firm behavior (and thus on tax revenues).

We will let $\hat{\mathcal{S}}$ denote the set of observations that the firm makes of its tax status, with

$$\hat{\mathcal{S}} = \{\hat{V}_1, \dots, \hat{V}_5, \hat{O}_1, \dots, \hat{O}_5, \hat{N}_1, \dots, \hat{N}_5\}. \tag{8}$$

Thus, at time k, the firm observes $\hat{s}_k \in \hat{\mathcal{S}}$, based on which it must decide on its course of action a_k. When the firm is in an audit state, it is of course aware of that situation and thus

$$P\left(\hat{s}_k = \hat{V}_i | s_k = V_i\right) = 1, \tag{9}$$

while in the case when the firm has an O_i or N_i status,

$$P\left(\hat{s}_k = \hat{O}_i | s_k = O_i\right) = z_O, \quad P\left(\hat{s}_k = \hat{N}_i | s_k = O_i\right) = 1 - z_O, \tag{10}$$

$$P\left(\hat{s}_k = \hat{N}_i | s_k = N_i\right) = z_N, \quad P\left(\hat{s}_k = \hat{O}_i | s_k = N_i\right) = 1 - z_N, \tag{11}$$

where z_O is the probability of correctly distinguishing an O_i state from its N_i counterpart, while z_N is the probability of correctly observing N_i versus O_i.

Given the state evolution equation (3) and observations $\hat{s}_k$ we can construct the firm's belief as a probability distribution over its states, which is to be updated with every new observation made by the firm [17]. It is important to note that the uncertainty in the firm's observations has to do solely with its tax status, i.e., the first element s_k of the state vector $x_k = [s_k, h_k^T]^T$. The rest of the state vector, h_k, is the firm's tax history which is of course always known to the firm. In light of this, we may define the observations in the entire state space to be

$$\hat{x}_k = [\hat{s}_k, h_k^T]^T,$$

and Equations (9)–(11) determine the observation probabilities of all states, so that it is sufficient to consider the belief $b(s)$ as a probability distribution on $\mathcal{S}$, instead of $b(x)$ on the entire state space. The firm's belief $b_k(s_k)$ after taking action a_k and observing $\hat{s}_{k+1}$ will then be updated to

$$b_{k+1}(s_{k+1}) = c\ P(\hat{s}_{k+1}|s_{k+1}) \sum_{s_k \in \mathcal{S}} Pr(s_{k+1}|s_k, a_k) b_k(s_k), \tag{12}$$

where c is a normalizing constant, and in (12) we have used the fact that observations depend only on the state the firm is in and not on its actions.

Given the firm's current belief, $b_k(s_k)$, as to its tax status, known tax history h_k, and action a_k, we can calculate its expected reward over the belief distribution, based on Equation (7):

$$R(b_k, h_k, a_k) = \sum_{s_k \in S} g\left([s_k, h_k^T]^T, a_k\right) b_k(s_k). \tag{13}$$

Assuming an infinite time horizon, the firm is then faced with the problem choosing its closure and tax-evasion decisions $a_k = \pi(b_k, h_k) = (v_k, u_k)$, so as to maximize, over the policy π, its discounted expected reward:

$$J^\pi(b_0, h_0) = \sum_{k=0}^{\infty} \gamma^k \mathop{\mathbb{E}}_{w_k} \left\{ g\left([s_k, h_k^T]^T, a_k\right) | b_0, h_0, \pi \right\} \tag{14}$$

subject to the dynamics (3) and observations (12), where $\gamma \in (0, 1]$ is a discount factor. The optimal value function J^* in Equation (14) obeys the well-known [8, 17] Bellman equation

$$J^*(b_k, h_k) = \max_{a_k} \left[R(b_k, h_k, a_k) + \gamma \sum_{\hat{s}_{k+1} \in \hat{\mathcal{S}}} P(\hat{s}_{k+1}|b_k, a_k) J^*(b_{k+1}, h_{k+1}) \right] \tag{15}$$

with

$$P(\hat{s}|b_k, a_k) = \sum_{s, s_k \in \mathcal{S}} P(\hat{s}|s) Pr(s|s_k, a_k) b_k(s_k).$$

2.5 Solving for the Firm'S Optimal Policy

Solving Equation (15) for the firm's optimal policy is generally difficult given (i) the fact that the number of states that the firm can occupy is not countable (recall that the firm's state includes its history of tax-evasion decisions, which are real numbers in the unit interval), and (ii) the difficulties associated with partial observability of the state. As is often the case when it comes to optimizing POMDPs, we will be able to make progress only by approximating the optimal value function, by means of an iterative process [24, 27].

Because we have assumed that the firm is risk-neutral, the reward function (7) is linear in the fraction of the profits to be concealed, u, making J^* linear as well. This implies that J^* will be maximized at the boundary of u's feasible region, meaning that we only need to consider the two extreme values $u = 0$ or $u = 1$ each year (i.e., be completely honest or conceal as much profit as possible—see [11] for a fuller discussion on this point). This yields a significant reduction in computational complexity, because it will be sufficient to consider $h_k \in \{0, 1\}^5$ and solve the Bellman equation only for a finite set of $|\mathcal{S}| \cdot 2^5 = 480$ states. This is to be compared to the 869 states in [11]. Of course, our model is more challenging computationally, because of the uncertainty in state observations.

To solve for the firm's optimal policy, we used the PERSEUS point-based value iteration algorithm [28]. Point-based algorithms became popular as methods of approximating POMDP policies [22, 23]. They rely on the fact that performing many fast approximate updates to a set of policy/value samples often results in a more useful value function than performing a few exact updates [16, 26]. The algorithm in [29] differs from other point-based algorithms in that at each iteration it "backs up" only a random subset of points in the firm's belief space. Doing so leads to computational savings so that the method can afford to use a larger number of samples compared to other point-based methods, and obtain better accuracy [29, 32]. In the next section we present a series of numerical experiments, where we optimize the firm's policy using [28], and discuss the results.

3 Results and Discussion

To begin with, we would like to determine the optimal policy and expected firm revenue depending on the availability of closure and the probability of the firm correctly observing its own state. We considered three cases with respect to the closure option: (i) available every fiscal year, (ii) never available, and (iii) given randomly, with a fixed probability of $q = 0.2$. With respect to the firm's observation probabilities, we considered the cases when the firm is (i) fully aware of the government's decision with respect to closure (i.e., it can correctly observe whether it is in a N state or an O state; this will allow us to validate our model against known solutions), or (ii) has no information on the government's intentions and must guess as to its state (i.e., $z_O = z_N = 0.5$). The remaining tax parameters were set as in [11] using Greece as a case study for the sake of being specific and to facilitate comparisons with previous work, namely, a tax rate of $r = 0.24$, penalty rate $\beta = 0.24$, cost of closure $\ell = 0.023$, and discount factor $\gamma = 0.971$ corresponding to a 3% annual rate of inflation.

3.1 Model Validation: The Case of "Perfect" Observations

In the case where the closure option is given randomly but the firm can make perfect observations of its state, the probability of a "correct" tax status observation in Equations (11)–(10) is $z_O = z_N = 1$. Then, the firm's optimization problem reduces to a Markov decision process (MDP) similar to that in [11] which can be solved easily via value iteration. Furthermore, the cases where the closure option is either always or never available also imply perfect observations because then the firm (who is not being audited) can safely conclude that it is in an O_i or N_i state, respectively. Table 1 shows the firm's optimal discounted expected reward in each case. As we can see, the reward is highest when the government offers the closure option with probability 1. In that case, the optimal policy was a constant $a = (1, 1)$, i.e., for the firm to always use the closure option every single year and to conceal all profit. In fact, from the point of view of government revenue, it is best to never offer the closure option. The firm's optimal policies agreed with the exact solutions obtained by [11], as did the firm's expected rewards (to within 0.11%).

Table 1 Comparison of the firm's discounted expected reward, under perfect state observation, with $r = 0.24$, $\beta = 0.24$, $\ell = 0.023$, and a 5% overall audit probability. Numbers are in % of the firm's annual profit Π, with a discount factor γ corresponding to a 3% annual rate of inflation

Closure option	Always available	Available with 20%	Never available
Exp. reward	3354.3	3307.1	3254.5

3.2 The Role of Uncertain Observations

Next, we investigate whether the firm's uncertainty as to its tax status may affect its tax-evasion behavior and, ultimately, its expected revenue. To this end, we must choose a specific probability for the closure option being offered at any year. We set $q = 1/5$, which is in some agreement with empirical data from Greece [11]; this rate corresponds to the government attempting to collect some revenue—in the form of closure fees—from firms which it does not have the resources to audit and whose tax filings are about to pass beyond the five year statute of limitations.

We solved for the firm's optimal policy and computed its expected revenue under: (i) perfect observations, where the firm always knows its tax status (equivalently, the government's decision to offer closure or not, $z_O = z_N = 1$), (ii) observations which are 50% correct, i.e., the firm knows nothing about the government's decision and is merely guessing ($z_O = z_N = 0.5$), and (iii) observations which are 90% correct because, for example, the government may have hinted at its intentions regarding closure ($z_O = z_N = 0.9$). Table 2 shows the firm's discounted expected revenues for these three cases. We observe that the firm's revenues are unchanged when the probability of a correct tax status observation changes from 50% to 90%, with a slight difference compared to the 100% (prefect observations) case. This seems counterintuitive, because one would expect the firm to be able to use any information on its true state to its advantage. The explanation, however, can be found in the firm's optimal policy which is constant and identical for all cases in Table 2, namely, that the firm's best action in the beginning of each fiscal year is to apply for the closure option ($v = 1$) and declare to the government as little profit as possible ($u = 1$) regardless of its belief of being in an O_i vs an N_i state.

Previous work [11] has shown that the tax parameters in the range used in Greece in recent years are such that they encourage tax evasion. What we find here is that the incentive they create for tax evasion is such that the firm's uncertainty as to its state is unimportant because the optimal policy—even in the face of that uncertainty—is to always "cheat." However, this might change—and should, from a policy viewpoint—for other combinations of tax parameters (e.g., higher penalties). Numerical experimentation shows that there is a range of values for r, β, and ℓ where state observation probabilities do become important with regard to the firm's optimal reward and policy. For example, for a tax penalty of $\beta = 9$ (that is, 9 times any tax that went unpaid due to tax evasion) and closure cost $\ell = 0.04$ (or 4% of the firm's profit) we notice a difference in expected rewards between the cases of 50%

Table 2 Comparison of the firm's expected reward when the closure option is offered with a probability of $q = 0.2$

Prob. of correct observation	100%	50%	90%
Exp. reward	3307.1	3309.7	3309.7

The experiments were run with $r = 0.24$, $\beta = 0.24$, $\ell = 0.023$, and a 5% overall audit probability. Numbers are expressed in % of the firm's annual profit, discounted at a 3% annual rate of inflation

Table 3 Comparison of the firm's expected rewards with 50% vs. 90% probability of correctly observing its tax status when the closure option is available with probability $q = 0.2$

Prob. of correct observation	50%	90%
Exp. reward	2630.6	2648

The tax parameters were set to $r = 0.24$, $\beta = 9$, $\ell = 0.04$, and a 5% overall audit probability. Rewards are expressed in % of the firm's annual profit, discounted at a 3% annual rate of inflation

and 90% probability of correct tax status observation, as shown in Table 3 where the firm obtains higher rewards through tax evasion when more certain of its tax status.

The difference in discounted expected rewards for the two cases of Table 3 varies for other combinations of tax parameters, and—although it is beyond the scope of this paper—it would be of interest to "map" the (r, β, ℓ) space in order to quantify the increase in expected reward that the firm could obtain as a function of the "amount of uncertainty" in its state observations.

With respect to the closure option (and other such amnesties), our results suggest that unless the "surrounding" tax environment is sufficiently strict in terms of penalties and cost of the option, the latter should not be offered because it helps the firm achieve a higher expected reward through tax evasion (which means that tax revenues are reduced). If the tax parameters are appropriately set, then the closure option could be useful if taxpayers can be kept from knowing about it in advance.

3.3 The Role of Statute of Limitations

Our prototypical tax system has thus far included a 5 year "window" within which the government is allowed to audit the firm's tax statements. This has been the legal window in Greece, for example. However, the effect of extending this statute of limitations has not been investigated. To quantify the effect of such an extension, we considered the set of audit probability p and tax penalty β combinations, and determined those values for which the firm's optimal policy is to behave honestly (i.e., use $u = 0$ in every state). As a result of the firm's risk-neutrality (and consequent linearity of the reward function) there exists an "honesty boundary" in the form of a curve in the (p, β) space, above which (high penalties) the firm's decision is to never conceal any profit, and below which (lower penalties) the firm conceals its profits in at least one state.

Any point on the boundary can be computed relatively easily by fixing p and using bisection on β, each time calculating the firm's optimal policy and checking whether it is completely "honest." For the case where there is no uncertainty in state observations this can be done via simple value iteration. We thus modified our model to increase the statute of limitations, from the nominal $L = 5$ to $L = 6, \ldots, 10$ years, and computed the corresponding honesty boundaries for comparison with the 5-year case. The various boundaries are shown in Figure 3 for the case where

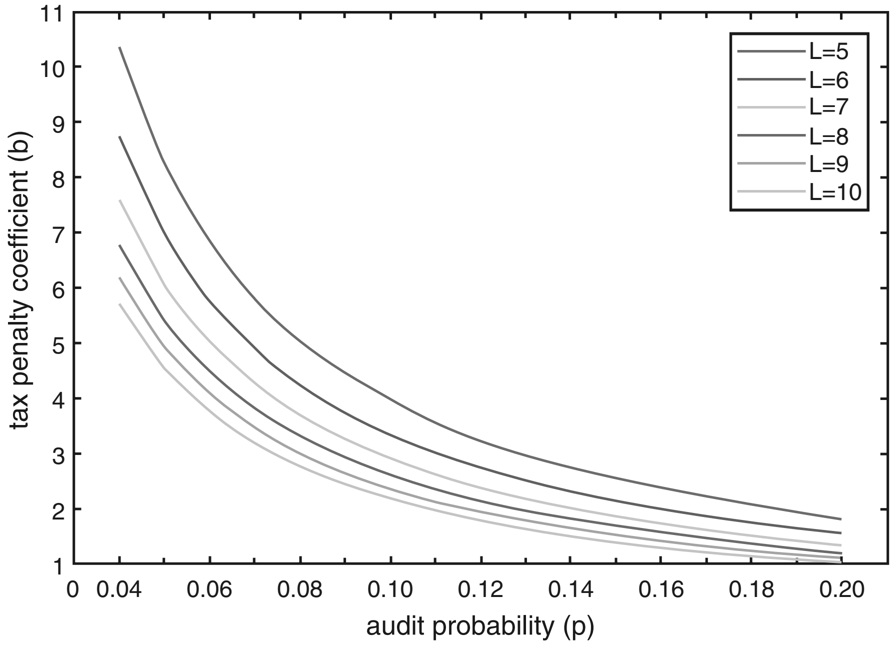

Fig. 3 Firm's "honesty boundaries" in the p-β space when the closure option is never available, as the statute of limitations on tax audits increases from 5 to 10 years. Above each curve the firm's optimal policy is $u_k = 0$, i.e., to always declare all profit

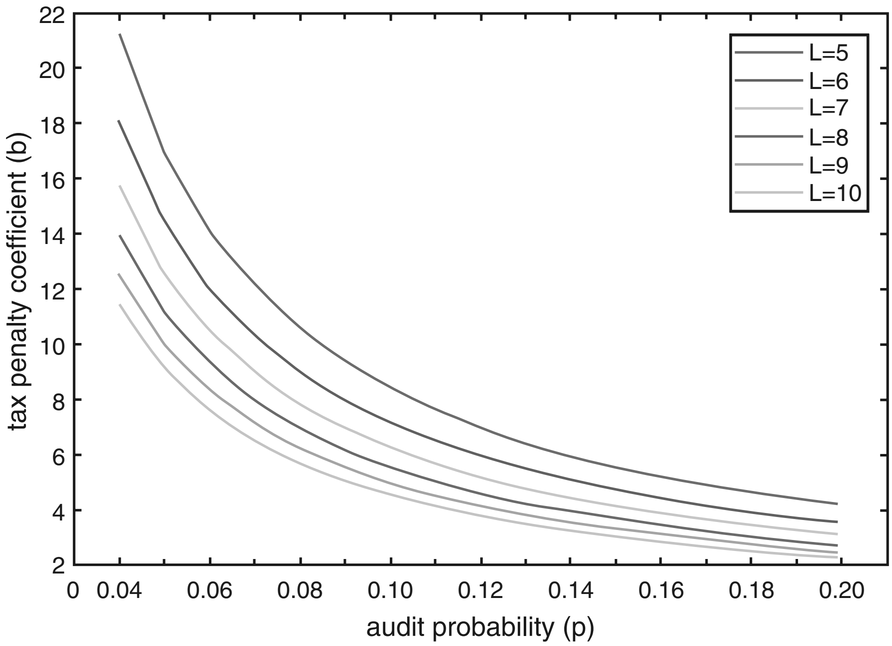

Fig. 4 Firm's "honesty boundaries" in the p-β space when the closure option is available with probability $1/L$ each year, as the statute of limitations L on tax audits increases from 5 to 10 years. Above each curve the firm's optimal policy is $u_k = 0$, i.e., to always declare all profit

Table 4 Average percentage reduction in the tax penalty β which is necessary to eliminate tax evasion, as the statute of limitations on audits, L, increases

Year-to-year/closure prob	q=0	q=1/L
5–6	15.67%	15.10%
6–7	12.87%	12.88%
7–8	10.61%	11.7%
8–9	8.91%	10.28%
9–10	7.63%	8.85%
Overall 5–10	44.74%	46.6%

the closure option is never available, and in Figure 4 for closure available with a fixed nonzero probability. As the time horizon on audits increased, we adjusted the probability of the closure option being offered to $q = 1/L$ so that the firm has, on average, one chance at closure within the statute of limitations period, regardless of that period's length. The first thing to note is that as the audit probability increases there is a lower tax penalty required to enforce complete honesty, and that the marginal effect is most pronounced for low p. Furthermore, as the statute of limitations increases, the entire honesty boundary moves downward, so that again lower penalties are required to remove the incentive for tax evasion. This is to be expected, because the government is given more opportunities to catch the firm at cheating and thus the firm behaves more honestly. Thus, from a technical point of view, the government could extend the auditing period and be able to use more reasonable penalties to deter tax evasion, although unless the audit rates are also increased, those penalties will still have to be very high (a factor of almost 10 for a 5% audit probability even for $L = 10$). In Figure 3, where there is no probability of amnesty, the honesty "threshold" is always lower—and involves less exorbitant tax penalties—than that in Figure 4.

To quantify the reduction of the honesty threshold for β as L varies, we calculated the threshold's year-to-year average percent variation, and also the total variation between the 5- and 10-year cases. The results are shown in Table 4. We note that each yearly extension of the statute of limitations reduces the average tax penalty coefficient required to make tax evasion unprofitable by a significant margin. Combined with our previous results, this points towards (i) a long statute of limitations, (ii) an increased audit probability, and (iii) avoidance of tax amnesties, as useful policy tools.

4 Conclusions

Motivated by the problem of tax evasion and the need for computational tools that may be used to elucidate the behavior of tax entities, we have described a POMDP which models the behavior of a self-interested risk-neutral firm in a tax system which includes random audits, penalties, and occasional amnesties. Using our model, together with a point-based approximation technique, we were able to compute the firm's optimal decisions regarding whether or not to apply for amnesty (in case it is given) and how much of its profit to conceal from authorities.

Our model is more realistic than previously proposed Markov-based models of firm behavior [11, 12] in that the firm is uncertain as to its tax status and must file its taxes before it knows whether its latest filing will be subject to amnesty or not. Using the Greek tax system as a case study, we confirmed previous results suggesting that for range of tax parameters commonly in use there, the firm's optimal policy is to always opt for amnesty and conceal as much of its profit as possible. This is a consequence of the firm's risk neutrality, and the fact that the audit probability and tax penalty coefficients are too low to be effective.

With respect to the role of the firm's conditional observations, we saw that if the tax parameters do not have sufficient deterrent value, the firm's uncertainty as to its true tax status is unimportant. Among other things, this means that the government need not worry about information "leaks" regarding the tax amnesty which, however, is rendered less effective. If tax revenues are to increase (in part by keeping firms in the dark about an upcoming amnesty), some combination of the tax penalties, audit rates, and cost of amnesty must rise significantly from the levels studied here.

Finally, we identified the set of tax rates and tax penalties which eliminate tax evasion in our model. The resulting curves in the tax rate—tax penalty space, show, among other things, how frequent tax audits must be in order for the government to be able to discourage tax evasion using realistic tax penalties. We also found that the "border" between tax evasion and honest behavior shifts significantly towards lower penalties as the statute of limitations on tax audits is extended.

Opportunities for future work include a computational study to explore the space of tax rates, penalty coefficients [31], and audit rates in order to find the settings in which the firm's uncertainty as to its true state has the greatest impact on its optimal long-term revenues. From a policy viewpoint, this would suggest parameter values for which government revenues increase as long as the government can keep any upcoming amnesty a surprise. Finally, it would be interesting to extend the model presented here by considering a risk-averse firm and explore ways of solving resulting POMDP in the presence of the nonlinearity that risk-aversion introduces in the reward function, in the spirit of [12].

Appendix

Transition Matrix when Firm Applies for Closure

Markov transition matrix for the case where the firm asks to use the closure option. The statute of limitations is assumed to be $L = 5$ years

$$T_1 = \begin{bmatrix} p_V & p_V & p_V & p_V & p_V & p_O & p_O & p_O & p_O & p_O & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{Nf} & p_{Nf} \\ q & q & q & q & q & q & q & q & q & q & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q & q \\ p_{IV} & p_{IV} & p_{IV} & p_{IV} & p_{IV} & p_{IO} & p_{IO} & p_{IO} & p_{IO} & p_{IO} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{IN} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{IN} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{IN} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{INf} & p_{INf} \end{bmatrix}.$$

Here, p is the fraction of tax filings which the government can audit each year regardless of firm status. For our purposes p is nominally set to 0.05 (5%). A 20% of those audits (1%) are spread over firms that have been unaudited for up to 3 years, and the remaining 80% of audits (nominally 4%) are for those that have gone unaudited for 4 or more years; q is the probability of the government offering the closure option; $p_V = (p - 0.8p)/4$ is the probability of a repeat audit (i.e., for the second year in a row); $p_O = (p/5)/4$ is the probability of being audited after using closure; $p_N = (p/5)/4$ is the probability of an audit if the firm has been unaudited for 1–3 years; $p_{N_f} = (p/5)4$ is the probability of being audited if the firm has been unaudited for 4 or more years (and thus some of its tax statements are about to pass beyond the statute of limitations); and $p_{IV} = 1 - q - p_V$, $p_{IO} = 1 - q - p_O$, $p_{IN} = 1 - q - p_N$, and $p_{INf} = 1 - q - p_{Nf}$.

Transition Matrix when Firm Declines Closure

Markov transition matrix for the case where the firm forgoes the closure option. The statute of limitations is assumed to be $L = 5$ years

$$T_2 = \begin{bmatrix}
p_V & p_V & p_V & p_V & p_V & p_O & p_O & p_O & p_O & p_O & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_N & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{Nf} & p_{Nf} \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
p_{lV} & p_{lV} & p_{lV} & p_{lV} & p_{lV} & p_{lO} & p_{lO} & p_{lO} & p_{lO} & p_{lO} & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{lN} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{lN} & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{lN} & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & p_{lNf} & p_{lNf}
\end{bmatrix}.$$

Again, p is the fraction of tax filings which the government can audit each year regardless of firm status; $p_V = (p - 0.8p)/4$ is the probability of a repeat audit (i.e., for the second year in a row); $p_O = 3[(p/5)/4]$ is the probability of being audited after using closure; $p_N = (p/5)/4$ is the probability of an audit if the firm has been unaudited for 1–3 years; $p_{N_f} = (p/5)4$ is the probability of being audited if the firm has been unaudited for 4 or more years (and thus some of its tax statements are about to pass beyond the statute of limitations); and $p_{lV} = 1 - p_V$, $p_{lO} = 1 - p_O$, $p_{lN} = 1 - p_N$, and $p_{lNf} = 1 - p_{Nf}$.

References

1. Allingham, P., Sandmo, H.: Income tax evasion: a theoretical analysis. J. Public Econ. **1**(6), 988–1001 (1972)
2. Alm, J., Beck, W.: Tax amnesties and tax revenues. Pub. Fin. Rev. **18**(4), 433–453 (1990)
3. Alm, J., Rath, D.M.: Tax policy analysis: the introduction of a Russian Tax Amnesty. Technical Report, Georgia State University, Andrew Young School of Policy Studies (1998)
4. Baldry, J.C.: Tax evasion and labour supply. Econ. Lett. **3**(1), 53–56 (1979). ISSN: 0165-1765
5. Bloomquist, K.M., Koehler, M.: A large-scale agent-based model of taxpayer reporting compliance. J. Artif. Soc. Soc. Simul. **18**(2), 20 (2015)
6. Chari, V.V., Nicolini, J.P., Teles, P.: Optimal Capital Taxation, University of Minnesota and Federal Reserve Bank of Minneapolis, Federal Reserve Bank of Minneapolis- Universidad Di Tella, Universidad Autonoma de Barcelona, Banco de Portugal, Catolica Lisbon SBE, CEPR (2016)
7. Das-Gupta, A., Mookherjee, D.: Tax amnesties in India: an empirical evaluation. World Dev. **23**(12), 2051–2064 (1995)
8. Dutech, A., Scherrer, B.: Partially Observable Markov Decision Processes. Wiley, Hoboken (2013)

9. Gao, S., Xu, D.: Conceptual modeling and development of an intelligent agent-assisted decision support system for anti-money laundering. Exp. Syst. Appl. **36**(2), 1493–1504 (2009)
10. Garrido, N., Mittone, L.: Tax evasion behavior using finite automata: experiments in Chile and Italy. Exp. Syst. Appl. **39**(5), 5584–4492 (2012)
11. Goumagias, N.D., Hristu-Varsakelis, D., Saraidaris, A.: A decision support model for a tax revenue collection in Greece. Decis. Support. Syst. **53**(1), 76–96 (2012)
12. Goumagias, N.D., Hristu-Varsakelis, D., Assael, J.: Using deep Q-learning to understand the tax evasion behavior of risk-averse firms. Exp. Syst. Appl. **101**, 258–270 (2018)
13. Greek Ministry of Finance, Law N.3259/2004 (pol.1034/2005) (in Greek) (2004)
14. Greek Ministry of Finance, Law N.3697/2008 (pol.1130/2008) (in Greek) (2008)
15. Greek Ministry of Finance, Law N.4337/2015 (pol.4337/2015) (in Greek) (2015)
16. Kaplow, R.: Point-based POMDP solvers: survey and comparative analysis. Master's thesis, McGill University (2010)
17. Krishnamurthy, V.: Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing. Cambridge University Press, Cambridge (2016)
18. Kydland, F.E., Prescott, E.C.: Dynamic optimal taxation, rational expectations and optimal control. J. Econ. Dyn. Control **2**(1), 79–91 (1980)
19. Levagi, R., Menoncin, F.: Optimal dynamic tax evasion: a portfolio approach. J. Econ. Dyn. Control **37**(11), 2157–2167 (2011)
20. Liu, A.A.: Tax evasion and optimal environmental taxes. J. Environ. Econ. Manage. **66**(3), 656–670 (2013)
21. Martinez-Vazquez, J., Rider, M.: Multiple modes of tax evasion: theory and evidence. Natl. Tax J. **58**(1), 51–76 (2005)
22. Pineau, J., Gordon, G., Thrun, S.: Point-based value iteration: an anytime algorithm for POMDPs. In: Proceedings of International Joint Conference on Artificial Intelligence, Acapulco, Mexico (2003)
23. Pineau, J., Gordon, G., Thrun, S.: Anytime point-based approximations for large POMDPs. J. Artif. Intell. Res. **27**, 335–380 (2006)
24. Poupart, P.: Exploiting structure to efficiently solve large scale partially observable Markov decision processes. Ph.D. Thesis, Department of Computer Science, University of Toronto (2005)
25. Ross, J.M., Buckwalter, N.D.: Strategic tax planning for state tax amnesties evidence from eligibility period restrictions. Public Finance Rev. **4**(3), 275–301 (2013)
26. Shani, G., Pineau, J., Kaplow, R.: A survey of point-based POMDP solvers. Auton. Agent Multi-Agent Syst. **27**(1), 1–51 (2012)
27. Sondik, E.J.: The optimal control of partially observable Markov decision processes. Ph.D. Thesis, Stanford University (1971)
28. Spaan, M.T.J.: POMDP solving software - perseus randomized point-based algorithm. Universiteit van Amsterdam (2004)
29. Spaan, M.T.J., Vlassis, N.: Perseus: randomized point-based value iteration for POMDPs. J. Artif. Intell. Res. **24**, 195–220 (2005)
30. Tatsos, N.: Economic Fraud and Tax Evasion in Greece (in Greek). Papazisis Publishings, Athens (2001)
31. Vasin, A., Vasina, P.: Tax optimization under tax evasion, the role of penalty constraints. Econ. Educ. Res. Consortium **1**, 1–44 (2009)
32. Vlassis, N., Spaan, M.T.J.: A fast point-based algorithm for POMDPs. In: Benelearn 2004: Proceedings of Annual Machine Learning Conference of Belgium and the Netherlands, pp. 170–176, (2004). Also presented at the NIPS 16 workshop 'Planning for the Real-World', Whistler, Canada, Dec. 2003
33. Yitzhaki, S.: Income tax evasion: a theoretical analysis. J. Public Econ. **3**(2), 201–202 (1974)