A Markov-Based Decision Support Model for Tax Evasion in Greece

by

Goumagias D. Nikolaos

B.S., Aristotle University of Thessaloniki (2005)

M.S., Lancaster University (2006)

Submitted to the Department of Applied Informatics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

at the

UNIVERSITY OF MACEDONIA

October 2013

Author ....

Department of Applied Informatics

October 21, 2013

Certified by ....

Dr. Dimitrios Hristu-Varsakelis

Associate Professor

Thesis Supervisor

Certified by ....

Dr. Dimitrios Papanastasiou

Professor

Thesis Committee Member

Certified by ....

Dr. Stavros Tsopoglou

Professor

Thesis Committee Member

A Markov-Based Decision Support Model for Tax Evasion in Greece

by

Goumagias D. Nikolaos

Submitted to the Department of Applied Informatics
on October 21, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis proposes a Markov-based decision support model which captures the behaviour of a typical firm in Greece, vis-a-vis tax evasion. To maximize its long-term wealth, the firm has two options at its disposal. First, it can manipulate the percentage of its profits that it will disclose (and consequently be taxed on). Second, there is a type of optional tax amnesty that may be offered to the firm by the Greek government. This option, termed "closure", allows the firm to pay a lump-sum tax, based on its gross income or stated profits, in exchange for eliminating the possibility of an audit against the firm's past income declarations. We describe a dynamical system, aimed at predicting the actions of the average Greek firm, and at evaluating tax policies before they are implemented. The basic model evolves through several iterations, some computationally challenging, depending on the firm's attitude towards risk and the availability of closure. The model proposed in this thesis represents an innovative step towards bridging the gap between classical macroeconomic and game-theoretic approaches on the subject. It takes into consideration the sequential nature of the firm's decisions regarding tax evasion, together with its attitude towards risk, and allows us to: i) analyze the effectiveness of the closure option, ii) show that in the current environment a rational enterprise has no incentive to disclose its profits, iii) identify "virtuous" combinations of tax parameters which lead to full disclosure of profits and iv) estimate a firm's risk-aversion. Our analysis can be used to evaluate the effectiveness of various taxation schemes, potentially benefiting both firms and government.

Thesis Supervisor: Dr. Dimitrios Hristu-Varsakelis
Title: Associate Professor

Acknowledgments

I would like to express my utmost gratitude to my supervisor, professor Dimitrios Hristu-Varsakelis for his continuous counseling and guidance, from initial advice of conceptual inception to ongoing encouragement to this day.

A special thank goes to my colleges and academic staff of the University of Macedonia for their support, advice and cooperation through all of these years of research.

I would also like to thank my parents, Athena and Dimosthenis, and my brother, Hristos, for their undivided support and encouragement without whom I would have been unable to complete it. Last but not least a special thank goes to my friend, Haralampos Triandafillidis, whose belief in me and encouragement was more than valuable during this journey.

Finally, I would like to thank my friends whose support through the last years was more than valuable.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Faced with perhaps its most serious debt crisis in modern history, Greece is currently implementing a series of austerity measures and reforms. One of the central components prescribed in the "rescue package" overseen by the EU and the IMF calls for a dramatic increase in tax revenues and the minimization of tax evasion, the latter being one of the country's most serious and persistent problems. The basic components of the current tax system for incorporated entities are a flat tax rate on profits, random audits for identifying tax evaders, and monetary penalties for under-reporting income. Currently the tax rate is set to 26% on profits and 10% on dividends. These rates vary, possibly quite frequently, depending on current legislation.

Typically, the government does not have adequate information on a firm's profits, which may be manipulated through a variety of methods. Two of the most often used include i) manipulation of financial statements to under-report income, and ii) invoices (often issued by another, usually short-lived firm) that document supposed expenses and are used to offset profits. Penalties for tax evasion depend on the amount of unreported income, and the time elapsed since the offense took place. Specifically, a firm found to have concealed income must pay any tax originally due on that income, plus a 2% monthly penalty on that tax. Thus, "older" tax evasion decisions are potentially more costly than recent ones. The total penalty amount is subject to a 2/3 "discount" for prompt settlement once the evasion is discovered, and is capped at twice the original tax owed. The firm's[1] true profit may be revealed by performing an audit. Because of scarce resources, Greece can only audit a limited number of cases each year, estimated at approximately 5%. Thus, in an effort to collect revenue and promote full disclosure, the government retains the right to audit businesses "retroactively" for up to five years in the past. Any tax evasion activity beyond that horizon is essentially capitalized by the firm. Because of this, the audit probability is comparatively higher for firms which have not been audited for the last four years.

A somewhat unusual feature of the Greek tax system is that the government periodically offers businesses the option to "close" past tax declarations to any audits, for a fee which is to be paid for each tax year a business would like to be exempt from possible audits. Because the statute of limitations on tax declarations is usually five years, the government has in the past offered this option in roughly five-year intervals. This "closure option" can be viewed as a kind of middle ground: it may allow an entity to cover-up past transgressions, at some cost, but it also provides the government with some tax revenue (if a sufficient number of businesses opt to use it), although that revenue may be less than what is properly owed. For our purposes, the option works roughly as follows. The government declares that closure will be available in the current fiscal year and will cover a given number of years in the past. The firm files this year's tax statement as usual, and declares some nominal profit. It pays any tax owed on that profit, plus a fixed amount for each fiscal year it wants to cover under the closure option. In exchange for that additional amount, the government agrees to consider the past tax statement(s) as truthful and never audit them[2]. If a firm does not avail itself of the closure option, it may find itself with a higher probability of being audited, as many of its peers effectively "remove" themselves from the audit pool. The most recent closures have been occurring roughly every 4-5 years during

[1]Here, we will use the word "firm" to mean incorporated entities in Greece, or elsewhere, that operate following the international accounting standards, commonly known in Europe as S.A. (Société Anonyme).

[2]There are certain safeguards in place to ensure that, for example, a firm must declare some minimum profit if it wants to "close", or must calculate its closure cost as a fraction of gross sales instead of net revenues. The precise amount is determined by the government each time the option is offered.

1998-2006 (e.g., [1] [2]). It is clear that Greece considers closure to be an integral part of the tax system, and the latest took place in 2010. Currently, the effectiveness of the closure option is unclear, and there is a widespread feeling that the auditing system and level of penalties are not adequate to prevent tax evasion.

The purpose of this thesis is to describe a decision support model which incorporates the salient features of the Greek tax system currently in place, as it pertains to firms. Our model, in the form of a dynamical system with inputs, is mainly designed to explore an enterprise's propensity to "cheat" under various scenarios, as the latter seeks to maximize the present value of its long-term expected profits. The main parameters of the model will be the tax rate, the probability of the firm being audited, the probability of the government offering the closure option in any given year, the cost of the option to the firm, and the penalty for unreported profits.

We will describe the firm's evolution within the tax system by means of a Markov decision process. Our main goal is to compute the optimal behaviour expected of a "typical" rational enterprise and identify the states in which tax evasion is an optimal policy for the firm. This will allow us to i) "chart" in our parameter space the region(s) which lead to honest behaviour (i.e., full disclosure of profits) and maximization of government revenues, and ii) evaluate the closure option as a revenue-collecting measure and determine whether it promotes or deters tax evasion. We are also interested in knowing the extent to which a firm's decisions in the current year depend on past decisions, e.g., its tax evasion policy within the last five years. We expect that for certain parameter values, which we would like to compute, the firm's optimal decisions will be independent of past behaviour.

We will begin our analysis by assuming a risk-neutral firm, but will subsequently lift this assumption by re-casting the firm's situation regarding tax evasion as a simplified portfolio optimization problem, in the absence of the closure option. This will lead us to i) a map of the optimal policy of the firm in the audit probability - tax penalty space and an estimate of the maximum utility under different risk-aversion coefficient and ii) an implicit estimate of the true level of the average firm's risk-aversion. Finally, we will consider a more complete model, which includes risk-aversion on behalf of the firm as well as the closure option. The complexity of that model will require that it be solved using an approximation method, in our case Approximate Dynamic Programming (ADP), implemented using a combination of techniques from reinforcement learning and function approximation. The proposed model may be useful as a tool for gauging the effectiveness of the current system, and for guiding future tax policy. Besides evaluating tax policies before they are implemented, our model can help identify those which are both fiscally responsible and business-friendly, in the sense that they are harsh enough to make tax evasion unprofitable, but no harsher.

The remainder of this thesis is structured as follows. Chapter 2 contains a detailed literature review. The basic model is described is Chapter 3; numerical results and discussion regarding the firm's optimal behavior are provided in Chapter 4. In Chapter 5 we revisit the problem, but in this case as a portfolio optimization problem excluding the closure option. In Chapter 6, we reintroduce the closure option in the model described in Chapter 3, this time assuming that the firm is risk-averse and acts according to a (nonlinear) constant relative risk aversion utility function. We solve this more complex model via Approximate Dynamic Programming, and discuss the results in Chapter 7.

Chapter 2

Literature Review

In this Chapter we review past work concerning tax optimization and tax compliance. Due to the nature of the problem, approaches vary widely depending on the researchers' area of expertise e.g public finance, microeconomics, macroeconomics and applied optimization. Furthermore, tax systems[1] are complicated and vary significantly among countries. However, there are some core similarities which allow us to discuss some salient features. Taxation is usually applied via two approaches: direct taxation, which is assessed directly on a taxpayer's declared income, and indirect taxation (i.e., VAT and commodity taxes) on consumption of goods and services. It is unanimously accepted by researchers that direct taxation is more just than the indirect variety because it takes into consideration the taxpayers' income. On the other hand, indirect taxation can only be avoided by a taxpayer through a restriction of his or her consumption. Here we will be mainly interested in the problem of tax evasion as it pertains to *direct* taxation. It should be noted that by "tax evasion" we mean the illegal minimization of tax liabilities by concealing taxable income or otherwise providing false information to the authorities. This is to be distinguished from tax avoidance, which is the minimization of tax liabilities through legal means but with full disclosure of income. Tax evasion can be achieved by concealing a percentage of the true income (although the means of doing so may vary considerably).

[1] According to the Oxford Dictionary of Accounting the term "tax system" refers to the means by which taxes are raised and collected in accordance with the relevant legislation

with respect to indirect taxation, the consumer can evade the VAT, for example, by negotiating with the a seller or provider of services to "conceal" the transaction by not issuing any invoice or receipt.

Arguably, the complexity of a tax system is such that a description of its features in one unified model is far from trivial, if not impossible. In this thesis, we are interested in capturing the salient features of direct taxation, and in particular income taxation. As a result, the following review will be mainly focused on the research regarding optimization of income taxation. Welfare theory and optimal taxation theory assume no tax evasion and consequently will not be dealt with here. The structure of the literature review that follows is based on [3]. For the sake of clarity and in order to avoid any notation overlap, we will use that work's mathematical notation when appropriate.

2.1 Theoretical Models

According to [4], the term "optimal taxation" is used in the literature to signify different things, including the minimization of resources used to collect taxes, issues of "fairness and justice" with respect to taxation and economic efficiency in the context of tax collection. Early approaches on the subject began in the 19^{th} century. However, interest concerning optimal taxation was minimal until 1956 when [5] presented a unified view of the "theory of second best" in the case where "Pareto" conditions are not met, and applied that view to the so-called "theory of tariffs". In that work the authors employed a three-commodity model (two imported and one domestic) and estimated the optimal level of tariff imposed by the state, and its effect on total welfare. Early work on direct tax evasion begins with [6], who introduced a theoretical macroeconomic equilibrium model in the form of an optimal portfolio allocation problem. In [6], the authors argued that it was unrealistic to assume no tax evasion was taking place, as was done up to that point. The general setting was one where a risk-averse taxpayer (also termed "agent") decides the allocation of his or her wealth between a risk-free asset (which implies disclosing his or her true income) and a risky asset of “partial declaration” of his or her income, facing a potential penalty in case of an audit. The taxpayer’s goal is to maximize the expected utility of his or her portfolio:

$$E(U(X)) = (1 - p)U(W - \theta X) + pU(W - \theta X - \pi(W - X)), \tag{2.1}$$

where:

- $E(U(\cdot))$: expected utility
- p: audit probability
- W: taxpayer’s true income
- θ: tax rate
- X: declared income

This early model provided some important insights into tax evasion. In [6], the authors were mainly interested in the effects of the tax rate on the level of tax evasion. Those effects are twofold. First, an increased tax rate lowers the after-tax income of the agent, which in turn decreases his or her absolute risk aversion, making the agent less inclined to under-report. On the other hand, by raising the tax rate and keeping the tax penalty fixed, the potential payoff of tax evasion increases, and consequently tax evasion is encouraged, via a so-called “substitution effect”. The question of which of the two contradictory effects persists depends on the relative risk aversion of the individual. In [6], the authors proposed the introduction of labour supply into the model so as to identify connections between incentives to evade taxes and work supply, as well as between savings and portfolio decisions. They also encouraged future empirical research to assess the relative efficiency of tax deterrents such as tax rates, audit probabilities, and tax penalties.

According to [6], it is easy for the government to assure tax compliance for risk averse individuals by setting the tax penalty sufficiently high. However, governments avoid setting penalties too high, to prevent potential bankruptcies and to preserve the private sectors' welfare in general. Studies that examine the optimal levels of penalties include [7], [8], [9], [10] and [11], among others. We will discuss papers that examine the effect of this and other parameters on tax evasion later in this chapter.

Several studies followed [6], including [12], who examined two contradictory effects of tax rates, namely the effect on income, and the substitution effect as described earlier. That work argued that when the tax penalty is proportional to the tax evaded, taxpayers tend to comply. The taxpayer again decides the optimal allocation his of her gross income between a risky asset (the undeclared income) and a risk-free asset (the disclosed income) and it was predicted that tax evasion decreases as the tax rate increases, based on the fact that the expected value of consumption decreases, and consequently, the variance of the disclosed income decreases along with tax evasion.

Starting with the models developed by [12], [13] incorporated labour supply and examined the effects of the probability of detection assuming the following relationship:

$$\frac{d\rho^*}{d\pi} = \left(\frac{d\rho^*}{d\pi}\right)_{\bar{Y}L} + \frac{wdH^*}{d\pi}\left(\frac{d\rho^*}{d\bar{Y}}\right)_L + \frac{dL^*}{d\pi}\left(\frac{d\rho^*}{dL}\right)_{\bar{Y}}, \tag{2.2}$$

where

- ρ^*: the optimal level of tax evasion,
- π: the probability of detection,
- Y: net income,
- L: leisure,
- $\bar{Y}$: gross income,
- H: hours worked,

The first term in Eq. 2.2 is called the "portfolio effect". If the taxpayer's utility is an increasing function of net income and leisure, then the effect of an increase in probability of detection is negative under conventional preference assumptions. The second term is called the "income effect" and is meant to explain the effect of the detection probability through income. Finally, the third term is "the leisure effect" which describes the effect of the probability of detection on tax evasion through leisure consumption. It is argued, however, that the incorporation of labour effect in the model does not have a significant impact on tax evasion.

In [14], the author examined the robustness of [6] with respect to the relationship between tax evasion and tax rates, using three modifications of the original model. The modifications involved the introduction of a non-linear tax scheme, a more sophisticated tax penalty equation, and the incorporation of labour supply. The utility that the taxpayer enjoys is taken to be a function of both income (y) and work hours (h):

$$U = U(Y, h). \tag{2.3}$$

The taxpayer's income, when under-reporting, is given by:

$$Y^0 = Z(h) - S + \tau y^\sigma, \tag{2.4}$$

where:

- $Z(h)$: the taxpayer's true income as a function of the working hours,
- S: welfare payments to individuals who would otherwise have no income,
- y: the income the taxpayer reports to the authorities,
- $0 < \tau \leq (S + y)/y^\sigma$,
- $0 < \sigma \leq 1$.

The parameters σ and τ govern the relationship between changes in reported income and changes in tax payments. On the other hand, the taxpayer's income, when caught "cheating", is given by:

$$Y^c = Z(h) + S - \tau y^\sigma - \begin{cases} \tau(Z^\sigma - y^\sigma) - F \\ \lambda\tau(Z^\sigma - y^\sigma), \end{cases} \tag{2.5}$$

where:

- $\tau(Z^\sigma - y^\sigma)$: taxes involuntarily paid by the taxpayer,
- F: the fine imposed, and
- λ: a multiplier that penalizes the taxpayer for under-reporting.

The goal of the taxpayer is to maximize the expected utility

$$E(U(Y,h)) = \pi U(Y^c, h) + (1-\pi)U(Y^0, h). \tag{2.6}$$

In [14] it is argued that under the above model, increased tax rates make taxpayers more honest. However this is yet to be confirmed through empirical evidence.

Further work concerning the incorporation of labour supply into a tax scheme includes [15], whose major concern is the optimal allocation of working hours among several risk-free and risky activities that are about to be taxed by the government. Although [15] does not focus on tax optimization or tax evasion per se, it does extend the classical labour supply model. An individual has the option of engaging in $m+2$ activities. Activity 0 is "regular work" with a known wage, w_0, that depends on the time-units, h_0, that the individual invested in that work. Activities $1, \ldots, m$ are characterized by random rewards w_j that depend on the hours h_j invested in these activities with $j \in 1, 2, \ldots, m$ and finally ℓ, are the hours spent on leisure time.
Consequently the constraints are:

$$h_0, h_1, h_2, \ldots, h_m, \ell \geq 0 \tag{2.7}$$

and

$$\sum_{j=0}^{m} h_j + \ell = 1. \tag{2.8}$$

The individual's consumption c is a function of the tax, v_j, that the government imposes on the income of each activity j, and a lump-sum grant, B, that government provides. The tax parameters are known *a priori* and the consumption is given by

$$c = B + \sum_{j=0}^{m} v_j w_j h_j, \tag{2.9}$$

while the net tax paid is

$$T = -B + \sum_{j=0}^{m} (1 - v_j) w_j h_j. \tag{2.10}$$

The individual's utility to be maximized, is a function of both consumption and leisure time:

$$E(U(c, h)). \tag{2.11}$$

In [15] the author argues that, as expected, labour supply decreases as the marginal tax rate increases. By reducing the tax imposed on risky activities, the amount of "safe" work is decreased as well. The term safe work is used to describe the official (legal) registered work, in contrast to the illegal one, cloaked as leisure. If the penalty for tax evasion is increased, and there is a decreasing absolute risk aversion, there is a switch from illegal to legal work. A relative increase in capital gains increases regular employment. Consequently, safe work always increases, at the expense of the risky one, when taxes on risky income become more progressive.

The work in [16] focused on determining the optimal level of tax rates and audit probability after incorporating labour supply into [6]. The authors derived formulas for the relationship between tax penalties and audit probabilities. Taxpayers were divided into two categories: the non-evaders, whose income is known to the government, and the evaders, whose true income is to be determined through an audit. The preference ordering of the non-evaders' group is given by the following utility function:

$$U^n = U(C^n, L^n), \tag{2.12}$$

where:

- n: the individual n of the N group of non-evaders,
- C: the consumption function, and
- L: labour supply.

The budget constraint of the non-evaders, after incorporating the wage, w, a marginal tax rate of t, and a lump sum α, is given by

$$C^n = w^n L^n (1 - t) + \alpha. \tag{2.13}$$

The non-evaders' aim is to maximize Eq. 2.12 subject to Eq. 2.13. On the other hand, the preferences of the evaders and their utility function is

$$U^e = U(C^e, L^e + E), \tag{2.14}$$

Labour supply consists of the hours worked, L^e, and and hours, L, of hidden work. The registered hours are known to the government and taxed as usual. The hidden hours can only become known to the government via investigation. Consequently, the evader can either be caught, with his or her consumption being C_1^e, or be unaudited, with his or her consumption being C_2^e. The budget constraints of the evaders' group are

$$C_1^e = w^e L^e (1 - t) + a + w^e E, \tag{2.15}$$

and

$$C_2^e = w^e L^e (1 - t) + a + b + w^e E(1 - \theta). \tag{2.16}$$

Based on the above, assuming that the taxpayer is utility maximizer and estimates the audit probability, p, the expected utility is

$$\bar{U}^e = (1-p)U(C_1^e, L^e + E) + pU(C_2^e, L^e + E) \tag{2.17}$$

In [16] it is argued that the utilitarian approach to determining an individual's behaviour concerning tax evasion is not "sufficient", because it focuses solely on the economic system without taking into consideration the underlying social system. This should lead policy-makers to more stern decisions concerning tax-evaders than the utilitarian approach implies. But the author also concludes that tax penalties and rates should not exceed a certain boundary in order to avoid extreme solutions that jeopardize the welfare of individuals.

The theoretical approaches to tax evasion that we have presented up to this point have led to a series of critical responses such as [17] and [18] who argued that, based on empirical evidence, a portion of taxpayers never engage in tax evasion. Tax evasion seems to increase with tax rates and tax evasion decisions are temporally interdependent. These two works introduced pecuniary costs to attempt a connection between theory and evidence. Taxpayers were separated into different groups according to their morality, and the empirical evidence of the positive relationship between tax rates and number of tax evaders was explained by introducing an endogenous reputation cost.

The model from [17] can be summarized in the following maximization problem. The preference of a taxpayer is a function of consumption, C, and the extent of his or her honest behaviour, H. The honest behaviour, H, is measured as minus the income concealed from the government, E.

$$u\{C, H\} \tag{2.18}$$

Consumption is given by:

$$C = Y(1-\tau) + x\tau E, \tag{2.19}$$

where:

- Y: the taxable fixed income,
- τ : the uniform tax rate,
- E: income concealed from tax authorities,
- s: tax penalty,
- $x \in \{1, -s\}$ with probability p and $1 - p$ respectively, and
- p: audit probability.

Based on Eq. 2.19, the equation 2.18 can be rewritten as

$$UC - vE. \tag{2.20}$$

The term vE represents the taxpayer's moral cost of tax evasion. The distribution of v among taxpayers is given by $F(v)$. The taxpayer's goal is to maximize his or her utility function as in Eq. 2.18.

Note that in Eq. 2.19 the effect of x (the random audit) is to make consumption a random variable. The above model can effectively explain why there is always a group of taxpayers that refuse to engage in a favourable gamble by evading taxes, and how a positive relationship exists between tax rates and tax evasion. Effectively, [17] endogenised the dis-utility of engaging in dishonest behaviour in order to explain the time-interdependence of tax evasion decisions, thus providing a reasoning behind the positive relationship between tax rates and tax evasion in the long run.

Further research that incorporates "morality" into a theoretical equilibrium model includes [19], where the morality of a "tax inspector" was introduced. The authors pointed out that, although there has been considerable research on the "immorality" of taxpayers, "tax inspectors" had been assumed to be moral up to that point. They identified three manifestations of this phenomenon: extortion, remuneration and distributional aspects. Their model aimed to maximize the government's wealth while minimizing tax misreporting. They concluded that poor taxpayers are more vulnerable to the immorality of tax inspectors and as a result this has an impact on the distributional effects of the tax system. Moreover, the introduction of premiums paid to the tax auditors (as an incentive to behave "morally") may be costly, and to some extent it is optimal to let a certain level of tax evasion to persist. Although there is considerable literature on corruption and how it affects the government's revenues, we will not include it in this review because it refers mainly to environmental regulation rather than tax optimization. Representative papers include [20], [21], [22] and [23].

Among the drawbacks of the "classical" model presented in [6] and its "descendants", is that model's static form fails to take into account the repetitive nature of income statements which create tax compliance motivations of a cumulative nature. There are only few works, including [24], [25], [26] and [27], that examine the *dynamics* of tax compliance and it is precisely this gap that this thesis aims to fill.

2.2 Game Theoretic and Principal - Agent Approaches

One characteristic of the models discussed up to this point is that the audit probability is fixed for all time and for all taxpayers. This assumption may be practical but not realistic. In an effort to lift that assumption, two approaches were proposed depending on the tax authority's ability to commit to an "audit rule" before or after income declaration. Here, the term audit rule refers to the probability distribution of auditing a taxpayer, or group of taxpayers based on their income. We proceed to review each approach.

2.2.1 Principal-agent framework

In works belonging to the so-called principal-agent approach, the government or tax agency, determines the audit rule *a priori* and commits to it after income declaration on behalf of the agents is complete. The agency dilemma describes the problem of providing incentives for one party (agent or taxpayer) to act in the best interests of the other (principal or the government). One problem this framework aims to solve is the existence of asymmetric information between parties. Usually, the agent is better informed than the principal. The principal cannot know whether or not the agent acts in the principal's best interest due to conflicts of interest and moral hazard (sometimes called "agency costs") resulting in lower overall welfare.

The pioneer behind the introduction of the principal-agent approach in the context of tax evasion was [28]. In that work, the income of the taxpayer is an independent random variable I. The distribution is given by $G(I)$, where $g(I) = G'(I) > 0$ for all $I \in [0, \infty)$. The government observes I with a cost $c > 0$. The taxpayer reports to the government his or her income, x, each year. If the government does not audit the taxpayer, the tax paid to the government is given by $t(x)$. If the tapayer is audited, the government's revenues are $f(x, I)$. The expected net income of the taxpayer, $r(x, l)$, is given by:

$$r(x, I) = [1 - p(x)][I - t(x)] + p(x)[I - f(x, l)]. \tag{2.21}$$

The taxpayer's goal is to maximize $r(x, I)$ given $t(x)$, $f(x, I)$, and the audit probability $p(x)$. Assuming that the optimal policy of the agent is $\phi(I)$, the authors allowed the principal to choose an alternative policy in $\Phi(I)$, based on the assumption that if the agent is indifferent among a set of policies, he or she chooses the one most favourable to the principal. In this setting, the principal's expected net revenues are

$$R(x, I) = [1 - p(x)]t(x) + p(x)[f(x, l) - c]. \tag{2.22}$$

The expectation is based on the probability distribution of $G(I)$. The principal's goal is to choose the optimal policy regarding audit probability, tax rate and tax penalty that maximizes the expected net revenues $R(x, I)$. In this context, given the set of all feasible policies, P, available to the principal, there is another policy that forces taxpayers to behave honestly and consequently the authors focused on this set of feasible policies.

The principal-agent framework developed in [28], was generalized further in [29]. Consider a set of agents (taxpayers) with average income i. Their income is not known to the government and it is a random variable in the interval $[l, h]$. The distribution of the taxpayers with income lower than the average is given by $F(i)$ with $F'(i) > 0$. The principal's goal is to determine an audit rule that maximizes tax revenues. The tax policy of the principal consists of setting the audit probability (p), tax penalty (π) and budget constraints (B). The true income of the taxpayer can be determined via an audit at a cost of $c > 0$. The function that determines the tax based on the income reported is given by $t(\cdot)$. The taxpayers' expected revenues are

$$t(r) + (1 + \pi)[t(i) - t(r)], \tag{2.23}$$

with no reward for over-reporting. The reported income is r. The principal's goal is to find the proper taxation mix that maximizes the after-tax utility distribution of taxpayers. The utility function, $U(A, R) = A + \phi(R)$, depends on the taxpayers' after-tax income, A, and the monetary value, $\phi(\cdot)$, of public expenditures, R. If v is the after-tax utility distribution, the government aims to maximize:

$$\int W[U(A, R)]dv, \tag{2.24}$$

where W is a twice differentiable function. More detailed descriptions of the principal-agent approach can be found in [30], [31] and [32].

2.2.2 Game-theoretic framework

In game-theoretic approaches to tax modeling, the government decides the audit rule after the taxpayers' income declarations have been completed. Each taxpayer has true income y and should pay a tax of $t(y)$. The income distribution of the taxpayers lies between $\underline{y}$ and $\overline{y}$ according to a density function $f(y)$. The government expects the taxpayer to declare his or her income and pay the corresponding tax. The declared income, x, is less than or equal to the true income, y. Expecting that, the tax agent decides to audit a group of taxpayers. Only then is the real income, y, revealed to the agent, the additional tax, $t(y) - t(x)$, is paid and a penalty, $\theta[t(y) - t(x)]$, with $\theta(\cdot)[q] > 0$ is imposed. The taxpayer's expected utility is given by:

$$EU(x) = (1 - p(x))U(y - y(x)) + p(x)U(y - t(y) - \theta(t(y) - t(x))). \tag{2.25}$$

In Eq. 2.25 the audit probability, $p(\cdot)$, is a function of the declared income. At this point the tax agent applies an audit rule, which is the same as "selecting $p(\cdot)$", for the purpose of maximizing the expected net revenues of the government, and in some cases, social welfare. The following articles are grouped according to the government's propensity to commit to an audit rule before or after income declaration, as discussed earlier.

2.2.3 The audit rule is determined before income declaration

In the simplest setting, the government's audit policy is a simple cut-off rule [3]. The tax authority determines an income threshold, ω, below which the taxpayer faces an audit probability p, uniformly. The tax authority has to determine p so that all taxpayers reporting below ω declare their true income, while the number of audits is subject to budget constraints. On the other hand, taxpayers with income above ω are not audited based on the notion that lower-income taxpayers are more prone to tax evasion. This argument is supported by [10], which will be further discussed later in this section.

One of the earliest such models was developed in [28], who relied on the concept of sequential equilibrium as expounded in [33], [34]. More specifically, the authors formulated a game theoretic equilibrium model with incomplete information (i.e., the taxpayer's true income). Assuming a limited audit budget, a cut-off rule was introduced, under which taxpayers should be audited with a probability p if they declared income below the threshold, while the entire audit budget is to be spent. Although this setting is not very realistic, it comes about as a result of the fact that taxpayers are treated as a single group and not separated into various groups according to their income. The condition of optimality relies on the assumed risk neutrality of taxpayers and the linearity of tax rates and tax penalties. The authors examined the robustness of their results and found that equilibrium does not hold when risk aversion is introduced.

The cut-off rule proposed in [28] was further generalized in [29], allowing the tax authorities to choose both tax rates and penalties before income declaration. A hierarchical model of tax compliance was developed, after taking into consideration the potential conflict between the government and the tax collection agency's incentives. The tax agency usually aims to maximize collected taxes but the government's goal could be the maximization of social welfare. The basic assumptions (risk-neutral taxpayers and fixed income distribution) remained the same as in [28]. When the taxpayers' income distribution is known to the government, the optimal audit policy requires taxpayers to be divided among three classes based on declared income. The findings presented in [28] regarding the greater audit probability required for lower income individuals are further supported by [29], where it is also argued that an increased audit budget yields increased government revenues. In [29] and [35] it is also argued that higher-income classes are more prone to evasion, and consequently effective tax rates are more regressive than nominal ones. In [36] labour supply was introduced into the model of [29], and the optimal audit policies were estimated when a taxpayer could not tax-evade but could reduce labour supply. Generally, there is a consensus among researchers that the reporting process is sensitive to the probability of audit. At the same time, the introduction of a budget constraint in the model developed in [29] could make tax evasion insensitive to the tax-rates and other enforcement parameters [37].

In [10] the effects of taxpayer risk aversion are explored. In that work, the taxpayers' income is divided into discrete classes, and the qualitative characteristics of an optimal audit policy are examined. The authors move one step forward by introducing rewards for taxpayers that truthfully disclose their income. Finally, [10] supports the argument made in [28] that if there are only two classes of income, "high" and "low", then the high income class should not be audited. For more than two classes, the audit probability may not be monotone in income, in contrast with [29], who discussed what happens in the case of risk neutrality.

The authors of [38], [39] were the first to assume a risk-neutral tax authority and risk-averse taxpayer based on the model from [29]. The taxpayer declares his or her income, and the tax authorities follow by determining which taxpayers will be audited, at a given cost per audit. These two aforementioned works aimed to improve the incentives for authorities to correctly verify the taxpayers' income (avoiding, for example, bribery). However, the improvement of taxpayers incentives to declare their true income is not examined.

The work in [40] examined the dual effect of information on the audit probability distribution, based mainly on [28], [37] and [41]. That work considered a zero-sum game with two players: the taxpayer, who aims to minimize the tax due by declaring a lower income, and the tax agency, which aims to maximize tax revenues. To achieve that, the agency pays a cost to obtain information such as the taxpayer's income from salary and investments, the distribution of recorded income losses and, finally, the ability of the taxpayer to deduct losses from his or her income. Using this information to decide who to audit, has four effects: i) tax evasion decreases, ii) there is no effect on the expected level of gross government revenues, iii) expected audit costs are also increased, and iv) the optimal level of the amount of money "invested" in obtaining the above information, does not vary by tax rate, penalty rates or audit costs. In [40] it is also argued that equality and enforcement cannot always coexist. Consequently, the government needs to accept a certain level of tax evasion in order to maximize the collected tax revenues. It is also suggested that risk-neutrality should be dropped in order for more realistic models to be developed.

A further contribution regarding the construction of an optimal audit policy can be found in [42]. That work derived an optimal audit rule that maximizes social welfare via costly audits while ensuring tax compliance. It also distinguished between tax authorities and government, as in [29]. The authors argued that maximizing social welfare involves tax equality, while maximizing tax revenue's net of costs is based on tax compliance and on whether the tax authority can commit to an audit rule, as in [28]. The government, on the other hand, corrects the tax authority's behaviour to ensure tax equity.

In the quest for a "general" optimal audit strategy, [43] argues that such a strategy cannot consist of a collection of optimal strategies each optimized for a different income class, because of the interactions that exist among the classes. That work's authors suggested that maximizing gross tax revenues is more appropriate than net tax revenues. However, for simplicity reasons, [43] does not include into the analysis any incentive scheme for tax authorities. A more thorough analysis of the differentiation of objectives between government and tax authorities is presented in [29] which provides a discussion on the fact that the government should hand a smaller budget to the tax authority than the amount regarded as optimal (based on maximizing net tax revenues) because otherwise social welfare may be reduced.

2.2.4 The audit rule is determined after income declaration

There are several works that propose general game-theoretic and principal-agent models, where the tax authorities do not commit to an audit rule until after income declaration. As a result, in [28] and [44] the problem of tax evasion takes on the form of a sequential equilibrium game. In this context, all players are willing to "cheat" for the "right price". This approach is presented in a comprehensive manner in [3]. Assuming that tax penalties and tax rates are linear, the true income, $y(x)$, is thought to be a function of reported income, x. In equilibrium, it is assumed that all taxpayers under-report by the same amount:

$$\frac{\lambda c}{(1+\theta)t}, \tag{2.26}$$

where:

- c: the audit cost,
- λ: the shadow price of the audit budget constraint,
- θ: the tax penalty, and
- t: the tax rate.

The true income, y, lies between an upper bound, $\overline{y}$, and lower bound, $\underline{y}$, and the reported income, x, between $\overline{x}$ and $\underline{x}$ respectively. The relationship between the lower bounds are labeled by

$$\underline{x} = \underline{y} - \frac{\lambda c}{(1+\theta)t} \tag{2.27}$$

while for the upper bounds:

$$\overline{x} = \overline{y} - \frac{\lambda c}{(1+\theta)t}. \tag{2.28}$$

The audit probability, $p(x)$, is taken to be zero for income that lies in the interval $[\overline{x}, \overline{y}]$, and strictly positive for values of reported income in $[\underline{x}, \overline{y}]$. For any reported income x, the tax authority can estimate the true income from

$$x + \frac{\lambda c}{(1+\theta)t}. \tag{2.29}$$

In [3], the author argues that under the setting just described, tax authorities are indifferent regarding whether or not to audit the declared income because the cost of the audit, λc, is equal to the earned revenue $(y-x)(1+\theta)t$ after the audit. The audit probability, $p(x)$, has to be determined in a way that provides the taxpayers with the incentive to follow the above strategy, assuming that the taxpayers can forecast the audit probability correctly, depending on their income value, and given that the tax penalty is linear. More specifically, the audit probability is expected to be zero when

$$x = y - \frac{\lambda c}{(1+\theta)t}. \tag{2.30}$$

On the other hand, the value of λ should be determined so as to satisfy the budget constraints set by the tax authority. In other words, all taxpayers that fulfill the criteria of the audit rule must be audited.

In [45], in the case of risk-averse taxpayers, it is suggested that the level of tax evasion varies among taxpayers with different levels of true income. This issue of separating taxpayers into different "classes", is addressed in [44], where a new sequential equilibrium model was introduced, classifying taxpayers into three groups, according to their probability of being audited. The author determines a series of assumptions that yield an equilibrium between each group of taxpayers and the tax agency. One of the assumptions, however, is that tax agency "threatens" to audit certain taxpayers, which is criticized in [3], along with the purported equilibrium.

There are a number of critical assumptions incorporated into the principal-agent and game-theoretic approaches. The most obvious drawback is that they are poor descriptions of the real world [3]. This criticism is based on the assumptions themselves, as well as on the fact that the results obtained are not compatible with empirical evidence. According to [3], the cut-off rule used in the principal-agent approach implies that all taxpayers that are audited will be found to report honestly, because they are aware of the audit probability they face. Moreover, within an audit class, all taxpayers with higher true income are found to report at or near the cut-off threshold. However, it is difficult to validate these conclusions through empirical research.

On the other hand, game-theoretic approaches produce more realistic predictions as far as taxpayers' behaviour is concerned. The sequential equilibrium approach and, more specifically, the partially separating equilibrium approach, as in [3], imply that not all taxpayers follow the same tax evasion strategy. Instead, a group of taxpayers remains honest no matter what. For this reason, in [46], the authors extended the model described in [3]. In [46], there exists in every income level, y, a fraction of taxpayers, Q, that always declare their true income. On the other hand, a fraction, $1-Q$, always try to evade taxation. Going further, [45] presented a more detailed effect of this separation and produced more realistic results, according to empirical evidence but at the same time complicated the analysis significantly. That work suggested that the two factors that affect the government's revenues are the magnitude of under-reporting and the fraction of taxpayers that under-report. It also overcame many drawbacks of the first game-theoretic model developed in [46]. The major difference between [45] and [46] is that, in the latter work, a group of taxpayers always reports its true income.

2.2.5 Comparison between Principal-agent and Game-theoretic approaches

As far as the audit policy is concerned, the principal-agent method produces more realistic results when compared with game-theoretic approaches. When the audit policy takes the form of a cut-off strategy, reports of lower income are audited with a positive probability, while higher reported incomes are not audited at all. According to [3], many tax agencies, including the IRS, employ this strategy. Game-theoretic approaches produce less realistic results. At equilibrium, the tax authorities are indifferent towards auditing in the reporting range. As a result, they do not employ a cut-off rule and the equilibrium models predict that higher income classes will not be audited as well. The probability of audit decreases monotonically with declared income.

In the principal-agent framework, the tax authorities determine the true income of taxpayers reporting income lower than the threshold but not the income of those reporting higher. On the other hand, in the game-theoretic framework, the authorities cannot determine the true income of any taxpayer without performing an audit, which agrees with empirical evidence. Of course, a critical evaluation of the two approaches should not only be based on empirical evidence but on their reasoning as well. Both assume that taxpayers can easily deduce the taxation scheme (audit probabilities) *a priori* and "tune" their behaviour accordingly; they are also aware of their true tax liability. See [47] for an example where those assumptions are lifted. This information asymmetry can be only partially overcome by allowing the existence of honest taxpayers who do not need to know the audit strategy. Moreover, the cost of auditing affects both taxpayers and tax authorities, both directly and indirectly (bureaucratic costs, alternative costs, professional assistance, etc.). These undocumented costs may result in different equilibria. Lastly, all the above frameworks assume that taxpayer information is disclosed only via income declaration. In reality, however, the taxpayer includes in the report several other useful pieces of information, such as sources of income. There has been limited research concerning the effect of such information on the game equilibrium, [48], [49].

A deviation from the two "traditional" approaches of principal-agent and sequential equilibrium can be found in [43]. That work assumed that the tax agency cannot commit to an audit strategy before income declaration and developed a model that bridged the gap between principal-agent and sequential-equilibrium, by allowing the tax authorities to delegate the task of tax collection to an external agent. It was shown that with a proper incentive scheme, this delegation maximizes both net revenues and social welfare. The incentive schemes can take on a salary-based or budgetary form. The authors also argue that the maximization of gross collections, instead of collections minus audit costs, may be a more appropriate objective due to the different objectives between government and the tax agency. However, other parameters should also be taken into consideration, such as salaries, ancillary costs, moral hazards, or even the probability of an untruthful audit. More specifically, there are historical examples, [50], which suggest that simple incentive schemes may easily lead to corruption, posing additional difficulties for the analysis.

2.3 Empirical Research

In the previous section we described the main classes of tax evasion models developed thus far, the assumptions used, and a critical analysis of the results. However, during the past decades, empirical research on tax evasion has increased significantly. We will present a review of that literature, examining the parameters that affect tax evasion and the effectiveness of several audit policies.

Early empirical research begins with [51] and [52], who estimated the level of tax evasion in different economic sectors. That research is beyond the scope of the current review, which is more focused on the effect of the parameters of the tax system on tax evasion. The greatest challenge empirical researchers face is how to measure tax evasion, which as an offense, and distinguish it from usual "honest" mistakes where no penalty is imposed. One of the first to examine the effect of tax rates on tax evasion empirically was [53]. In that work, the author used 1969 TCMP[2] data from the United States' IRS and estimated a tobit[3] model of evasion. The independent variables were: after-tax income, the combined state and federal marginal tax rate, and a number of dummy variables. In [53] it was argued that both marginal tax rates and the level of after-tax income had a significant impact on tax evasion. Factors that act as tax evasion deterrents included wages, interest and dividends. The rationale behind this effect is that all these three must also be reported by a firm or employer, and consequently taxpayers are reluctant to under-report.

In instances where both tax rates and income are considered dependent variables, identifying their relationship may be difficult. This problem was overcome in [55] by using a fractional detection model first proposed in [56], [57]. The model consisted of two equations: the first, a tobit form as in [53], and the second a model of the detection process. The model assumes, more realistically, that the IRS can identify only a fraction of the total tax evasion. The TCMP data used are from 1982 and 1985, and the authors aim to identify the relationship between tax evasion and income, marginal tax-rates and various socioeconomic characteristics. Contrary to [53], the authors argued that no statistically significant relationship exists between income and tax evasion and that there is a negative relationship between marginal tax rate and tax evasion. The results are consistent with [12].

To circumvent the difficulties of direct observation of tax evasion, a game-simulation is used in [58], to generate data and examine the relationship between tax evasion and tax rates. The experiment involved a questionnaire which was completed by 15 students. The findings of that work suggest that there is a positive relationship between tax rates and that tax evasion and that tax penalties are more efficient tax deterrents compared to frequent audits. However, tax evasion varies considerably among individuals. It is suggested that other factors should be taken into consideration such as the social stigma attached to tax evasion, the effect of perceptions regarding the fairness of the tax scheme, and cultural factors. Further research based on lab-generated

[2]IRS Taxpayer Compliance Measurement Program

[3]A statistical model developed in [54]

data includes [59]. The experiment described therein involved student subjects who received income and were asked to report it. Based on the collected taxes, public good became available to them. The subjects were organised into groups according to their reported income. At each round, a random income was assigned to each student and they were audited randomly to verify their income. The results remained in favour of a positive relationship between tax rates and tax evasion but, in contrast to [58], higher tax penalties are seen as ineffective tax evasion deterrents.

Further empirical research regarding marginal tax rates and tax evasion can be found in [60], where microeconomic and US Tax Treasury data were used to test for the presence of tax evasion. The authors identified a small but significant group of tax evaders which is sensitive to parameters such as probability of detection and income reported instead of the tax paid. As a result, the identified group of evaders may be underestimated. These arguments are further supported by [61].

2.3.1 Tax penalties and Audit probabilities

Other influential parameters that have been examined in the context of tax evasion models are tax penalties and audit probability. As discussed earlier, the results concerning the direct effect of these factors are ambiguous. In [58], it is argued that increased audit probabilities are far more effective than tax penalties, while [59] argued the opposite. In [62], an econometric analysis is performed for: risk-neutral vs risk-averse taxpayers. In the case of risk-neutral taxpayers, increased tax penalties and audit probabilities led to higher income declarations on behalf of the taxpayers, while the tax rate had no effect on income reporting. On the other hand, under risk-aversion, tax rates had only a minimal effect on income reporting. Taxpayers were aware of the audit probability distribution which had been set by the government regardless of the reported income. Furthermore, [62] strongly suggested the incorporation into the analysis of the role of tax practitioners vis-a-vis the information distribution among the taxpayers.

Examining the effect of audit probabilities and tax penalties is not a trivial matter because under the principal-agent and game-theoretic approaches both audit probability and tax penalty are considered to be endogenous variables. One way to deal with this difficulty is via controlling the environment artificially through computer-generated data. Examples include [63], who used 1969 TCMP data and, via an econometric model, treated the audit probability as an exogenous variable, and [64], which examined the effect of tax rates after separating audit classes using a three-digit zip code. The authors found a significant positive correlation between tax penalties, audit probabilities and tax evasion. However, in [63] the audit probability was found to be effective in one out of seven audit classes. In [65], the same 1969 TCMP data analysis is conducted through a three-equation econometric model. The findings of that work suggest that audit rates are given endogenously and that their effect is only minimal. The authors also concluded that audits are ineffective as tax-deterrents and remain statistically insignificant among audit classes. Finally, other sociodemographic variables are explanatory concerning the contradictory empirical evidence. In [66], the research is based on Swiss cantons, and the aim was to capture the effect of tax penalties and audit probabilities on tax evasion; the authors concluded that tax penalties have a positive but insignificant effect while the effect of audit probabilities was marginally more significant.

Several researchers, regard audit probabilities as a function of the budget devoted by the government. In [63], the budget available to tax authority is treated as being uncorrelated with the taxpayers' perception about audit probability in order to ensure consistency of the estimates of tax evasion. However, taxpayers' decisions are affected by their perception about the budget allocation of IRS and the IRS' budget allocation is affected by the level of tax evasion, and consequently cannot be used as a proxy for tax evasion. The proxy which is used in [65], according to [3], suffers from the same drawback. The number of information returns about other sociodemographic variables per tax return is an explanatory variable because it potentially has an impact on income reports. Finally, past audit actions on behalf of the government may affect current compliance, as in [66]. According to [3], no matter the caveats of the empirical research, a consistent qualitative picture is presented of the tax penalty's and audit probability's effectiveness as deterrents. However, the actual effect is not easily quantified.

Further research concerning tax penalties and audit probabilities includes [67], who argue that tax evasion on certain line items is more likely to be discovered. The perceived probability on behalf of the taxpayer is positively related to tax evasion and negatively related to the cost of audit. That research is based on 1982 TCMP data published by the IRS. The above negative relationship was further examined in [68], who argued that third-party information about reporting requirements may yield a reduction in cost of audit and, as a result, increase audit probability. That research was also based on 1982 and 1985 TCMP data and the findings suggest that under reporting on several line items can be reduced significantly. Finally, [69] performed a cross-country analysis, including the United States, Spain and Sweden, among others, using a model simulating a small business. Their findings suggest that the perception of a decreased audit probability reinforces tax evasion.

As discussed earlier, the taxpayers' perception of audit probabilities is based on past audits. This relationship is examined in [70], who proposed a two-equation econometric model. The first equation, of tobit form, is based on 1982 tax returns and the second equation, of torbit form, is based on 1980 data. The findings suggest that prior experience of audits does not statistically affect future behaviour, but the results were sensitive to sample selection. Further experimental approaches include [71], who conducted an experiment with forty-six undergraduate students. The students were separated into two groups with different tax schemes and they were asked to make economic decisions (acquiring information, advertisement, etc.) and state what kind of information they may require during the fiscal year. Through two mixed analyses of variance, audit probabilities and personal differences appear to be good predictors of tax evasion and support the results of [72]. The major difference between [71] and [72] lies in the fact that the participants did not know *a priori* the purpose of the experiment. On the other hand, [72] examined the assumptions that a taxpayer's perception of the audit probability is the same as that of other taxpayers, and that the perceived audit probability increases after an audit. The methodology was similar to [71]. The authors of [72] failed, however, to spot any connection between the behaviour of a single taxpayer and information about other taxpayers' behaviour. Their findings also suggest that whenever an audit occurs, taxpayers tend to overestimate audit probabilities and that tax evasion is reduced.

An explanation of the limited effectiveness of audits as tax deterrents is provided in [3]. Taxpayers tend to over-estimate the possible outcome of an audit, or an audit may fail to disclose the full extent of tax evasion and consequently impose decreased penalties. Alternatively, taxpayers may react against an audit with future tax evasion in order to retrieve any "losses".

2.3.2 The role of taxation schemes

A significant factor that determines the audit probability distribution is the structure of the audit rule. There is a growing literature concerning the development and analysis of different taxation schemes. Early work begins with [73], which developed a model of tax-return audits by the IRS analyzing 1969 data grouped by three-digit area codes. The data were aggregated into seven classes of taxpayers. It was found that the IRS significantly supports tax compliance and intentionally uses its budget to meet this objective. However, the main objective of the IRS is revenue maximization mainly because revenue is effectively traceable by the U.S. Congress. Also, audit strategies followed for low-income non-business taxpayers differ from those for other taxpayer classes.

Further work about audit strategies includes [74]. A model of individual tax compliance is developed, taking into account tax rates, payroll tax contributions and benefits with audit probabilities and tax penalties as parameters. Taxpayers seem to be less affected by large tax penalties than lower tax rates. Large tax penalties must be used in isolation because other factors may lead to a decrease in the tax base. In [75], the effect of the decline of audit rates between 1977 and 1986 in the U.S is examined. Based on IRS data, it is argued that the decline of audit rates has considerably limited potential revenues for the U.S. government. The decrease of audit rates may have been concealed by the stagnation of revenues throughout that period. Finally, in [76], data from emerging economies are analysed in order to examine the determinants of individual audit selection. More specifically, the authors examined data of self-employed Jamaican taxpayers from the years 1980 and 1982. A three-stage estimation procedure was employed to determine the strategic behaviour of the taxpayers. In the first two stages, a tobit analysis was employed to determine the parameters affecting individual audit and tax compliance probabilities. In the final stage, a regression analysis was performed using income and tax evasion as dependent parameters. The findings suggest that the audit procedure relies heavily on information derived from past taxpayers' behaviour. As a result, the audit probability should not be regarded as an exogenous variable.

Tax evasion and tax rates seem positively related. Although there is a positive relationship between audit probability and gross tax income, [65] points out that the tax liability is not increased by the same amount for all taxpayer groups. The authors used tax returns as an audit probability proxy, based on TCMP published data, and discussed the inability of the U.S. IRS to effectively allocate the available budget to maximize potential tax revenues. On the other hand, in [77], the IRS is considered effective in targeting the audits to the areas that present increased potential for tax evasion.

The incentives for taxpayers to report honestly are examined in [78]. A field experiment was performed with random audits. Taxpayers belonging to the lower and middle income classes seemed more sensitive to letters issued by the tax authority, warning of a potential audit. On the other hand, higher income classes were found to be less intimidated by such actions. The authors attributed this behaviour to the higher income class' view that the process is a kind of negotiation rather than a potential illegal action. As far as tax penalties are concerned, the effectiveness of tax penalty is questioned in [77], as is its ability to limit tax evasion. However, in [26] it is shown that tax penalties are severely "handicapped" as tax deterrents, especially when they are not backed up by an effective audit mechanism. For more research in that area, [79] argued that the level of compliance decreases when tax penalties increase, especially when audit probabilities are not high enough to support the tax collection mechanism. The information the taxpayers use to estimate the audit probability is usually based on the opinion of tax practitioners (accountants, internal auditors, etc.).

2.3.3 The role of tax practitioners

The practitioners' role on tax evasion and information asymmetry is examined in [80] where it is argued that they are important because they limit potential errors in reporting and, as a result, reinforce tax compliance. Practitioners also help decrease the audit probability and, consequently, the cost of auditing. Compliance is promoted for taxpayer classes characterized mainly by wage and salary reporting. However, in [80] it is also argued that the practitioners' role is ambiguous in cases where creative accounting may be employed. In [3], significant differences are found to exist between taxpayers that choose to be represented versus those that file their own tax reports.

Also on the role of practitioners, [81] jointly analysed their tactics and tax evasion based on data published by the IRS. The findings suggest that the choice of representation is based on the source of the income rather than the income level per se. Moreover, representation enhances tax evasion in cases where creative accounting is employable.

2.3.4 Tax evasion detection

Yet another challenging question researchers are called to answer is how tax evasion can be detected. In [55], the authors examined TCMP data between 1982 and 1985, published by the IRS, and introduced an econometric model that examined the compliance behaviour by using the compliance level as a dependent variable. The model consisted of two equations, the first is estimated in both cases of compliance and non-compliance, while the second equation only in the presence of non-compliance. The explanatory variables included mainly sociodemographic parameters. One drawback of the methodology is that the sociodemographic parameters, due to the nature of the analysis,were only rough estimates of their true values.

Further study on the detection problem and its implications on audit strategies was taken up in [82]. An econometric model was proposed aiming to examine the relationship between tax authority and government tax audit rules, and the correlation between tax authority and government compliance. According to that work, differences exist between the tax authority's and government's optimal budget allocations, and revenues can potentially increase by sharing information between tax agents. A different approach for determining tax evasion was introduced in [83]. Controls for non-random allocation of taxpayers through an endogenous switching specification were used. Two separate measures of compliance detection were taken into consideration. A detection correction measure was used to take into account the failure of tax agents to fully uncover non-compliance. The findings suggested that although tax penalties and regular audits are significant tax evasion deterrents, the most important attribute that increases compliance is the taxation system structure itself. Finally, in [84], the increased importance of Generally Accepted Accounting Principles (GAAP) as a means of tracking taxable income was discussed. The authors of that work proposed that information about taxable income is concealed from outsiders and that tax evasion detection is extremely difficult to track. Moreover, the multinational presence of some firms was considered to be another factor that encourages tax evasion.

2.3.5 Morality

Taxpayer's morality is an important factor that determines tax compliance, and there is a growing literature regarding this matter. Most researchers "grouped" morality effects into three categories: i) the individual's personal morality that dictates his or her behaviour, ii) how fair the taxation system is regarded to be by the taxpayers, and iii) welfare maximization. The introduction of "morals" in tax evasion was originally proposed in [85]. The relationship between the psychological factors of shame and guilt with tax evasion was examined, along with the taxpayer's expectations regarding the audit probability. The authors used IRS data from the state of Oregon and showed that expected utility tends to "overprice" the level of tax compliance. However, considerable heterogeneity exists among taxpayers along with an upwards bias for the audit probability. Furthermore, the explanatory power of shame is much stronger than that of guilt.

An experimental approach examining the relationship between the fairness of the tax system and tax evasion was taken up by [86]. Their hypothesis of a positive relationship between unfairness and tax evasion was shown to be robust; however, the actual effect on tax revenues could not be estimated. A similar experiment was presented in [87], where no significant effect on tax evasion of the tax rates was found.

In [17], a dynamic aspect of morality among the taxpayers was examined. The authors argued that morality varies among taxpayers and several equilibria in tax compliance among taxpayers exist. As a result, tax evasion levels change with time. In [59] it is evident that tax compliance increases with the expectation of better future public goods offered. These results are further supported by [87]. It is suggested that the lawfulness of a taxpayer increases tax compliance significantly. The question of how the above empirical evidence could be potentially incorporated into a theoretical model is examined in [88], where a dynamic analysis is presented between the three moral factors presented above (fairness, justice, lawfulness of taxpayers), and tax evasion.

According to [3], most researchers assume that the taxpayer is able but potentially not willing to disclose the full level of his or her income. However, it could be the complexity of the tax system that prevents them from reporting their true income. This may be caused by the complexity or vagueness of laws governing tax reporting, for example. This problem also extends to auditors and jurisdictions as well. As a result, a portion of tax evasion may be unwitting, and in some cases undetected by tax authorities. In other words, the taxpayer is unable to forfeit the risk of under-reporting. This situation is referred to in [3] as "random enforcement". In case of a miss-audit (the tax authority miss-verification of the taxpayer's true income), the taxpayer is refunded the excess tax. Consequently, tax authorities deliberately nurse randomness to further enhance tax compliance and simultaneously keep audit costs at bay. The effect of tax complexity and its relationship with randomness is also the subject of [89], where a positive relationship between the two exist. Some taxpayers will seek costly advice, while others will not do so and will file their report nonetheless. Complexity and randomness are optimal in the sense that they enhance tax revenues considerably, because non-informed taxpayers will pay more and audit costs will be transferred to gamblers.

Another significant factor that affects tax evasion is the "quality" of taxation system as viewed by taxpayers. In particular, tax compliance and eventual tax evasion is highly affected by the conception that prevails among taxpayers about tax system's fairness [77]. In [90], the authors failed to detect a significant relationship between the tax system fairness and tax evasion. However, the same work argues that this result is not robust due to the variations in perception of fairness. In [91], the tax revenues' destination (e.g war expenditures, national security etc.) are found to be important in forming taxpayers' opinions on the tax system.

2.3.6 Tax Amnesty

A significant factor that affects tax compliance is tax amnesty. As discussed in [3], tax amnesty was a common part of the U.S. taxation system during the 80s, which provided the option for taxpayers to pay taxes for any previously non disclosed income, with the promise not to be subject to an audit. However, this option has no effect on non-compliance because it actually provided no incentive for taxpayers to use it. Instead, it provided better opportunities for taxpayers to calibrate their tax evasion strategies. Research about amnesty has been undertaken in [92] based on data across the U.S. The authors' findings suggest that by including accounts receivables into the estimates, tax revenues can be considerably boosted. However, the long-term effect of amnesty on tax compliance remains to be investigated. Two studies, [93] and [26], examined the long-term effect of amnesty in India and Greece respectively. They both argued that the long-term effect of amnesty in tax revenues is considerably negative. Furthermore, [94] and [76] suggest that even the short-term effects of tax amnesty is non-existent.

2.3.7 Separation between government and tax authority

A growing portion of the literature regarding tax evasion, and tax compliance in general, focuses on the separation of incentives between management and ownership of a corporation. The question that arises, posed in [84], is how these incentives can be constructed in order to be tax efficient, meaning that managers should still be able to maximize shareholders' wealth. As in the government-tax authority distinction, the researchers' aim is to identify the potential conflicts in incentives between management and owners that can lead to increased tax evasion.

In [95], the risk-neutral shareholders of a firm are "separated" from the risk-averse management, assuming that the management, aiming for short-term profit maximization, pursue practices whose incremental benefit will exceed their incremental cost. In [96], the potential challenge on behalf of the shareholders to create incentives for management to increase tax compliance is discussed. The efficiency loss of such a separation, in terms of tax evasion, is examined in [97]. In that work, a principal-agent model is developed to estimate the loss, thus providing a theoretical platform for future research regarding the conflict of interest between management and ownership. Moreover, [98], whose work is also based on the principal-agent framework, includes in the analysis of incentives the compensation schemes of managers and aims to measure the efficiency of tax compliance in corporate governance. Finally, in [99] the tendency of managers to construct firms in a way that diverts tax liabilities and keeps them to a minimum is discussed. The authors encourage an increased monitoring of managers from both shareholders and government.

Chapter 3

Model

We consider a firm which, at the end of each fiscal year, must declare its net profit to the government or tax authority. We proceed to describe the core components of our model, in the form of an Markov decision process which captures the salient features of the Greek tax system. We will make use of the following notation. The integer $k = 0, 1, 2, ...$ will denote discrete time, and x_k will be the value of the quantity x at time k. Individual elements of a vector, x, or matrix M, will be indicated by $[x]_i$ and $[M]_{ij}$, respectively. Finally, $0_{i \times j}$ will denote a i-by-j matrix of zeros.

3.1 State space

We will let $s_k \in \mathcal{S}$ be the tax status of a representative firm in year k, with

$$\mathcal{S} = \{V_1, ..., V_5, O_1, ..., O_5, N_1, ..., N_5\}, \tag{3.1}$$

where

- V_i: the firm is being audited so that its true income for the last $i = 1, ..., 5$ years is verified.

- O_i: the firm has decided to use the closure option and has neither employed closure nor been audited in the past $i = 1, ..., 5$ years,

- N_i: the firm's last audit or closure was $i = 1, ..., 5$ years ago. Thus, the firm has i unaudited tax years and will now make its $(i+1)$-st consecutive decision since its last audit/closure.

At the same time, $c_k \in \mathcal{C}$ will be the status of the closure option in year k, where

$$\mathcal{C} = \{\textit{option available}, \quad \textit{option not available}\} \tag{3.2}$$

The elements of $\mathcal{S}$ and $\mathcal{C}$ were labeled as above mainly for the purpose of facilitating the discussion. However, for the sake of notational convenience, we will sometimes refer to them by integer, in their order of appearance in $\mathcal{S}$ or $\mathcal{C}$, i.e., $V_1 \to 1$, $V_2 \to$ 2,...,$N_5 \to 15$ for states in $\mathcal{S}$, and *option available* $\to 1$, *option not available* $\to 2$ for $\mathcal{C}$.

Each year, k, the firm makes its decisions in the form of a two-element vector, u_k whose first element $[u_k]_1 \in [0, 1]$ is the fraction of its profits to conceal, whereas $[u_k]_2 \in \{1, 2\}$ corresponds to its selection on whether to use the option (if available). Based on the above, we define the firm's state vector at time k to be

$$x_k = [s_k, c_k, h_k^T]^T, \tag{3.3}$$

where $s_k \in \mathcal{S}$, $c_k \in \mathcal{C}$, and $h_k \in [0, 1]^5$ will contain a history of the firm's latest five decisions with respect to tax evasion. We will refer to s_k as the firm's "Markov state" to distinguish it from the state (vector) proper, x_k.

3.2 State Evolution

Each year, the firm's status will evolve in $\mathcal{S} \times \mathcal{C} \times [0, 1]^5$ according to a Markov decision process, with transition probabilities that depend on whether the government audits the firm or offers the closure option, and on whether the firm decides to use the option. Specifically,

$$x_{k+1} = Ax_k + Bu_k + n_k, \quad x(0) \text{ given}, \tag{3.4}$$

where

$$A = \begin{bmatrix} 0 & & \\ & 0 & \\ & & H \end{bmatrix}, \quad H = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad n_k = \begin{bmatrix} w_k \\ \epsilon_k \\ 0_{5\times 1} \end{bmatrix} \tag{3.5}$$

and the terms $\epsilon_k \in \mathcal{C}$, $w_k \in \mathcal{S}$ are (independent) random variables whose distributions are discussed next.

Assuming that the government offers the option with a fixed probability, p_o, each year, we can write

$$Pr(\epsilon_k = i) = \begin{cases} p_o & \text{if } i = 1 \quad \text{(option available)} \\ 1 - p_o & \text{if } i = 2 \quad \text{(option not available)} \end{cases} \tag{3.6}$$

where, for notational convenience we have labeled the elements of $\mathcal{C}$ by integer.

The term w_k in Eq. 3.4 determines the first element of the state vector (i.e., the firm's "next" Markov state in $\mathcal{S}$). Before writing down the probability distribution for w, it will be helpful to have some intuition as to what kinds of transitions are possible in $\mathcal{S}$. Based on our description of the Greek tax system, there will be three possible transition diagrams: i) $c_k = 1$ (the option is offered) and $[u_k]_2 = 1$ (the firm takes the option), ii) $c_k = 1$ and $[u_k]_2 = 2$ (the firm declines the option), and iii) $c_k = 0$ (option not available). The situation is illustrated in Figure 3-1,

where cases ii) and iii) described previously have transition diagrams which are structurally identical but differ in their transition probabilities. For example, if the option is available and the firm has decided to use it (Figure 3-1-A) then in the next fiscal year it will transition to an option (O_i) state with probability 1. If there is no option available, or if there is but the firm has declined it (Figure 3-1-B) then in the next year the firm moves to either a N_i state (i.e., "accumulates" one more tax statement which may be audited in the future) or a V_i state (it is audited).

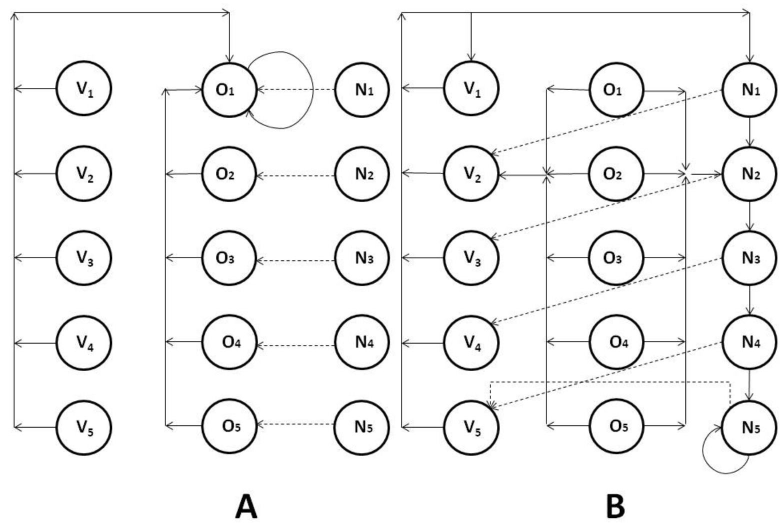

Figure 3-1: A: transition diagram in the case where closure is available and the firms uses it. B: transitions when closure is available but the firm declines, or is not available at all (these two cases have structurally identical diagrams and are differentiated only by their transition probabilities). Each arrow represents a non-zero transition probability between a pair of Markov states in $\mathcal{S}$. The transition probabilities are omitted to avoid clutter, but are included in the Appendix.

The transition from N_5 to itself indicates the fact that, because of the statute of limitations, a firm's tax history can usually only be scrutinized for the last fiscal year plus another four years in the past. For each transition diagram we can write down a corresponding Markov matrix whose (i, j)-th element represents the probability of transitioning from the j-th to the i-th element of $\mathcal{S}$. We will denote these matrices (see Appendix A) by M_{no}, for the case where the option is *n*ot *o*ffered ($c_k = 2$), M_a for the case where $c_k = 1$ and the firm *a*ccepts, and M_d if $c_k = 1$ but the firm *d*eclines the option.

Returning now to Eq. 3.4, the above discussion implies that w_k's distribution will depend on x_k and $[u_k]_2$, because the random transitions that the firm undergoes in $\mathcal{S}$ depend on its existing state as well as its decision to accept or reject the option (if it is offered). In particular,

$$Pr\left(w_k = i | x_k = [j, q, h_k^T]^T, [u_k]_2 = m\right) = P_{qij}(m), \quad i, j \in \{1, ..., 15\}, q \in \{1, 2\} \tag{3.7}$$

where, for q and u fixed, the $P_{qij}(m)$ form one of the Markov matrices M_{no}, M_a, M_d, governing the firm's transitions in $\mathcal{S}$:

$$P_{qij}(m) = \begin{cases} [M_{no}]_{ij} & \text{if } q = 2, \ \forall\, m \quad \text{(no option)} \\ [M_a]_{ij} & \text{if } q = 1, \ m = 1 \quad \text{(option taken)} \\ [M_d]_{ij} & \text{if } q = 1, \ m = 2 \quad \text{(option declined)} \end{cases} \tag{3.8}$$

3.3 State rewards and Optimal Value function

Let R denote the firm's annual profit, r the nominal tax rate (currently at 0.24), β the annual penalty rate for past uncollected taxes applied in the event of an audit (currently at 0.24 as well) and, finally, ℓ the cost of closure as a fraction of the firm's profits. Based on our earlier discussion of how tax penalties are determined (Chapter 1), the firm's reward associated with making a decision u_k while at state x_k is

$$g(x_k, u_k) = g\left([s_k, c_k, h_k^T]^T, u_k\right) = R \cdot \begin{cases} 1 - r + r[u_k]_1 & s_k \in \{11, ..., 15\} \\ 1 - r + r[u_k]_1 - \ell(s_k - 5) & s_k \in \{6, ..., 10\} \\ 1 - r + r[u_k]_1 - r\sum_{i=1}^{s_k}[h_k]_{6-i} & \\ -\frac{3}{5}\beta r \sum_{i=1}^{s_k} i[h_k]_{6-i} & s_k \in \{1, ..., 5\} \end{cases} \tag{3.9}$$

where we have again labeled elements of $\mathcal{S}$ and $\mathcal{C}$ by integer. The top term in the right-hand side of Eq. 3.9 corresponds to the reward obtained if the firm conceals an amount of $R[u_k]_1$. The second term is the reward 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 owed and the penalty.

Assuming a time horizon of N years, the firm is then faced with the problem of choosing its policy, u_k, so as to maximize its discounted expected reward:

$$\max_{u_k} \mathop{\mathbb{E}}_{w_k, \epsilon_k} \left\{ \sum_{k=0}^{N-1} \gamma^k g(x_k, u_k) \right\}, \tag{3.10}$$

where $\gamma \in (0, 1]$ is a discount factor. Here, we will be interested mainly in the case where the firm assumes its economic lifetime will be infinite and acts accordingly, thus we will not be concerned with possible end-of-horizon effects. However, such effects can easily be incorporated into the model, by adding a $\gamma^N g_N(x_N)$ term to Eq. 3.10, in order to capture, for example, a situation where the firm does not make a decision in its last year of economic lifetime.

To apply dynamic programming [100], let $J_k(x_k)$ be the optimal expected reward that can be obtained from time k onwards, starting from x_k. Then, the J_k will satisfy the following recursive equation:

$$\begin{aligned} J_k(x_k) &= \max_{u_k} \left\{ g(x_k, u_k) + \gamma \mathop{\mathbb{E}}_{w_k, \epsilon_k} J_{k+1}(x_{k+1}) \right\}, \\ J_N(x_N) &= g_N(x_N) \end{aligned} \tag{3.11}$$

subject to the state dynamics, Eq. 3.4.

Assuming that the firm expects to operate indefinitely ($N \to \infty$), its optimal decisions are then obtained by solving the stationary Bellman equation associated with Eq. 3.11, i.e.

$$J_\infty(x) = \max_u \left\{ g(x, u) + \gamma \mathop{\mathbb{E}}_{w,\epsilon} J_\infty(Ax + Bu + n) \right\} \tag{3.12}$$

Finally, given the probability distributions for w and ϵ (Eq. 3.6 and Eq. 3.7), and the state equation (Eq. 3.4), Eq. 3.12 can be re-written as

$$J_\infty(i, q, h) = \max_u \left\{ g(i, q, h, u) + \gamma \sum_{t=1}^{2} \sum_{j=1}^{15} P_{qji}([u]_2)\, Pr(\epsilon = t) J_\infty(j, t, Hh + e_5[u]_1) \right\} \tag{3.13}$$

where $e_5 = [0, 0, 0, 0, 1]^T$ and, for convenience, we have slightly abused the notation by writing the argument of J_∞ as (i, q, h) instead of $x = [i, q, h^T]^T$, and that of g as (i, q, h, u) instead of $\big(x = [i, q, h^T]^T, u\big)$, with $i = 1, ..., |\mathcal{S}|$, $q = 1, ..., |\mathcal{C}|$, $h \in [0, 1]^5$.

Notice that the reward function, Eq. 3.9, as well as the state transitions, Eq. 3.4, are linear in the fraction of profits to be concealed, $[u_k]_1$, for all k. Thus, using an argument similar to that from [101], we can conclude that all the J_k, as well as J_∞, will be linear in $[u_k]_1$. Consequently, J_∞ will be maximized at the boundary of $[u_k]_1$'s feasible region, and the firm should follow a "bang-bang" policy of either $[u_k]_1 = 0$ or $[u_k]_1 = 1$ each year (we will have more to say about these extreme values shortly). This implies a significant reduction in computational complexity, because it will be sufficient to consider $h \in \{0,1\}^5$, and calculate Eq. 3.13 only on a finite set of $|\mathcal{S}| \cdot |\mathcal{C}| \cdot 2^5 = 869$ states. The latter can be done in a straightforward manner using value iteration.

3.4 Parameter dependence

We are interested in determining the optimal expected reward's dependence on the main parameters of the tax system as described in the previous Section, including the tax rate, r, penalty factor, β, and closure cost, ℓ. For an annual closure probability $p_o \in (0,1)$, the reward function (Eq. 3.9), is decreasing in r, β and ℓ. Moreover, the parameters r, β and ℓ affect neither the state transitions, Eq. 3.4, nor the probability distributions with respect to which the expectation is taken in Eq. 3.11. We conclude that the term $\mathbb{E}J_{k+1}$ in Eq. 3.11 will be decreasing in r, β or ℓ, for any k, and thus, based on [101], the same property will be shared by J_∞.

The effect of the probability of the closure option being offered, p_o, on the value function depends on i) whether the firm's expected reward is higher if the option is offered, and if so, ii) whether it is better for the firm to use the option. If we assume the firm's state vector is $x_k = [j, 1, h_k^T]^T$, with $j = 11, ..., 15$ (i.e., the firm is not in a V or O Markov state, and closure is available), then the firm has a choice of either transitioning to a closure state (according to the probabilities in Eq. 3.7), or taking its chances and perhaps being audited in the next period, with higher probability than if the option was not available at all. One can then check that, based on the rewards (Eq. 3.9), and transition probabilities (Eq. 3.8), the firm's expected reward by taking the option at some time k will be higher than that obtained by declining it, iff the probability of an audit, given that the firm forgoes the option while in the

j-th Markov state, $P_{11j}(2)$, satisfies

$$P_{11j}(2) \geq \theta(j, h_k) \triangleq \frac{\ell \cdot (j-10)}{r \cdot \left(\sum_{i=1}^{j-10} [h_k]_{6-i} + \frac{3}{5}\beta \sum_{i=1}^{j-10} i[h_k]_{6-i} \right)} \qquad j = 11, ..., 15, \quad (3.14)$$

with $(j-10)$ being the number of years since the firm's last closure or audit event, and h_k the lower five elements of the firm's state vector. Furthermore, if $P_{21j} \geq \theta(j, h_k)$ as well (where P_{21j} is the probability of transitioning to an audit Markov state given that closure was not available), then saying "yes" to closure is preferable to not having the option at all. Because the audit probability is increased if a firm is given the option and declines it (relative to the case where there was no option at all), i.e., $P_{11j}(2) > P_{21j}$, it will always be advantageous for the firm to have (and use) the option if it would be willing to do so under the "usual" lower audit probabilities. On the other hand, if $P_{11j}(2)$ is below the threshold θ in Eq. 3.14 then so is P_{21j}, meaning that the firm's reward would be lower under the option if taking the option is more expensive than discarding it. In general, the threshold θ that makes the option desirable for the firm depends on the firm's state vector (i.e., on the Markov state, j, it is in, and on its past history of decisions, h_k. However, there are choices for β, ℓ, and r (including those in use today, discussed in the next section) which are of practical interest and lead to the option being to the firm's advantage uniformly, for any j and h_k.

Based on the above discussion, if $P_{21j} > \theta(j, h_k)$ along the optimal trajectory, then the firm will always say yes to closure, and having the option to do so is better than the alternative. In that case, the value functions J_k will be increasing in the closure probability, p_o, for all k. To see why that is, notice that the reward function, Eq. 3.9, is independent of p_o and of the actual availability of closure, c_k. Also, in terms of the state transitions, Eq. 3.4, p_o only affects $Pr(\epsilon = 1)$ (the probability of arriving at a state where the firm has the option to use closure) and none of the remaining elements of the state vector (in particular, s_k) on which the reward function depends. Thus, the term $\mathbb{E}J_{k+1}(Ax_k + Bu_k + n_k)$ in Eq. 3.11 will be increasing in p_o, because an increase in p_o simply corresponds to a higher probability of a more favorable outcome (namely closure). These facts imply [101] that the long-term optimal reward function J_∞ will be increasing in p_o as well.

Using similar arguments, one can show that if either $P_{11j}(2)$ or P_{21j} are below the threshold, Eq. 3.14, then the value function will be decreasing in p_o because in those cases p_o increases the probability of a lower-payoff event, while leaving the state transitions and reward function unaffected. Finally, in the extreme case where $p_o = 0$, the closure cost ℓ becomes irrelevant (since closure is never offered), while if $p_o = 1$ and $P_{21j} > \theta(j, h_k)$, then the value function will be independent of β (because the firm employs closure every year and will never be audited).

3.5 Assumptions and Parameter selection

Our model includes a few assumptions which require justification. For simplicity, we will assume that the firm's annual profit, R, is constant throughout its economic life. It is straightforward to allow R to rise at a steady rate, by manipulating the discount coefficient γ; the model could also be easily adapted to include any other pre-defined growth profile for R. In the following, we will assume an interest rate of 3%, corresponding to $\gamma = 1/(1 + 0.03) = 0.9709$.

Regarding the choices of R, $[u_k]_1$, and ℓ, we will always refer to "relative" amounts, so that, for example, $R = 100$, and $R \cdot [u_k]_1$ is the *percentage* of year k profit to be hidden from the authorities. We chose this approach for the following reasons. The firm's decisions are, of course, based on its true profit, which the government does *not* know. Expressing the cost of closure, ℓ, as a percentage of the firm's net profit, and $[u_k]_1$ as the fraction of profits to be hidden, will make it easier to draw conclusions as to the effectiveness of tax measures and behavior of the firm. Furthermore, given an estimate of the size of the country's "hidden economy" (studies such as [102] place it at around 40% for Greece), the quantities computed by the model can be converted to estimates of absolute amounts.

Regarding the range of values for $[u_k]_1 \in [0, 1]$, it may be impossible for a firm to systematically claim zero annual profits by overstating expenses and/or hiding income. There are several practical reasons for this, including pressure by shareholders or capital markets to demonstrate profitability, and other safeguards in the accounting system, so that at least some income will be documented (e.g., via sales invoices which some clients will likely demand in the course of business). There are at least two possible approaches here. One is to set some upper bound $u_{max} < 1$, so that a firm with profit R can never hide more than $u_{max}R$. This is meaningful in certain settings, but again requires knowledge - by the government - of the firm's true profit. Instead, here we will allow $u_{max} = 1$, and interpret the model's results in a "marginal" sense: $[u_k]_1 = 1$ simply means that the firm should hide as much of its profit as possible, or that the next euro that could be hidden, should be hidden.

Our model can easily be used to examine the effects of applied tax rates and audit probabilities, however, these quantities will be kept fixed to their estimated current levels. We do this in order to isolate the effect of tax-penalties and closure cost on the firm's behavior, and because, in the case of audits, an increase is not easy to implement (e.g., it may require hiring of new personnel, training, etc.). Because of space considerations, we discuss only income tax and ignore VAT collection and payments by the firm, which are subject to a separate mechanism and can be incorporated in the model at a later stage.

With respect to the audit probability distribution and closure-related data, there is a scarcity of official reports. In order to demonstrate our model, we have estimated the various parameters of interest using other sources, including reports in the Greek financial press, which suggest that audits can cover no more than approximately 5% of all firms in a given year, and that the cost of closure for the options offered during 1998-2008 was approximately 2-3% of profits for the average firm. Based on these, we assumed an overall audit probability of 0.05. This probability is distributed heavily (80%) towards firms with past tax declarations whose statute of limitations is about to expire, i.e., a 0.0025 probability that the firm audited is drawn from those with 1-4 years since their last audit or closure, and a 0.04 probability that it is one of those which have not been audited for five years. Of course, the model can be easily adjusted to different parameter values at the hands of government entities which would have more precise knowledge of the parameters.

As we have already mentioned, the probability of an audit increases when the firm rejects the option to use closure. That increase will depend on the number of firms which choose the option of closure, leaving the rest to increased scrutiny. There is little official data on this; here, we have used a rough estimate of 2/3 for the fraction of firms who opt to use closure, meaning that the audit probability is roughly tripled for those who don't. Arguably, the rate of participation is determined largely by expectations, whose dynamics are however, beyond the scope of this paper. Finally, we chose the annual probability of closure being offered to be $Pr(\epsilon = 1) = 0.2$, because that value is near the current average. We opted for a probabilistic treatment of closure mainly for two reasons. One was that it is easily implementable in practice and makes for a tractable model. The other reason is that we would like to make comparisons between random vs. periodic closure. In Greece, there is a history of closure being offered quasi-periodically, roughly every 5 years (as the statute of limitations on tax statements is about to elapse). One may hypothesize that this is anticipated by firms which may alter their policy to take advantage. As our numerical experiments will show, this can indeed happen and it is best for the government to not allow firms to anticipate when the option will be offered (in fact, it seems best not to offer it at all).

Chapter 4

Running the model: Results and Discussion

We implemented our model as a MATLAB-based application, and obtained results on a series of scenarios of practical importance, regarding the effect of tax penalties and the closure option on firm behavior and government revenues. The experiments described below are arranged based on whether the closure option is: i) stochastically available ($p_o = 0.2$), ii) always available ($p_o = 1$), iii) never available ($p_o = 0$), and iv) available every five fiscal years, where for the latter scenario the basic model was modified so that $Pr(\epsilon_k = i)$ was periodic in k. In each case, we kept the tax rate and audit probabilities fixed, and allowed the closure cost to vary in the range between $\ell = 0.023$ (a rather low figure, and our current estimated level) and $\ell = 0.50$, in increments of 0.01 (values beyond 0.50 are not considered realistic, given the current tax rate of $r = 0.24$). In that range, we determined a kind of boundary, (in the tax penalty - closure cost space) at which the behavior of the firm changes from being dishonest in every feasible[1] state, to being honest in a) at least one feasible state, and b) in all feasible states. We will refer to these as the *total* and *partial honesty* boundaries, respectively. We are also interested in the boundary at which the firm's policy changes from always using the option, to discarding it, in a) at least

[1]By "feasible" we mean a state which will be visited with non-zero probability under the firm's optimal decision policy, u_k.

one feasible state and b) in every feasible state, for a range of tax penalty coefficients. For simplicity, we absorbed the 3/5 "prompt payment" discount factor in Eq. 3.9 into the per annum tax penalty coefficient, β. Thus any tax penalty multipliers discussed henceforth are "net" (after discount) values. In order to study the effects of tax penalties over a wider range, we have "removed" the 200% upper bound on accumulated tax penalties which is in place today, as discussed in Chapter 1, and considered the range from $\beta = 0.144$ (the current level, after discount), to $\beta = 4$, or 400%, in increments of 0.05.

4.1 Stochastically available option

In this case, the firm does not know a priori whether closure will be available but does know the corresponding probability distribution, which is set to $p_o = 0.2$ each year. Figure 4-1 shows the output of our model, in the form of the total honesty

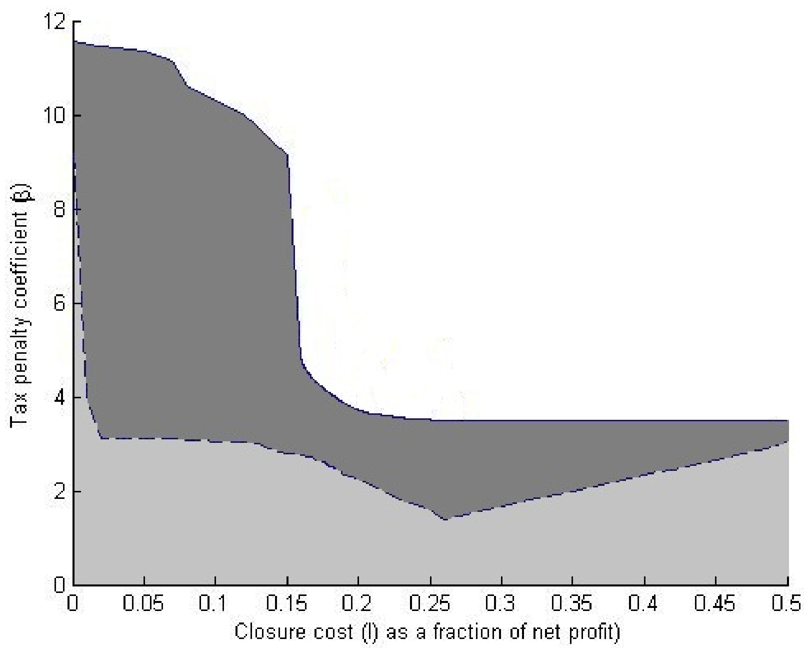

Figure 4-1: Partial (dashed) and total (solid) tax penalty, honesty boundaries when the closure option is available with probability 0.2 each year. The firm uniformly chooses a policy of maximum tax evasion in the light gray area, and always declares all profits in the white area. In the dark gray region the firm's policy is state-dependent.

boundary over which the firm declares 100% of its profit in all feasible states (white area) and the partial honesty boundary below which the firm declares as little profit as possible (light gray area). In the dark gray area between the two boundaries, the firm's policy is state-dependent. The first group of states at which the firm alters its behavior as the tax penalty rises from very low values are those for which $s = 15$ ans $s = 14$, i.e., the firm is in the N_5 or N_4 Markov state - unaudited for five and four years respectively, with no closure-option at her disposal. On the other hand, the last states to "switch" to honest behavior as we cross the upper boundary in Figure 4-1 are those with $s = 11$, where the firm is in N_1 - audited one year ago. In the area between the two boundaries, the firm is honest in at least one feasible state.

We notice that as the closure cost, ℓ, approaches the current tax rate, the tax penalty multiplier required to enforce total honesty declines from 11.8 to 3.9 (i.e., slightly above 300% of back taxes owed), and remains constant for higher values of ℓ. For the same range of ℓ, the partial honesty boundary declines from $\beta = 9.63$ to $\beta = 2.9$. We observe that both boundaries are situated *well above* the current net penalty coefficient (approximately 0.14), even for relatively high closure costs. This agrees with the widely accepted assessment that tax evasion in Greece is high in part because the current combination of ℓ, β, and audit probabilities are ineffective. In that case, the option provides firms with a less costly way of settling their tax obligations. We will have more to say about this in a moment.

Figure 4-2 illustrates the appeal of closure to the firm. Points below the lower boundary (light gray area) correspond to option cost/tax penalty combinations where it is optimal for the firm to use the option in every feasible state. Between boundaries (dark gray area), the firm ignores the option in at least one feasible state. Finally, above the upper boundary (white area), the firm never uses the option. The lower boundary, separating the areas of total vs. partial option acceptance, indicates the highest percentage of net profit that a firm would be willing to pay always in order to "lock in" past gains earned through tax evasion. At today's tax penalty of $\beta = 0.14$, the firm should be willing to pay up to approximately $\ell = 0.05$, or 5% of its net profit, i.e., well above the current estimate of 2-3%.

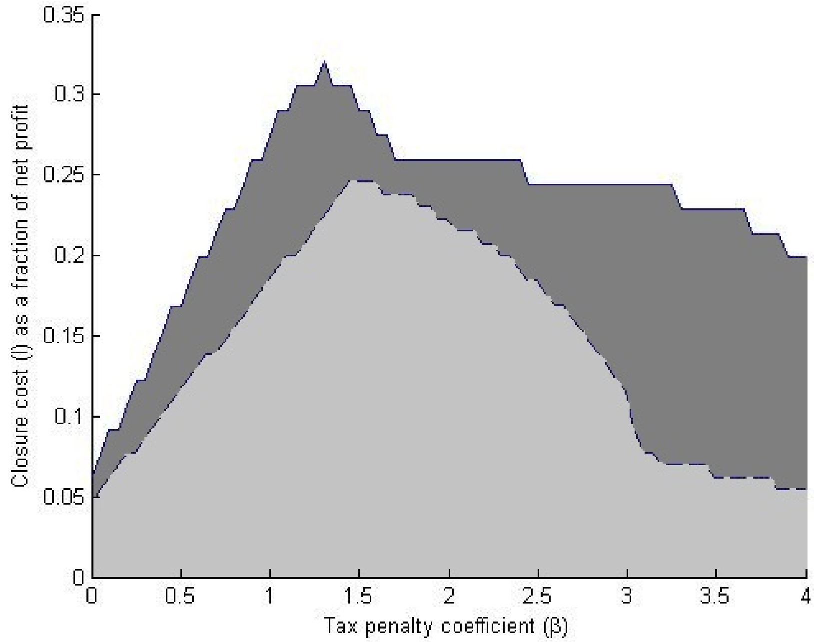

Figure 4-2: Partial (dashed) and total (solid) option usage boundaries when the option is available with probability 0.2 each year. The option is used always at points below the lower boundary (light gray area), never at points above the upper boundary (white area).

The upper boundary of Figure 4-2, beyond which the firm never makes use of the option, also rises in steps, from approximately $\ell = 0.07$ when $\beta = 0$, to almost 0.32 when β reaches 3.6. This indicates how much of a closure cost the firm is willing to accept as the tax penalty increases, as long as tax evasion remains its most profitable choice. In numerical experiments, we have observed that for even higher tax penalties the upper boundary declines similar to the lower boundary; this is because (see also Figure 4-1) tax evasion becomes profitable in fewer and fewer states, and thus the firm is willing to pay gradually less in order to avoid a possible audit.

4.2 Option available every fiscal year

There is a set of firms and freelancers in Greece, that settle their tax-obligation solely through closure every year. In this case the closure option is called "self-assessment". We examined the effectiveness of this policy, in terms of the government's and firm's expected earnings. In that case our model indicated that the partial and total honesty boundaries are situated at more than 12 times the uncollected taxes. Such a tax penalty seems unrealistic; it appears therefore that tax evasion cannot be curbed under this scenario unless the audit mechanism is significantly reinforced. This argues against the frequent use of the closure option as a revenue collecting mechanism.

The lower boundary of total-to-partial option appeal (see Figure 4-3) increases

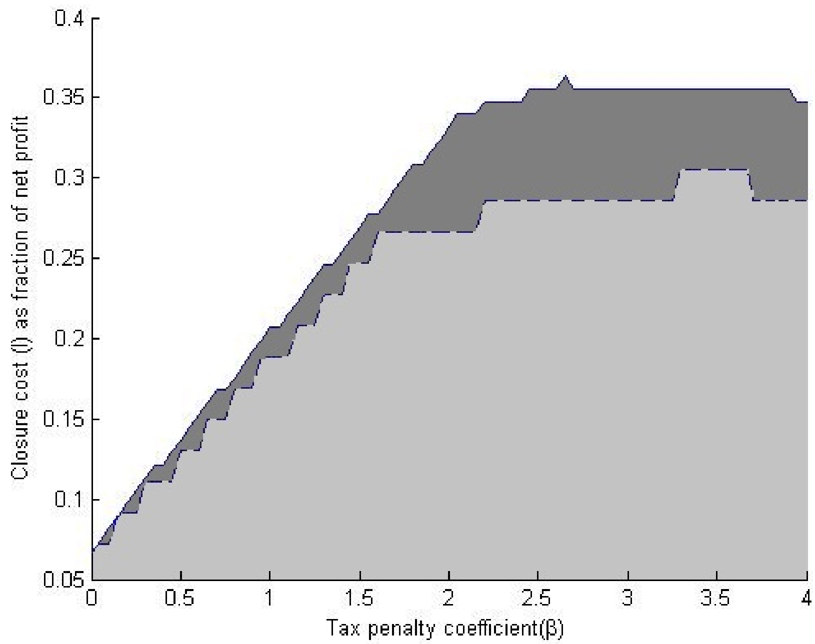

Figure 4-3: Partial (dashed) and total (solid) option usage boundaries when the option is available every year. The option is used always at points below the lower boundary (gray area), and never at points above the upper boundary (white area). In the dark gray region the firm's policy is state-dependent.

from $\ell = 0.06$ to $\ell = 0.30$ for a range of β between 0 and 2.95, followed by a slight decrease to $\ell = 0.28$ for $\beta = 4$.In the light gray region, it is optimal for the firm to use the option every year. The upper threshold, over which the firm refuses the option, rises in steps from $\ell = 0.071$ to $\ell = 0.35$ for the range of β between 0 and 3.9, followed by a decrease to $\ell = 0.34$ when $\beta = 4$. indicating that the firm is willing to accept higher closure costs as long as that cost remains lower than the current tax rate. Overall, our model suggests that if the option is offered every year, enforcing honesty via sufficiently high tax penalties may be infeasible, while for a modest range of tax penalty coefficients most firms would not make use of the closure option.

4.3 The option is never available

If the closure option is never offered by the government, tax evasion persists in every feasible state at today's tax penalty rates. In order to make "full disclosure" an optimal policy for the form in at least one feasible state, the tax penalty needs to be approximately $\beta = 1.7$ (i.e., slightly higher than a 400% annual tax penalty rate on back taxes owed). The first feasible states in which the firm turns honest as β rises are those with $s = 15$ and $s = 14$ (N_5,N_4 - five and four years unaudited respectively). When $\beta = 4.9$, the firm becomes honest in every feasible state, the last group of states to "switch" being those with $s = 11$ (Markov state N_1). Thus, in the absence of a closure option, the firm's optimal policy with today's parameters is to evade taxes, however the tax penalties required to change that are considerably lower than when the option is offered frequently.

4.4 Option available in five year time intervals

In the case where firms may guess that the government will offer the option periodically in an effort to collect revenue from past years whose statute of limitations is about to expire (as is likely to be the case in Greece, given the recent history), the situation is similar to that when the option is offered annually. In particular, the tax penalty thresholds for partial or total honesty, exceed 12 times the amount of uncollected taxes, and the firm's optimal policy is identical to that of Section 4.2, i.e., conceal profits when possible and use the closure option when available.

Figure 4-4 illustrates the appeal of closure to the firm. The percentage of net profit that a firm would be willing to pay for using the option at (at least) one feasible state increases from $\ell = 0.04$ to $\ell = 0.19$ for β ranges between 0 and 1.85, suggesting that the government cannot hope to collect amounts that are much higher than what it can obtain today. Above $\beta = 1.85$ the lower boundary declines to $\ell = 0.17$ for $\beta = 4$.

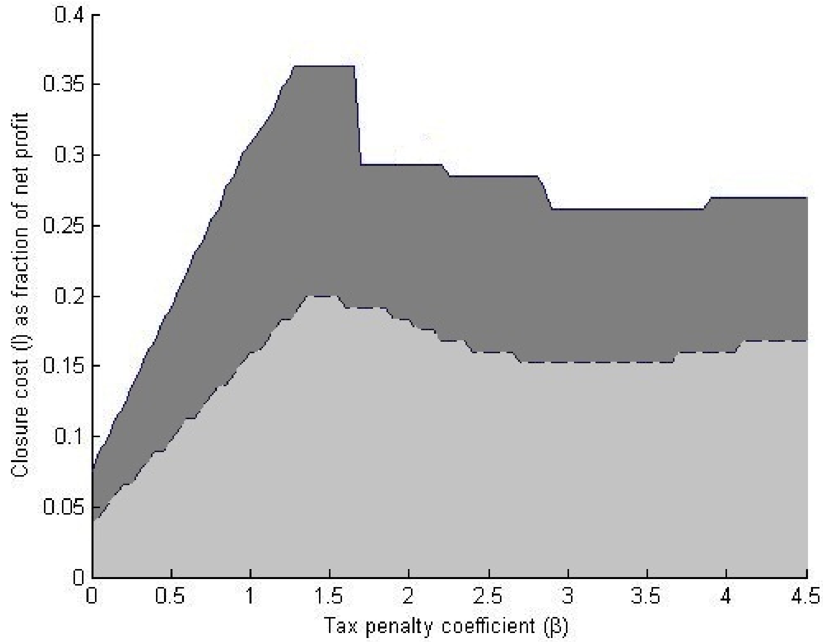

Figure 4-4: Partial (dashed) and total (solid) option usage boundaries when the option is available in five-year intervals. In the light gray area the option is used in every feasible state. At points above the upper boundary (white area) the option is never used. In the dark region the firm's policy is state-dependent.

In this set of states (those with $s = 15$), the firm is willing to pay a progressively higher closure cost (as long as that cost remains lower than the current tax rate). On the other hand, in order for the firm to stop employing the closure option, its cost must increase from approximately $\ell = 0.9$ to $\ell = 0.36$ for β ranges between 0 and 1.55. Above this range the upper boundary decreases gradually to $\ell = 0.27$ for $\beta = 4$. Those states determine the upper boundary of Figure 4-4.

4.5 Comparing Government vs. Firm expected revenues

Assuming that the firm is risk-neutral, with an infinite economic life, and that it employs an optimal tax evasion and option usage strategy, we computed (via Eq. 3.13) the present value of the expected firm revenues and government tax revenues per firm, for each of the scenarios discussed previously. The results are listed in Table 4.1, in terms of % of the firm's annual profit, and assume that future revenues are discounted assuming again a 3% rate of inflation.

Option availability	Expected firm profit (net)	Expected Government Revenue per firm
Never available	3254.6	178.7
Stochastic (20%)	3307.9	125.3
Always available	3358.3	74.8
Every five years	3313.9	120.2

Table 4.1: Comparison of expected firm and government revenues under different option availability scenarios, 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. The figures for the last scenario ("every five years") are for an initial state of $x = [11, 2, 0, 0, 0, 0, 1]^T$, i.e., the firm hid all its profit since its last closure, 1 year ago, and has no closure option at its disposal for the next 4 years.

We observe that the firm maximizes its expected revenues in the case where the option is available every year; that is the worst-case scenario from the point of view of the government. Tax revenue improves gradually if the option is offered 20% of the time, with the highest revenue collected when the option is not offered at all. Offering the option periodically (last line of Table 4.1, and thus allowing firms to anticipate closure, leads to a slightly higher tax revenue than if the option was offered each year. Intuitively, this is explained by the fact that the optimal policy in both cases is $u = [1, 1]^T$ uniformly, i.e., always take advantage of the option and hide all profit. Consequently, the firm will pay the same amount for closure in either case (either yearly or as a lump sum every five years). However, if closure occurs every five years, the firm is exposed to a small probability of an audit between closures, so its long-term expected income is slightly less than if it could use closure yearly and never be audited. These figures support the argument that, in the current state of affairs, the closure option is appealing to firms, and that it is not a productive revenue collecting mechanism: it limits the effectiveness of tax penalty as a tax-deterrent, and promotes tax evasion by providing a cheaper alternative to tax compliance.

The results for the "no-option" scenario suggest that tax penalties would be an important part of the tax revenue collected in that case, since it is optimal for the firm to conceal as much of its profit as possible. However, it may be argued that the government will sometimes encounter difficulties in actually collecting tax penalties and back taxes owed in the event of an audit (e.g., the firm may ultimately be unable to pay because of bankruptcy or other reasons). We adjusted the basic model so as to take this alternative into consideration. Although there are no official data on the percentage of audits which result in an inability to collect, we estimate this figure to be approximately 60%, based on recent discussions in the Greek parliament about the percentage of long-term overdue taxes the government considered uncollectable in 2010. In that case, when the firm applies its optimal policy as before, the expected government revenues in the right column of Table 4.1 decrease to 106.9 (option never available), to 92.99 (option given 20% of the time), 98.54 (option given every five years). The figure for the "option always available" scenario remains at 74.8, because in that case the firm uses the option every year, and is never audited. The best scenario for the government is still to *never* offer the option, followed by offering it on a probabilistic basis, while the worst-case scenario is again to offer it every year. The tax revenue gap between best (from the point of view of the government) two cases becomes narrower as the government is the percentage of uncollected back taxes and penalties increases. Once the effectiveness of the collection mechanism is reduced below 45.4%, then the highest tax revenue is obtained by always offering the option.

The fact that, in Table 4.1, firm revenues in increase for higher probabilities of the closure option being offered, is not coincidental. Based on the parameter values used here (as per our estimates discussed in Section 3.5), we verified that the audit probabilities with no option available, as well as when the option is given and the firm declines it, are greater than the threshold (Eq. 3.14), for all j and for any choice of disclosure history, h_k. Thus, (see discussion in Section 3.4), the firm's long-term expected reward will be increasing in the probability of closure being available, p_o.

Figure 4-5 shows the corresponding (decreasing) government tax revenues.

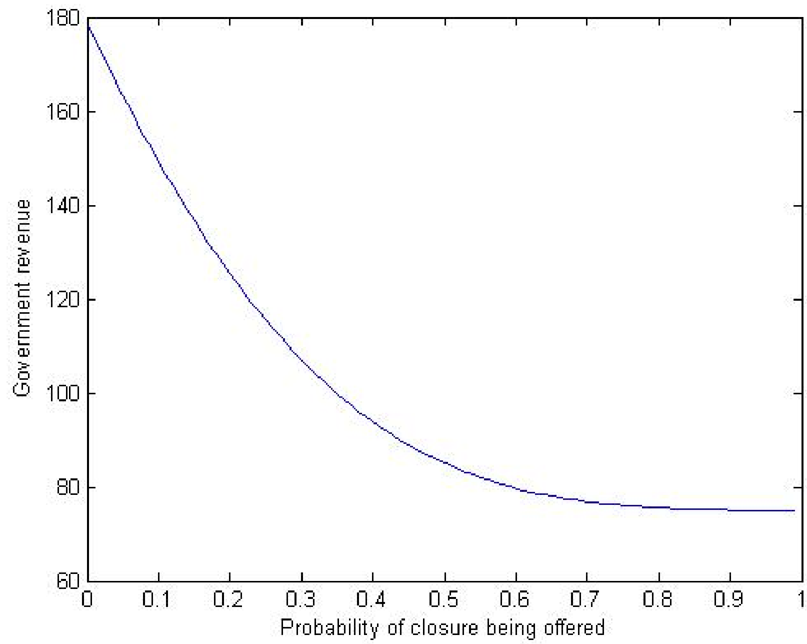

Figure 4-5: Expected tax revenues versus probability of closure being offered each year, assuming the firm employs its optimal policy regarding the fraction of its profits that it hides.

Finally, the problem discussed here can be viewed as a zero-sum game between the government (which may or may not offer closure) and the firm (which chooses whether to conceal its profits). In that setting, p_o and $[u]_1$ represent the government's and the firm's mixed strategies, respectively. Although a full game-theoretic analysis is beyond the scope of this paper, the results presented here are suggestive of an equilibrium at $p_o = 0$ (never offer the option) and $[u]_1 = 1$ (always conceal as much profit as possible) under the parameters in place today. That is because, at $p_o = 0$, the government cannot unilaterally improve its revenue by raising p_o when the firm acts optimally for itself (the firm's value function is increasing in p_o, so that raising p_o causes the government to loose some revenue); at the same time, maximal tax evasion is optimal for the firm when $p_o = 0$.

Chapter 5

Risk-averse firms with no closure option

In Chapters 3 and 4 we explored a Markov-based decision support model assuming risk-neutrality on behalf of the firm. We computed the firm's optimal policy regarding tax evasion and use of the closure option under four scenarios of closure availability, and identified the "virtuous" combinations of tax parameters (i.e., tax penalties and closure cost) that force the firm to behave honestly. The assumption of risk-neutrality may be computationally convenient but is not very realistic. In this Chapter we proceed to consider a firm that acts based on a constant relative risk aversion utility function. Then, if the government is willing to forgo the closure option, the firm's optimal behavior can be obtained by considering a relatively simple portfolio allocation problem.

5.1 The Model

We define a Markov process whose state set, $\mathcal{S}$, describes the possible tax status of a representative firm in any given year. In particular,

$$\mathcal{S} = \{A, N_1, ..., N_4\}, \tag{5.1}$$

where

- A: the firm is being audited,

- N_j: the firm's last audit took place $j = 1, ..., 4$ years ago,

When convenient, we will refer to elements of is by their order of appearance, i.e., the 1st through 5th states, instead of $A, ..., N_4$.

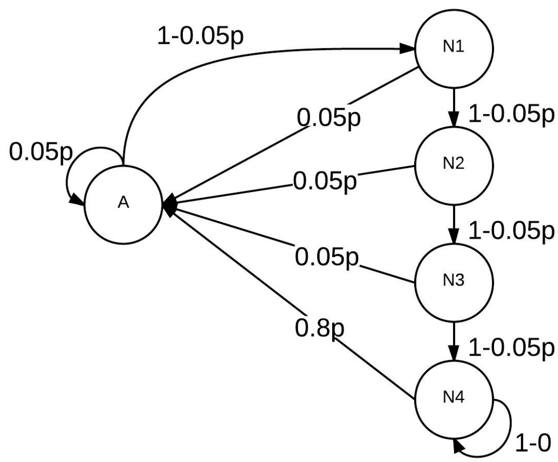

Figure 5-1: Markov transition diagram for our simplified model of the Greek tax system as it pertains to firms.

The transition probabilities p_{ij} with $i, j \in S$ of moving from state i to j are shown in Figure 5-1, where p is the overall 5-year audit probability (5%). While at state $j = 1, ..., 5$, the firm will have to decide what fraction, u, of its profits to attempt to conceal. We will consider the firm's tax evasion decisions in the context of a portfolio $W(u, j)$ containing i) a risk free asset, R, whose payoff is the declared fraction of profits minus the tax due and ii) a risky asset, $B(u, j)$, whose payoff is either the discounted undeclared fraction of profits or — in case of an audit — the tax penalty imposed:

$$W(u, j) = (1 - u)R - uB(u, j). \tag{5.2}$$

The value of the portfolio depends on the state j because the probability of being audited (and thus the value of the risky asset) depends on the number of years since the firm's last audit. If I is the firm's annual profit, r the tax rate and $\beta > 0$ the per-annum tax penalty coefficient, then based on our description of the tax system, we can write:

$$B(u,j) = \begin{cases} Iu & \text{initial state is } j \text{ and audit} \\ & \text{does not happen in } n \text{ years} \\ & \\ I(1 - r - 3/5n\beta r) & \text{initial state is } j \text{ and audit} \\ & \text{happens in } n \text{ years} \end{cases} \tag{5.3}$$

We assume that the firm has constant relative risk aversion, so that its utility function is of the form:

$$U(x) = \frac{x^{1-\lambda}}{1-\lambda}, \tag{5.4}$$

where λ is the risk aversion coefficient of the "average" firm and will have to be estimated. We will have more to say about this in the next Section. The firm's objective is to make the optimal choice regarding the fraction of the profits that it should conceal, so as to maximize the expected utility of the portfolio:

$$G(u,j) = \max_u \left\{ \mathbb{E}(U(W(u,j))) \right\}. \tag{5.5}$$

We note that, because of the five year statute of limitations on tax statements, the firm's current decision will affect its future cash flows for up to five years. We conclude that the firm's behavior will depend on the Markov process' first passage probabilities from any state to the first (audit) state, since it is those probabilities who will ultimately determine the expected utility of the firm's decision. More specifically let $f_{ij}^{(n)}$ stand for the probability that the firm, starting at a Markov state i will reach state j for the first time in n steps. In our case, we assume that the firm begins its economic life in the audit (A) state, without a history of past transgressions. It is well known that, as in [103], that the probabilities $f_{ij}^{(n)}$ satisfy the following system of linear equations:

$$\alpha_{ij}^{(n)} = \sum_{r=1}^{n} f_{ij}^{(n)} \alpha_{ii}^{(n-r)}, \tag{5.6}$$

where the $\alpha_{ij}^{(n)}$ are the elements of the n-th power of the Markov transition matrix corresponding to the process in Figure 5-1. The numerical values for the $f_{ij}^{(n)}$ are included in the Appendix C. Based on the above, Eq. 5.5 can be rewritten as follows:

$$\begin{aligned} G(u,j) = \max_{u} \Bigg\{ & \sum_{n=1}^{5} f_{1j}^{(n)} U(\gamma^n I(1 - r - 3/5n\beta r)) \\ & + (1 - \sum_{n=1}^{5} f_{1j}^{(n)}) \{U(\gamma^5 I u)\} \Bigg\} \end{aligned} \tag{5.7}$$

where $0 < \gamma < 1$ is a discount coefficient capturing the time value of money.

Remark: We are aware of the fact that it is technically possible for a negative wealth (argument of U) to occur, i.e., when the firm conceals all of its profit and is audited 5 years later with a high penalty factor β. A negative wealth value would be problematic in the context of the utility function chosen. To circumvent this problem, we make the — quite reasonable — assumption that positive wealth will always be preferred over negative wealth, and thus the firm would never consider values of u leading to negative wealth. After applying this restriction, that the maximum in Eq. 5.7 will always occur at some positive wealth value because there is always a u that produces positive wealth (e.g, $u = 0$ - declare all and keep the after-tax profit).

5.2 Assumptions and Parameter Selection

As in Section 3.5, we will assume that the firm's annual profit, I, is constant throughout its economic life. Regarding the choices of I and u, we will again refer to "relative" amounts, so that, for example, $I = 100$, and u is the *percentage* of current profit to be hidden from the authorities. The fraction of the profits that a firm can choose to hide can take values $u \in [0, 1]$. As before the audit probability will be 0.0025 for the firm to be audited in years 1-4 since its last audit, and 0.04 if it has not been audited for at least five years.

5.3 Risk-neutrality

Starting with a risk-neutral firm ($\lambda = 0$), we mapped the level of tax evasion vs. tax penalty and audit probability, identifying the (β, p) combinations that lead to full profit disclosure (see Figure 5-2). We observe that tax evasion cannot be eliminated, even with tax penalties of 10 times the tax owed, unless the audit probability is increased above 5%. Moreover, no matter the audit probability, the tax penalty should be increased to at least 1.5 instead of the current 0.24, in order for tax evasion to begin to decrease.

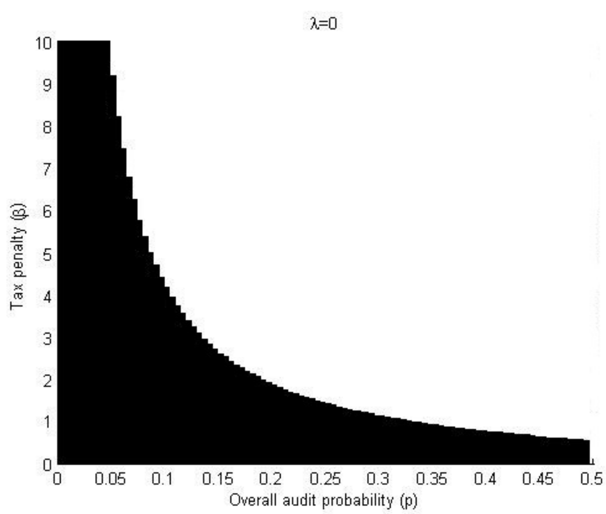

Figure 5-2: The level of tax evasion as a fraction of a risk-neutral firm's profits.

We also note in Figure 5-2 that under the current combination of audit probability (5%) and tax penalty $\beta = 0.24$, a typical risk neutral firm will choose tax evasion.

5.4 Risk-aversion

We estimated that if the current level of Greek "hidden" economy persists (40% based on [102]), the risk aversion level of a typical Greek firm should be approximately $\lambda = 6$, because it is for that value of the risk aversion coefficient in Eq. 5.4 that the current tax penalty coefficient and overall audit probability produce tax evasion of approximately 40%. In the following, we discuss the behavior of firms with $\lambda = 6$, as well as firms which are less risk-averse, i.e $\lambda = 3$.

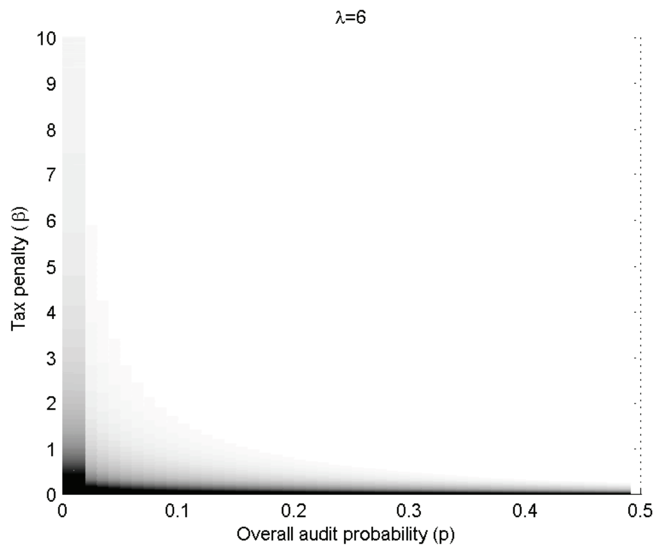

Figure 5-3: Firm's decision on the tax penalty - audit probability space. Black: 100% tax evasion ($u = 1$); white: full profit disclosure ($u = 0$); Gray levels: intermediate values of u.

The optimal level of tax evasion for a risk-averse firm with $\lambda = 6$ is shown in Figure 5-3, where gray levels represent the degree to which the firm chooses to disclose its profits. In that setting, tax evasion begins to decrease when the tax penalty reaches 34%, at an audit probability of only 1%. It is clear from Figure 5-3 that raising the audit probability is a much more effective deterrent compared to increasing penalties. These results are shown in Figure 5-3. Furthermore, under the current levels of tax penalty and audit probability, the level of tax evasion is approximately 41.35%, which agrees with the work in [102].

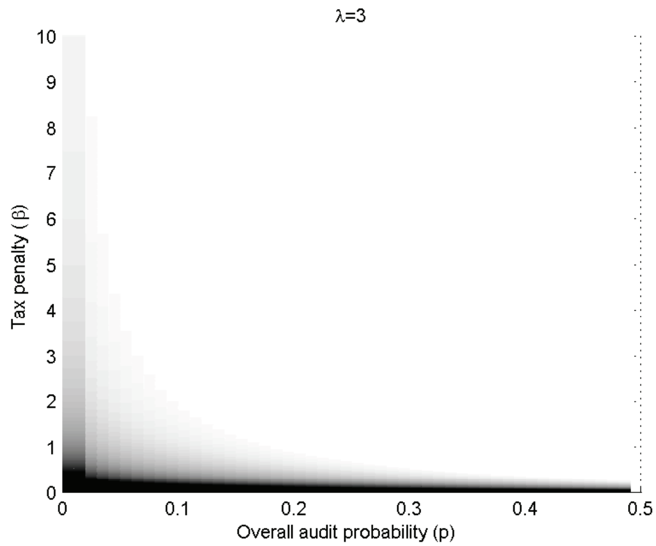

Figure 5-4: Tax evasion mapping for $\lambda = 3$. The black area represents full tax evasion and white area full profit disclosure. Gray levels represent intermediate choices.

Figure 5-4 illustrates the effectiveness of tax penalty increases for firms which are slightly risk-averse ($\lambda = 3$). It is still not possible to completely eliminate tax evasion using "reasonable" penalty factors when the audit probability is lower than 0.05 (however tax evasion is being limited at a tax penalty lower than 0.3, which is quite reasonable. For the current levels of $\beta = 0.24$ and $p = 0.05$ the firm would still choose to hide its profits as much as possible (i.e $u = 1$).

5.5 Government Revenues

We used our model to compute the expected government revenues under different assumptions concerning the firm's risk aversion. If, as the literature suggests, tax evasion in Greece is between 25% and 40%, then, based on our model, the firm's risk aversion coefficient must be between $\lambda = 6$ and $\lambda = 12$. Figures 5-5 and 5-6 show the corresponding government revenues per firm, as a fraction of the firm's net profits.

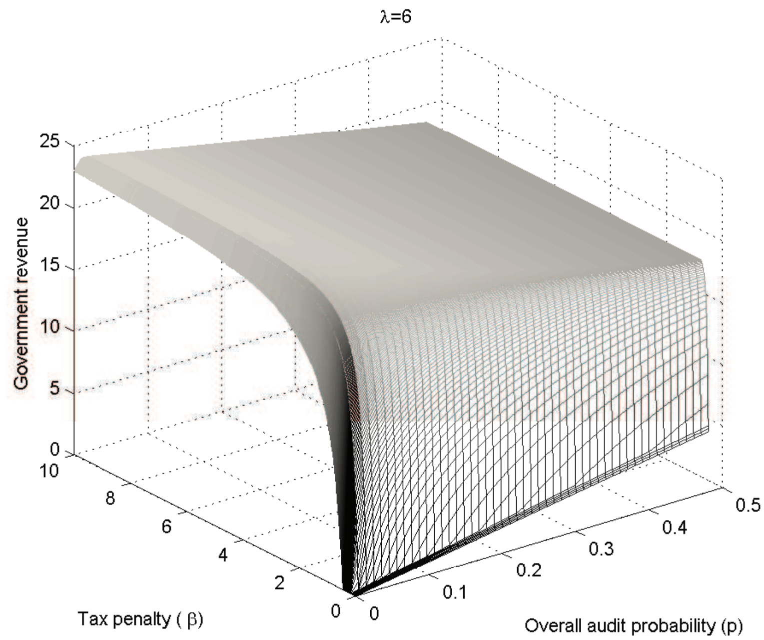

Figure 5-5: $\lambda = 6$. Government revenues as a function of audit probability and tax penalty coefficient.

We observe that the maximum expected revenue of the government cannot exceed the tax rate of 24%. This can be attributed to the fact that the best-case scenario for the government is for the firm to disclose the full amount of its profits and thus pay its taxes as it should. Consequently, the combination of tax penalty coefficient and audit probability has to yield an expected tax expense that does not exceed the normal amount of taxes due, assuming that the firm disclosed its profits fully. Moreover, there is a positive relationship between the tax penalty and government revenues, either by collecting penalties for tax evasion or by preventing tax evasion in the first place. For low audit probabilities the effect of the tax penalty is limited.

The differences in tax revenues between the $\lambda = 6$ and $\lambda = 12$ cases are highlighted in Figure 5-7, with a different view of the same surface shown in Figure 5-8. In those plots, we have zoomed on the area of tax penalties between 0% and 20% and audit probabilities between 0% and 10% in order to show more clearly the effect of these two parameters.

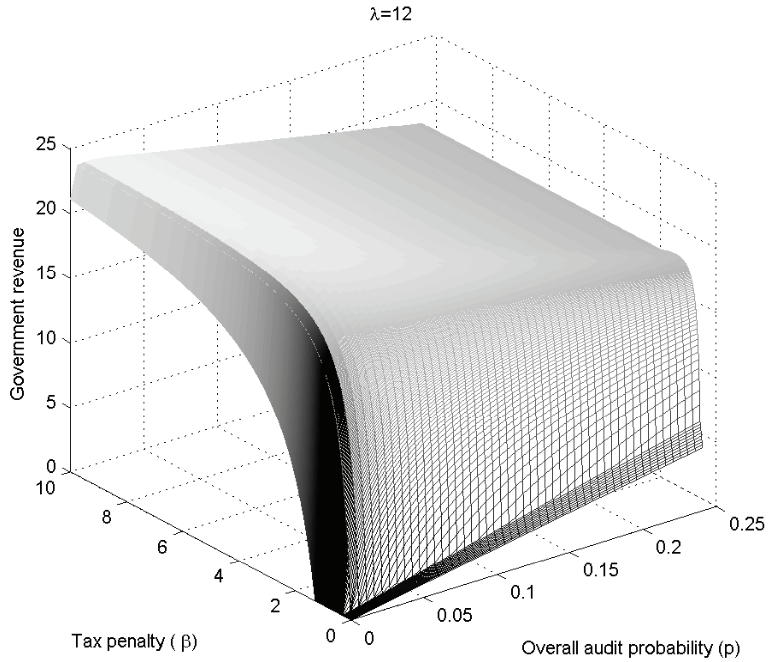

Figure 5-6: $\lambda = 12$. Government revenues as a function of the audit probability and tax penalty coefficient.

We note that, as risk aversion increases, government revenues increase as well in the region of tax penalties and audit probabilities that lead to tax evasion. When the audit probability or the tax penalty are near zero, there is no significant difference in government revenues for the two cases. It is also interesting to note the sharp, downward "fold" present in Figures 5-7 and 5-8 which shows that government misses out on the most tax revenue (from the less- vs. the more risk-averse firms) for a range of tax penalty β between 0.03 and 0.12, which (unfortunately) is close to the current penalty factor.

We have argued that according to the estimated current levels of tax evasion in Greece, the risk aversion coefficient was estimated to be approximately $\lambda = 6$. Table 5.1 presents the level of tax evasion on a ceteris paribus basis, by keeping the tax penalty stable at its current level of 0.24 and gradually increasing the audit probability. This facilitates the comparison of the relative effectiveness of the audit probability and tax penalty deterrents. One conclusion that can be safely drawn from

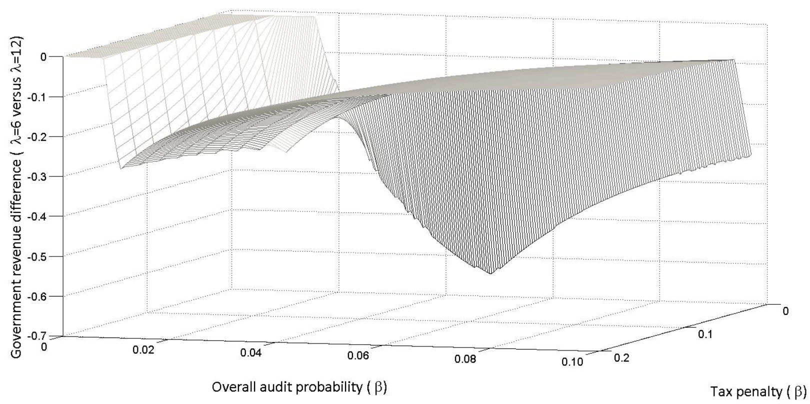

Figure 5-7: Government revenue collected from a firm with risk aversion $\lambda = 6$ minus that collected for $\lambda = 12$.

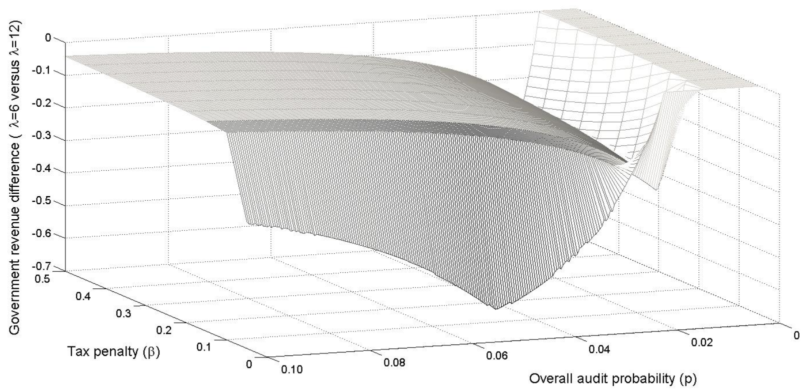

Figure 5-8: Difference in government revenues collected from a firm with risk aversion $\lambda = 6$ vs. those collected from a firm with $\lambda = 12$. This surface is the same as that shown in Figure 5-7, viewed from a different angle.

tables 5.1 and 5.2 is that both audit probabilities and tax penalties are relevant tax deterrent tools. Their relevance though does not remain constant across the board. It seems that government can restrain tax evasion more effectively through an efficient audit mechanism than by using higher tax penalties.

$\lambda = 6$, $b = 0.24$

p	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	0.10
u	1	0.7571	0.6511	0.5861	0.5384	0.5005	0.4689	0.4417	0.4177	0.3962
p	0.11	0.12	0.13	0.14	0.15	0.16	0.17	0.18	0.19	0.20
u	0.3767	0.3588	0.3422	0.3267	0.3122	0.2986	0.2856	0.2733	0.2616	0.2503

Table 5.1: Level of tax evasion ($u \in [0, 1]$) on behalf of a typical risk-averse firm ($\lambda = 6$). The tax penalty is fixed at 0.24, while the audit probability p varies between 0.01 and 0.20

$\lambda = 6$, $p = 0.05$

b	0	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90
u	1	1	0.5964	0.3964	0.2914	0.2300	0.1754	0.1466	0.1222	0.1051
b	1.00	1.10	1.20	1.30	1.40	1.50	1.60	1.70	1.80	1.90
u	0.0899	0.0776	0.6666	0.0584	0.0514	0.0455	0.0404	0.0359	0.0320	0.0286

Table 5.2: Level of tax evasion, ($u \in [0, 1]$), on behalf of a typical risk-averse firm ($\lambda = 6$), keeping the audit probability at 0.05 for different values of the tax penalty ($\beta \in [0, 1.9]$). In this setting, tax penalty coefficient above 2 had a negligible effect on tax evasion.

Chapter 6

Risk-averse firms with optional closure

In this section we revisit the model developed in Chapter 3. Its description will be given here again in as compact a form as possible, for the sake of readability and to facilitate the discussion. We use the following notation: The integer $k = 0, 1, 2, ...$ will denote discrete time, and x_k will be the value of the quantity x at time k. Elements of vector x, or matrix M, will be indicated by $[x]_i$ and $[M]_{ij}$, respectively and $0_{i \times j}$ will denote an i-by-j matrix of zeros.

The tax status of the firm will again be $s_k \in \mathcal{S}$, where

$$\mathcal{S} = \{V_1, ..., V_5, O_1, ..., O_5, N_1, ..., N_5\}, \tag{6.1}$$

The elements of $\mathcal{S}$ represent 15 discrete statuses that are possible in the Greek taxation system. When in V_i, the firm is being audited and its last five income reports are verified; while in O_i, the firm has just used the closure option and has not employed it, nor has it been audited in the past $i = 1, ..., 5$ years. Finally, when in N_i, the firm's last audit or closure was $i = 1, ..., 5$ years ago. Thus, the firm has i unaudited tax years and will now make its $(i+1)$-st consecutive decision since its last audit/closure. Each year the firm decides the fraction of its profit to conceal in the interval $[0, 1]$.

The availability of the closure option in year k (as decided by the government) is given by $c_k \in \mathcal{C}$

$$\mathcal{C} = \{option\ available, \quad option\ not\ available\} \tag{6.2}$$

In some cases, we will refer to elements of $\mathcal{S}$ and $\mathcal{C}$ by integer, in their order of appearance, i.e., $V_1 \to 1$, $V_2 \to 2$,...,$N_5 \to 15$ for states in $\mathcal{S}$, and *option available* $\to 1$, *option not available* $\to 2$ for $\mathcal{C}$.

The firm's state vector at time k is

$$x_k = [s_k, c_k, h_k^T]^T, \tag{6.3}$$

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

6.1 State Evolution

The system's dynamics are given by

$$x_{k+1} = Ax_k + Bu_k + n_k, \quad x(0) \text{ given,} \tag{6.4}$$

where

$$A = \begin{bmatrix} 0 & & \\ & 0 & \\ & & H \end{bmatrix}, \quad H = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad n_k = \begin{bmatrix} w_k \\ \epsilon_k \\ 0_{5\times 1} \end{bmatrix} \tag{6.5}$$

Where w_k and ϵ_k are random processes which correspond to government's actions as described next. We assume that the government offers the closure option with a fixed probability, p_o, each year

$$Pr(\epsilon_k = i) = \begin{cases} p_o & \text{if } i = 1 \quad \text{(option available)} \\ 1 - p_o & \text{if } i = 2 \quad \text{(option not available)} \end{cases} \tag{6.6}$$

The term w_k in Eq. 6.5 determines the first element of the state vector (i.e., the firm's "next" Markov state in $\mathcal{S}$). The random transitions that the firm undergoes in $\mathcal{S}$ depend on its existing state as well as its decision to accept or reject the option, $[u_k]_2$ (if it is offered):

$$Pr\left(w_k = i | x_k = [j, q, h_k^T]^T, [u_k]_2 = m\right) = P_{qij}(m), \quad i, j \in \{1, ..., 15\}, q \in \{1, 2\} \tag{6.7}$$

where, for q and u fixed, $P_{qij}(m)$ forms a Markov matrix governing the firm's transitions in $\mathcal{S}$:

$$P_{qij}(m) = \begin{cases} [M_{no}]_{ij} & \text{if } q = 2, \ \forall\, m \quad \text{(no option)} \\ [M_a]_{ij} & \text{if } q = 1, \ m = 1 \quad \text{(option taken)} \\ [M_d]_{ij} & \text{if } q = 1, \ m = 2 \quad \text{(option declined)} \end{cases} \tag{6.8}$$

and the M_{no},M_a,M_o are given in Appendix A.

6.2 Rewards and Value Function

The reward the firm enjoys for being in state x_k and following a policy u_k is given by

$$g(x_k, u_k) = g\left([s_k, c_k, h_k^T]^T, u_k\right) = \begin{cases} U\left(R \cdot (1 - r + r[u_k]_1)\right) & s_k \in \{11, ..., 15\} \\ U\left(R \cdot (1 - r + r[u_k]_1 - \ell(s_k - 5))\right) & s_k \in \{6, ..., 10\} \\ U\left(R \cdot (1 - r + r[u_k]_1 - r \sum_{i=1}^{s_k} [h_k]_{6-i} \right. & \\ \left. - \frac{3}{5}\beta r \sum_{i=1}^{s_k} i [h_k]_{6-i})\right) & s_k \in \{1, ..., 5\} \end{cases} \tag{6.9}$$

Where, R denotes the firm's annual revenues, r is the tax rate, ℓ the closure cost, β the tax penalty. $U(\cdot)$ is a constant relative risk aversion utility function,

$$U(g) = \frac{g^{1-\lambda}}{1-\lambda}, \tag{6.10}$$

where λ is the risk-aversion coefficient.

The first term in the right hand of Eq. 6.9 is the reward the firm gets when it conceals $R[u_k]_1$ of its revenue. The next term is the reward in case the firm uses the option of closure ($[u_k]_2 = 1$) and pays ℓ per year since its last closure or audit, up to a maximum of five years. Finally, the bottom term is the reward in case of an audit. The audit takes into account the firm's history since the last audit or closure. Note that the reward function is no longer linear as in [26]. This poses a significant challenge when it comes to solving the model, as we will discuss later.

The firm acts as if its lifespan is infinite and chooses its policy u_k so as to maximize the discounted expected reward:

$$\max_{u_k} \mathop{\mathbb{E}}_{w_k,\epsilon_k} \left\{ \sum_{k=0}^{\infty} \gamma^k U(g(x_k, u_k)) \right\}, \tag{6.11}$$

Where $\gamma \in (0, 1]$ denotes the discount factor. We must solve the stationary Bellman equation.

$$J_\infty(x) = \max_u \left\{ U(g(x,u)) + \gamma \mathop{\mathbb{E}}_{w,\epsilon} J_\infty(Ax + Bu + n) \right\} \tag{6.12}$$

Where the expectation operator of Eq. 6.12 is to be taken with respect to the distributions of w and ϵ. Using the state equation (Eq. 6.4) together with Eq. 6.6 and Eq. 6.7, we see that computing Eq. 6.12 is equivalent to computing

$$J_\infty(i,q,h) = \max_u \left\{ U(g(i,q,h,u)) + \gamma \sum_{t=1}^{2} \sum_{j=1}^{15} P_{qji}([u]_2)\, Pr(\epsilon = t) J_\infty(j, t, Hh + e_5[u]_1) \right\}, \tag{6.13}$$

where $e_5 = [0,0,0,0,1]^T$ and, for convenience, we have slightly abused the notation by writing the argument of J_∞ as (i,q,h) instead of $x = [i,q,h^T]^T$, and that of g as (i,q,h,u) instead of $x = ([i,q,h^T]^T, u)$, with $i = 1, ..., |\mathcal{S}|$, $q = 1, ..., |\mathcal{C}|$, $h \in [0,1]^5$.

6.3 Solving for the optimal policy

As we discussed before, the element $[u_k]_1 \in [0,1]$ of the control vector u denotes tax-evasion as percentage of the firm's annual revenues. It is obvious that the continuous nature of the control vector makes the state vector, x, continuous as well, because the firm's last five decisions are incorporated into the state. In Chapter 3, we circumvented the continuous nature of both state and control vectors, using the linearity of the reward function, and successfully applied DP after determining that $[u_k]_1$ should only take a bang-bang form. This time however, the cost-to-go function (Eq. 6.9) is non-linear. As a result, we can no longer assume that the optimal control is bang-bang and it is computationally difficult to apply DP as we were able to do in Chapter 3.

In order to circumvent the problem associated with the continuity of the control and state spaces, we will discretize $[u_k]_1 \in [0, 0.01, \ldots, 0.99, 1]$, assuming that tax evasion takes place in increments of 1%. However, even after discretizing the control and state spaces, (x, u), the number of states remains too large for DP to be applied effectively. More specifically, we are left with between $23 \cdot 2 \cdot 101^5$ and $23 \cdot 2 \cdot 202^5$ potential states (the number of the elements of the state vector x_k, times the number of control elements of u_k times all possible combinations of control for the past five years). It becomes clear that DP cannot be used because: i) the number of states is too large for DP to "visit" each state and update the value function J_k according to Eq. 6.13, and ii) it difficult even to store the updated values in tabular form. One way to go forward is to combine: i) an approximation method to estimate the value function J_k and ii) an approximate way of storing the optimal values of J_k, based on the optimal policy. To address the first issue we will use Approximate Dynamic Programming (ADP) and specifically SARSA(0), as described in [104]; for the second problem a Feed-Forward Neural Network will do the job, as we will discuss shortly.

6.3.1 Approximate Dynamic Programming

As per [104], we aim to estimate the value function for states, x_k, or for state-action pairs, (x_k, u_k), in order to see "how good" is a given state x_k in terms of expected reward. Future rewards depend upon the future controls (or actions) the firm applies. Consequently, the value function must somehow be estimated taking into account those the future controls. A policy, π, is a mapping from each state $x \in \mathcal{X}$ and control $u \in \mathcal{A}(x)$ to a control u when in state x, where $\mathcal{X}$ denotes the set of states and $\mathcal{A}(x)$ the set of controls when in state x. The value of a state x under policy π, denoted by $V^\pi(x)$, is the expected reward when starting at x and following π thereafter, and is given by

$$V^\pi(x) = \mathbb{E}_\pi\{G_k|x_k = x\} = \mathbb{E}_\pi\{\sum_{k=0}^{\infty} g_{k+1}|x_k = x\}, \tag{6.14}$$

where $\mathbb{E}\{\cdot\}$ denotes the expected value provided that the firm follows π. V^π is called the *state-value function of policy* π. We let $Q^\pi(x, u)$ denote the expected reward starting from x and using control u while following policy π thereafter:

$$Q^\pi(x, u) = \mathbb{E}_\pi\{G_k|x_k = x, u_k = u\} = \mathbb{E}_\pi\{\sum_{k=0}^{\infty} g_{k+1}|x_k = x, u_k = u\}, \tag{6.15}$$

The term $Q^\pi(x, u)$ is called *action-value function of policy* π and can be estimated from "experience": The firm follows π and keeps averaging, for each state-action pair that is visited, the actual returns that occurred every time the system visits that specific state-action pair. As the sample size approaches infinity, the average action-value converges to the real action-value, [104]. This method belongs to a group called *Monte-Carlo methods* and requires large samples of action-values to approximate the actual action-value of a state-control pair.

We will concentrate first on estimating the action-value $Q^\pi(x, u)$. This can be achieved by generating so-called "episodes", which are essentially simulations of the system's state evolution [104]. In order to estimate the action value, instead of state to state transitions, we must consider transitions from one state-control pair to the next, with the dynamics of these transitions being governed by Eq. 6.4. Over the course of the episode, the action values are updated according to:

$$Q(x_k, u_k) \leftarrow Q(x_k, u_k) + a[U(g_{k+1}) + \gamma Q(x_{k+1}, u_{k+1}) - Q(x_k, u_k)]. \tag{6.16}$$

This update is performed after every transition from a state-control pair to the subsequent one. This method is called "on-policy", because it estimates the action-value Q^{π} for the chosen policy π, and simultaneously improves Q^{π}. The general form of the SARSA(0) algorithm, [104], is:

Initialize all $Q(x, u)$ arbitrarily , set $\epsilon > 0$ "small" ;

repeat

  for each episode;

  Choose an initial state $x = x_0$ arbitrarily;

  With probability $1 - \epsilon$, chose the u available in x, that maximizes $Q(x, u)$;

  or, with probability ϵ, chose a random u;

  for *each step of the episode* **do**

    Apply u and observe the reward g and the next state x';

    Choose the next control u' from those available in state x', with probability $1 - \epsilon$;

    or, with probability ϵ, chose a random u';

    $Q(x, u) \leftarrow Q(x, u) + a[g + \gamma Q(x', u') - Q(x, u)]$ $x \leftarrow x'$; $u \leftarrow u'$;

  end

until x' *is terminal*;

Algorithm 1: SARSA(0)

In SARSA(0), x denotes the current state, u the current control, x' the subsequent state using control u from state x, and u' the control chosen in pair with x'. Note that when at state x, the control is chosen based on a $\epsilon - greedy$ policy. More specifically, with probability $1 - \epsilon$, the algorithm chooses the control that maximizes $Q(x, u)$. However, with probability ϵ, the algorithm chooses a random control to encourage exploration of the control-state space.

SARSA(0) is a form of ADP that attempts to circumvent the problem of having a large number of states. We now return to the problem of storing the values $Q(x, u)$. As we have already discussed, a tabular form is not appropriate because of space considerations. Alternatively, we can consider a function approximator. In [104] a number function approximation methods are proposed (i.e., Coarse Coding, Tile coding or Kanerva Coding). In this work we have opted for a Feed-Forward Neural Network (NN).

6.3.2 Function approximation via Neural Network

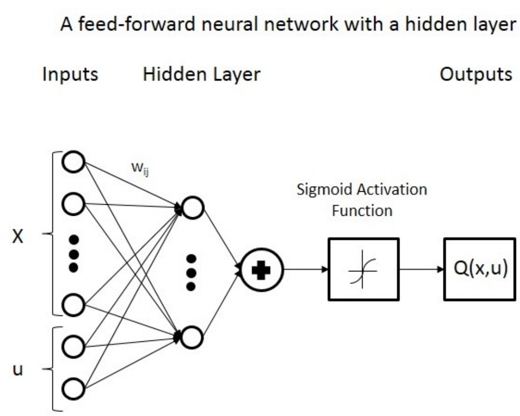

Figure 6-1: Schematic showing the architecture of the feed-forward neural network used in to store the values of $Q(x, u)$.

Consider a typical Feed-Forward NN, as shown in Figure 6-1. We will use the NN network to map state-control pairs to action-values $Q(x, u)$. The inputs of the NN are the state vector x_k, as in Eq. 6.3, and the control vector u_k. The output of the NN is (an approximation of) $Q(x_k, u_k)$ and there is a hidden layer with 120 neurons, this size having been arrived at after numerical experimentation, aimed at improving both computational efficacy and convergence rate. One could train a single such NN assuming that all possible controls are possible inputs. That would increase the number of inputs to 225 (the 23 elements of the state vector plus all 202 available controls, given the state x_k). However, based on [105], we decided to instead train

202 different networks, all identical in architecture, one for each possible control. By keeping the control u fixed for each network, we avoid provoking any unnecessary changes among its weights, caused by valuating other controls. In the next Chapter we present the results obtained using the ADP approach.

Chapter 7

Running the complete model: Results and Discussion

In this Section we present the numerical results obtained by solving the full model (including risk-aversion and closure) for the firm's optimal decisions. The results are separated into two categories. First, we assume risk-neutrality on behalf of the firm in order to validate the performance of our approach against the results in Chapter 4. We then proceed with the introduction of risk-aversion and estimate the risk-aversion level of a typical Greek firm assuming that the prevailing tax evasion level in Greece is approximately 40%, based on [102]. In Chapter 3 we discussed in detail our choices for the model's various parameter values of the parameters. To facilitate comparisons between the ADP and DP approaches, we kept the parameter values unchanged.

7.1 Risk-Neutral Firms

Figure 7-1 shows the results of the training process (see Chapter 6) for the 101 independent NNs, each corresponding to each of the available controls $[u]_1 \in \{0, 0.01, 0.02 \ldots 0.99, 1\}$. Note that the total available controls are not 202 because the option is not available ($\epsilon = 1$ and consequently $[u]_2 = 1$) in this scenario.

We observe a strong relationship between tax evasion, as a percentage of annual revenues, and long-term tax revenues collected. The (approximately) optimal control was computed as the average of the last 300 NN outputs which were produced during the learning process, after those outputs exhibited convergence. However, it is unclear which control u is optimal at higher levels of tax evasion (i.e., near control 101, or equivalently, $[u]_1 = 1$). The reason is that as tax evasion increases, the volatility of the estimated long-term revenues increases as well. The volatility of the NNs' output is shown in Table 7.1 and can be safely attributed to the increased penalization of tax evasion (there is a positive relationship between tax evasion and tax penalties) and the stochastic nature of audits.

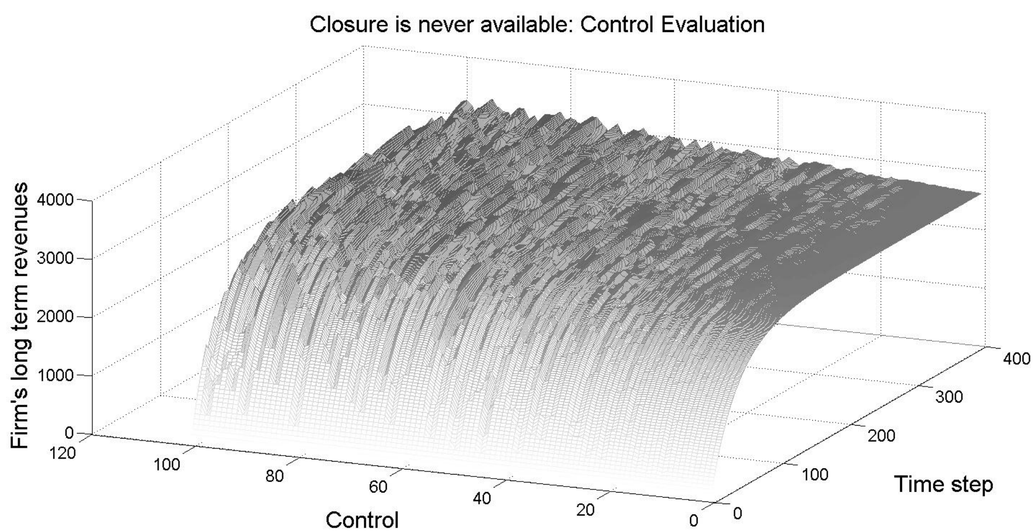

Figure 7-1: Control evaluation using 101 independent neural networks (horizontal axis) for each control ($[u_k]_1 = 0, 0.01, 0, 02 \ldots, 1$) when risk neutrality is assumed and closure is never available.

The volatility of the NNs' outputs is evident in Figure 7-2, where, for each NN (equivalently, each control input) we have plotted the time average of that NN's outputs during the last 300 steps of the learning process.. As we have already discussed in section 6.3.2, independent NNs were dedicated to learning the long-term impact of each available policy on the firm's revenues. To achieve precision in the policy iteration process, instead of using a single NN's output (e.g., the last one in each training episode), we used the average of the last 300 NN's outputs after the learning algorithm had converged. Consequently, Figure 7-2 shows that, in the long-run, the best policy for the firm to use, regarding tax evasion, is to conceal as much profits

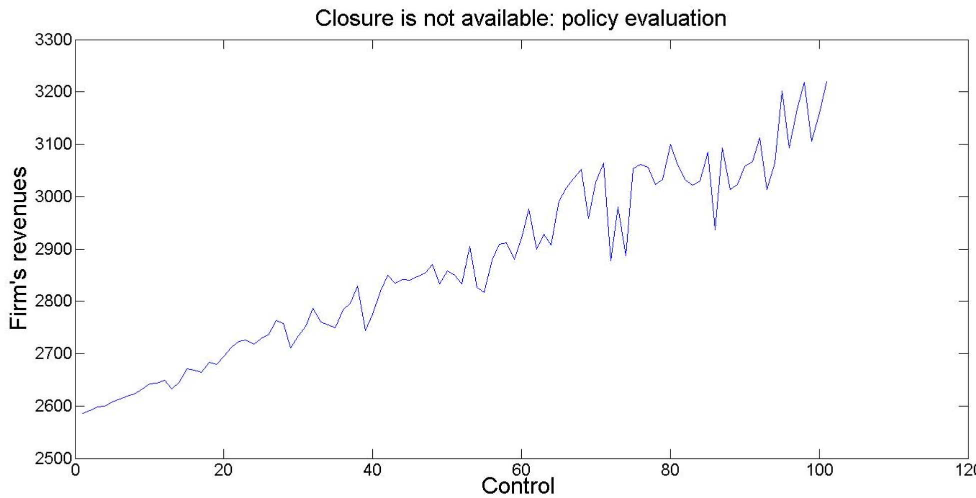

Figure 7-2: Average NN outputs (firm revenue) for the 101 NNs (one per control input) during the last 300 steps of the learning process. Firm is assumed to be risk neutral and closure is never available. One NN is dedicated to learning the long-term impact of each available policy ($[u_k]_1 = 0, 0.1, 0.2 \ldots, 1$), indexed from 0 to 101 in the horizontal axis) on firm's revenue.

as possible. The optimal policy's evolution begins at zero tax evasion because the learning process chooses an initial control arbitrarily, in our case $[u]_1 = 0$. The initial choice of control does not affect the learning process, as long as there are enough steps in the episode for the algorithm to converge, as in [104].

7.1.1 Risk-neutral firms with the option available annually

Figure 7-3 shows the results of learning process for all NNs, when closure is available annually. This hypothesis is rather realistic because the Greek government uses this taxation scheme for certain classes of freelance professionals (e.g., lawyers, engineers and doctors). The availability of closure means that the number of available controls increases to 202. Recall that for the first 101 controls, it is assumed that the closure is not used when available ($[u]_2 = 1$). For the remaining 101 controls, the closure *is* used when available.

It is evident from Figure 7-3, that the long-term rewards learned by the NNs for the first 101 controls where closure is not employed, remain volatile. The volatility of the long-term revenues is increased when compared with the case of closure not being

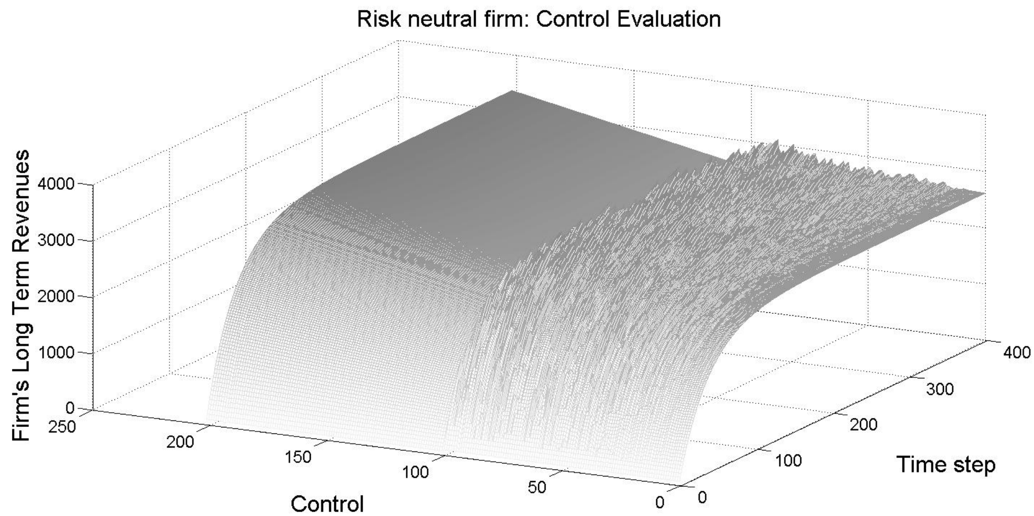

Figure 7-3: Control evaluation using independent neural networks, one for each control ($[u_k]_1 = 0, 0.01, 0, 02 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), when risk neutrality is assumed and closure is available annually.

available at all (Figure 7-1). This higher volatility can be attributed to the fact that the probability of an audit is tripled when closure is available and the firm chooses not to use it. On the other hand, for the last 101 controls (closure is available and the firm uses it), the firm always pays the cost of closure, $[u]_2 = 2$, instead of the normal taxes. In return, the probability of an audit becomes zero. Thus, any stochasticity is eliminated, and the NN estimates the long-term revenues for each control with precision. The results are in line with those in [26]. The firm's (optimal) behaviour can be attributed to three factors: i) the low cost of closure (ℓ), ii) the low tax penalty (β) and iii) the low audit probability.

Figure 7-4 illustrates the average output of each NN (equivalently, each available control input) over the last 300 steps of training. The volatility of the NNs' outputs when the firm does not use the option is evidently large (see also Table 7.1). The figure also shows that the firm's long-term revenues are maximal when closure is available and the firm decides to use it. We can thus argue that in this scenario it is not in the government's best interest to incorporate closure into the taxation system.

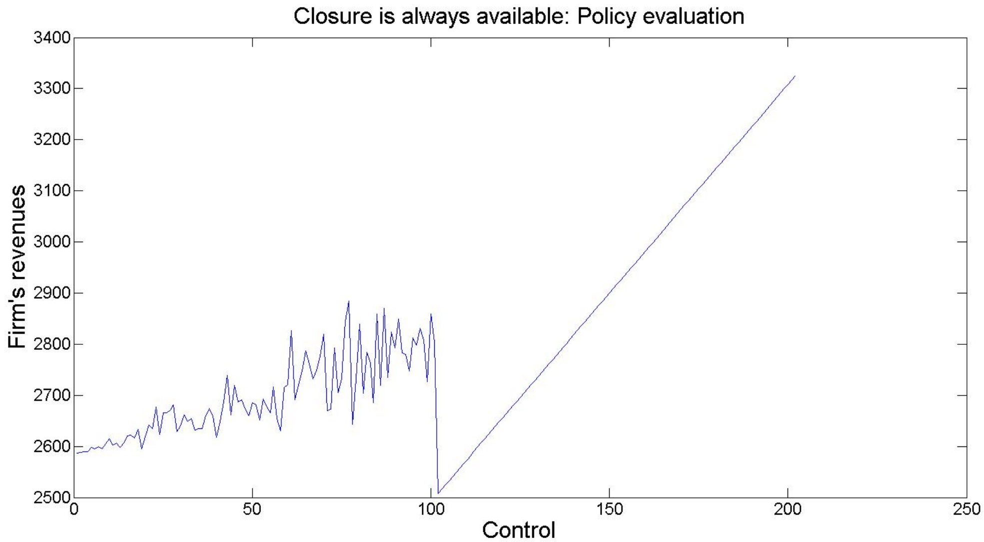

Figure 7-4: Average NN outputs (firm revenue) for the 202 NNs (one per control input) during the last 300 steps of the learning process. Firm is assumed to be risk neutral and closure is always available. One NN is dedicated to learning the long-term impact of each available policy ($[u_k]_1 = 0, 0.1, 0.2 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), indexed from 0 to 202 in the horizontal axis) on firm's revenue.

7.1.2 Risk-neutral firms with stochastically available closure

We continue with the scenario of a stochastically available closure option. The setting is motivated by the fact that the Greek government does not *a priori* disclose its intention to activate the option. Statistically, the probability of the option to become available in any given year is approximately 20%. The estimated long-term revenues obtained by training the NNs corresponding to the 202 available policies are shown in Figure 7-5. Although it is difficult to identify the optimal control at the end of each episode, we can still draw some conclusions from the graph. First, the volatility of long-term revenues produced by the set of the 101 controls which discards closure option when available, is noticeably higher when compared with the set that employs closure when available. The volatility of the long-term rewards, based on the optimal control, for different scenarios of closure availability is shown in Table 7.1.

Figure 7-6 shows the long-term tax revenue obtained using each available policy. Concerning $[u_k]_1$, the firm always conceals as much revenue as possible from the government. The reasoning behind this strategy can be found in $[u_k]_2$ because it is

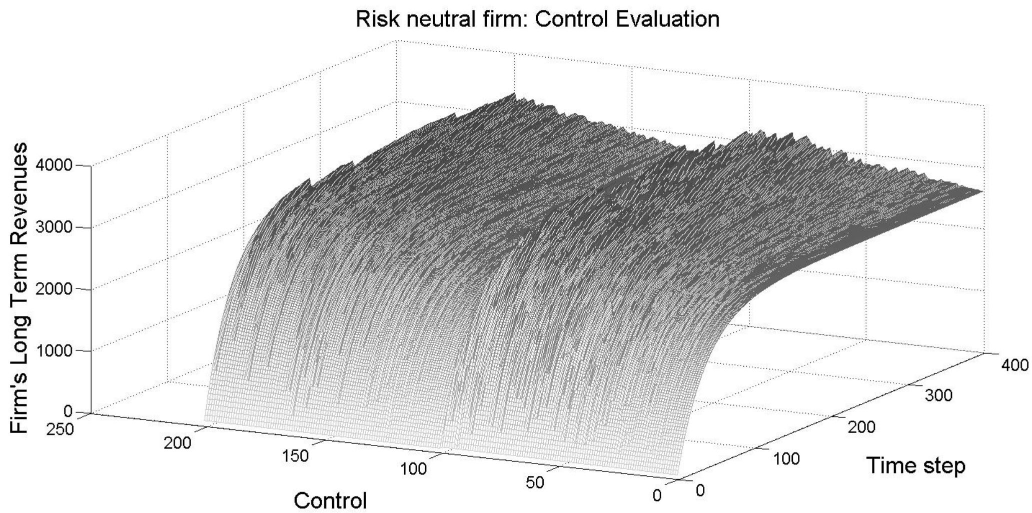

Figure 7-5: Control evaluation using independent neural networks, one for each control ($[u_k]_1 = 0, 0.01, 0, 02 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), when risk neutrality is assumed and closure is stochastically available.

best for the firm to always use the closure option when it becomes available. This strategy enables the firm to capitalize the revenues earned via tax evasion; the low cost of closure and probability of a potential audit make the risk "affordable" and even profitable, as we will discuss later. Consequently, closure encourages tax evasion, even when its availability is not known *a priori*. Furthermore, if we consider the volatility of the NNs' outputs as a proxy for the risk that the firm undertakes when closure is available and it decides to tax evade, it is evident that this risk increases, when closure is not used by the firm when available. This can be attributed to the increased audit probability the firm faces when the option is available but is declined, and to the fact that when closure is used the firm cannot be audited for any past decisions.

7.1.3 Risk-neutral firms with periodically available option

In this final scenario we examine the effects of the periodic availability of closure every five years (to coincide with the expiration of the statute of limitations on past tax declarations). This assumption is based on the fact that the Greek government has, in the past, offered the option in this manner, which can also be perceived as an admittance of its inability to battle tax evasion and as an aim to increase collected

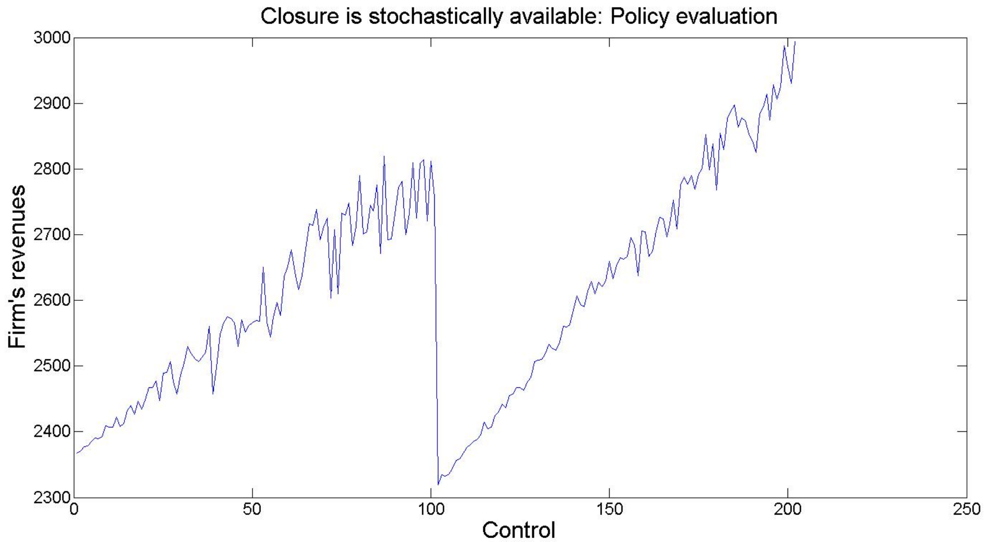

Figure 7-6: Average NN outputs (firm revenue) for the 202 NNs (one per control input) during the last 300 steps of the learning process. Firm is assumed to be risk neutral and closure is stochastically available. One NN is dedicated to learning the long-term impact of each available policy ($[u_k]_1 = 0, 0.1, 0.2 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), indexed from 0 to 202 in the horizontal axis) on firm's revenue.

tax-revenues in the short-term. Looking at Figure 7-7, it is evident that closure significantly limits the volatility of the NNs' outputs (long term revenues). Moreover, it is also interesting to ask in which scenario, between closure being stochastically versus periodically available, is the volatility of the NNs' long-term revenue outputs lower. We expect the volatility to be lower for periodically available closure option because that setting is deterministic. Naturally, the question that arises is: should the Greek government continue to use closure in the Greek taxation mixture. We will have more to say about this later in this section.

Figure 7-8 shows the long-term impact of each available policy on the firm's revenues ($[u_k]_1 = 0, 0.01, 0, 02 \ldots, 1$ and $[u_k]_2 \in [1, 2]$)for closure being available or not, which appear indexed from 0 to 202 in the horizontal axis of Figure 7-8). The optimal policy for the firm remains the maximum concealment of its revenues from the government and to use the closure periodically. The argument that closure encourages tax evasion seems to be strengthened because closure allows the firm to conceal as much revenue as possible and then use the option to avoid being audited, in exchange for a

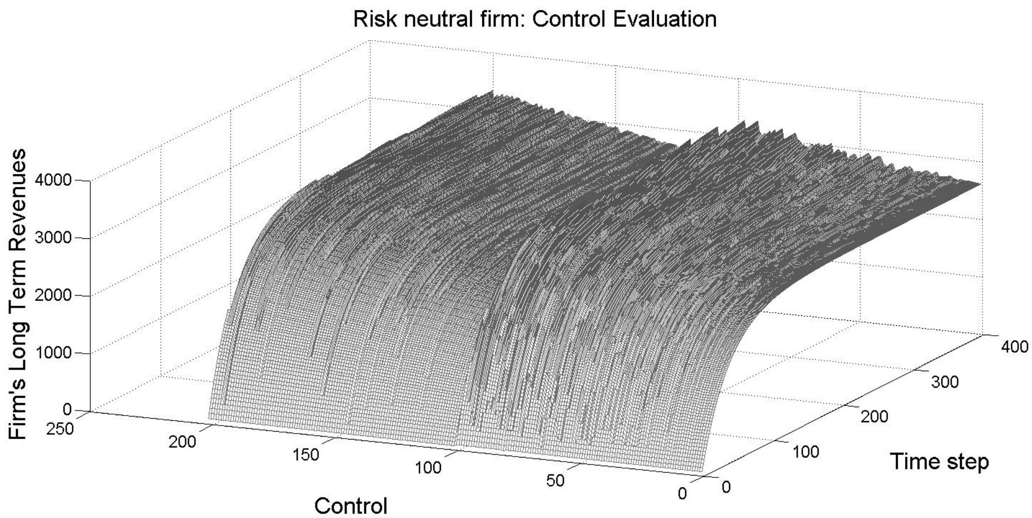

Figure 7-7: Control evaluation using independent neural networks, one for each control ($[u_k]_1 = 0, 0.01, 0, 02 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), when risk neutrality is assumed and closure is periodically available.

low cost. It is interesting to compare the evaluation of each available policy presented in Figure 7-8 with those in Figure 7-6 (when closure is stochastically available). It is evident that when closure availability is deterministic the volatility firm revenues decreases (Table 7.1), which implies that the firm faces a lower risk. Even though the difference in firm revenues in these two scenarios is small the risk the firm faces when it engages in tax evasion is considerably lower.

Looking at Table 7.1, one could argue that the option increases long-term revenues and provides safety to the firm. When closure is always available the firm can benefit in two ways: first it can eliminate the risk of a potential audit by paying a lump sum, and while this sum is lower than the normal tax, the firm will always opt for it. We notice that when closure is always available, and the firm uses it, the standard deviation of its revenues is 51.73. On the other hand, when closure is stochastically available, the standard deviation increases to 92.80 because there are two stochastic events that the firm needs to take into consideration: a potential audit, and closure availability. The nonzero standard deviation in the case of closure being

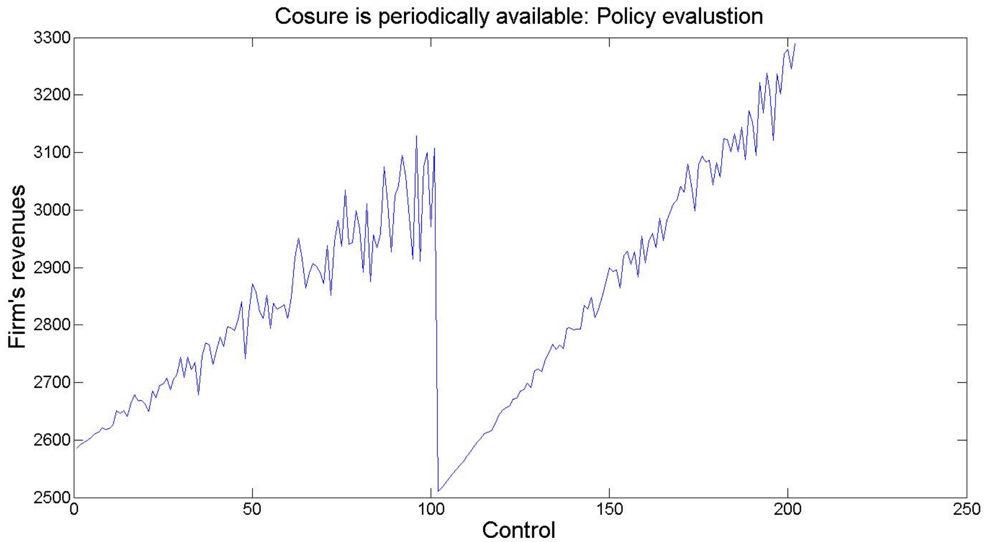

Figure 7-8: Average NN outputs (firm revenue) for the 202 NNs (one per control input) during the last 300 steps of the learning process. Firm is assumed to be risk neutral and closure is periodically available. One NN is dedicated to learning the long-term impact of each available policy ($[u_k]_1 = 0, 0.1, 0.2 \ldots, 1$ and $[u_k]_2 \in [1, 2]$), indexed from 0 to 202 in the horizontal axis) on firm's revenue.

always available (which would mean that the firm can always avoid an audit) can be attributed to the approximate nature of the SARSA/NN - based method used to optimize the tax evasion policy. The ranking of the four scenarios of closure availability has not changed between the two approaches.

7.2 Risk-Averse Firms

Up to now, we have studied the performance of our model under the assumption of risk neutrality and computed the optimal policy via SARSA/NN. We now proceed to re-introduce risk-aversion. The value of the risk-aversion coefficient that governs the firm's behaviour was chosen to be consistent with the literature findings on Greece, such as [102], which estimates that the total "hidden" economy to be approximately 40%. As a result we ran the model for different risk-aversion coefficients, and aimed to identify the one that resulted in an optimal policy ($[u_k]_1$) close to 40% (Figure 7-9). We found that the risk-aversion coefficient in Eq. 6.10 that results in an optimal

Option availability	Expected firm profit (net) via ADP	St.d. (ADP)	Expected firm profit (net) via DP
Never available	3,114.3	86.25	3254.6
Stochastic (20%)	3,241.5	92.80	3307.9
Always available	3,324.5	51.73	3358.3
Every five years	3,253.5	80.11	3313.9

Table 7.1: Expected firm revenues and standard deviation thereof under different option availability scenarios, with $r = 0.24$, $\beta = 0.24$, $\ell = 0.023$, $\lambda = 0$ and a 5% overall audit probability. Numbers are expressed in % of the firm's annual profit, discounted at a 3% annual rate of inflation. The figures for the last scenario ("every five years") are for an initial state of $x = [11, 2, 0, 0, 0, 0, 1]^T$, i.e., the firm hid all its profit since its last closure, 1 year ago, and has no closure option at its disposal for the next 4 years. The expected firm profits obtained via DP in Chapter 4 are included for comparison

policy near 40% was $\lambda = 2$.

After averaging the last 300 steps in each learning episode, for each available control, the average utility (see Fig. 7-10) indicates that the best policy for the firm is to conceal 40% of its profits, i.e., $[u_k]_1 = 0.4$. This can be attributed to the increased penalization (due to risk-aversion) of adverse scenarios, such as an audit, and the limited ability of the NN to capture a cost equation that is simultaneously non-linear and stochastic. As expected, the volatility obtained by the firm using the optimal policy is increased when closure is stochastically available, compared to the other deterministic scenarios shown in Table 7.2, where we list the maximum expected utility in each setting along with the standard deviation of the NN's outputs during the learning process assuming that the firm follows the optimal policy. Compared with the numerical results presented in Table 7.1, the introduction of risk aversion does not change the fact that the closure option remains significantly advantageous for the firm. Once again, the maximum expected utility when closure is always available is greater than when it is periodically available, which in turn is greater than when the option is stochastically available. Finally, the firm's expected utility is minimum when closure is not available. As in the case of risk neutrality, closure is appealing to the firm because it decreases both the risk of a potential audit and the effectiveness of

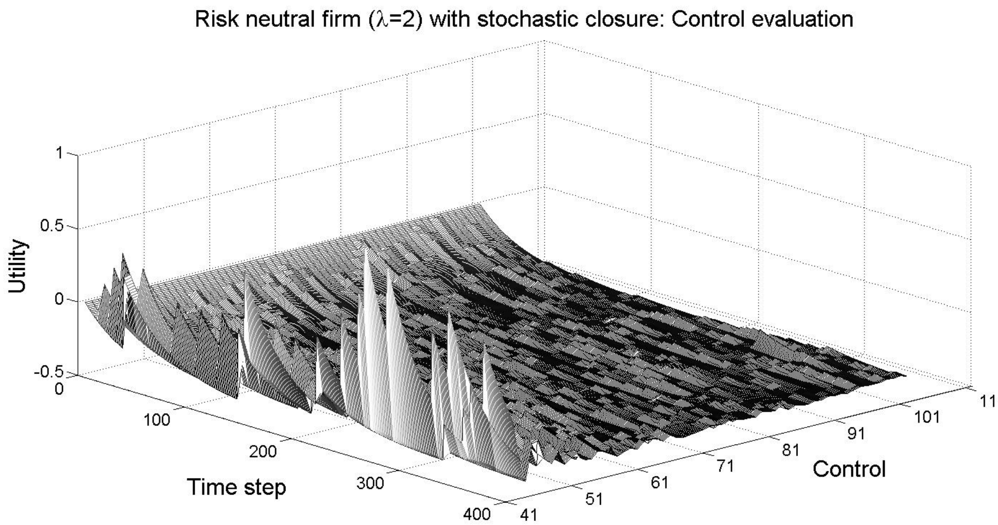

Figure 7-9: Control evaluation using independent neural networks for each control when risk-aversion ($\lambda = 2$) is assumed and closure is stochastically available (probability 20% per annum. The figure only includes the results for the learning process of the last 101 controls (the firm decides to use the option when available) for clarity reasons. The utility is maximized at 41st control (i.e., $u_1 = 0.4$ or 40% tax evasion).

the tax penalty as a deterrent. Finally, government revenues, being "complementary" to firm revenues, decrease with the use of closure in the tax "mixture".

Option availability	Expected utility (net)	St.d
Never available	0.17	0.32
Stochastic (20%)	0.37	0.46
Always available	1.17	0.37
Every five years	0.41	0.42

Table 7.2: Expected maximum utility and standard deviation under different option availability scenarios, with $r = 0.24$, $\beta = 0.24$, $\ell = 0.023$, $\lambda = 2$ and a 5% overall audit probability. The figures for the last scenario ("every five years") are for an initial state of $x = [11, 2, 0, 0, 0, 0, 0, 1]^T$, i.e., the firm hid all its profit since its last closure, 1 year ago, and has no closure option at its disposal for the next 4 years.

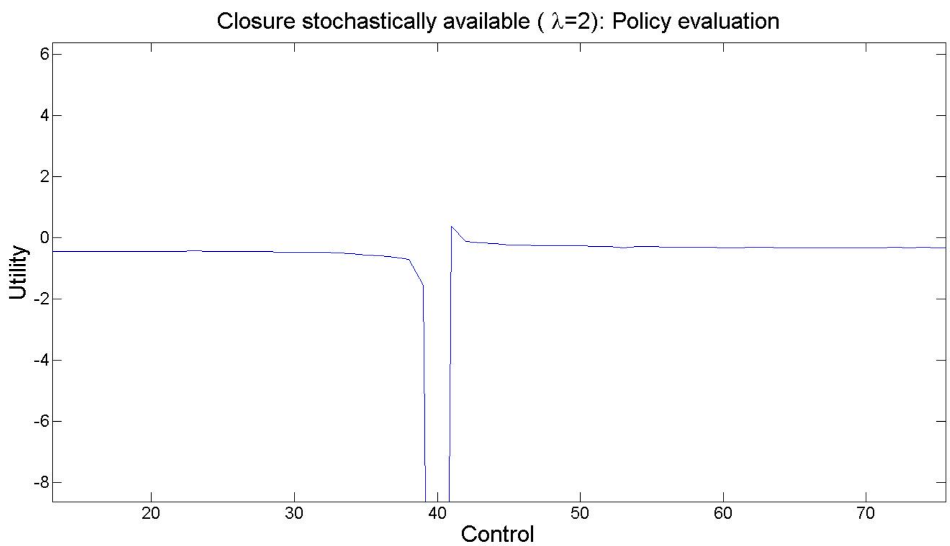

Figure 7-10: Average NN outputs (firm revenue) for the 202 NNs (one per control input) during the last 300 steps of the learning process. Firm is assumed to be risk averse ($\lambda = 2$ and closure is stochastically available. One NN is dedicated to learning the long-term impact of each available policy ($[u_k]_1 = 0, 0.1, 0.2, ..., 1$, indexed from 0 to 202 in the horizontal axis) on firm's revenue. Figure zooms in an interval of controls around 0.4, where utility is maximized.

Chapter 8

Conclusions and Opportunities for Future Work

This thesis describes a decision support model aimed at capturing the decision process of a typical firm which aims to maximize its long-term revenues through tax evasion. The proposed model is adaptable to different taxation schemes, both linear and non-linear, and can easily be "tuned" to reflect the values of various tax parameters presumably known to the government. Specifically, we have developed a Markov-based model which describes the firm's evolution within the Greek tax system. Our model incorporates an idiosyncratic form of optional tax amnesty used in Greece, termed "closure", which gives the firm the opportunity to pay a lump sum in exchange for tax amnesty over a given period of time, up to a maximum of five years in the past.

We computed the optimal tax evasion policy as a percentage of the firm's annual profits, and the optimal use of the closure option for the purpose of long-term revenue maximization. We initially assumed risk neutrality, which implies linearity of the firm's reward function. Doing so limits the optimal policies to a bang-bang form, and makes the optimum easy to compute via Dynamic Programming (DP). Our results support the argument that, under the current state of affairs, there is a strong incentive for firms to engage in tax evasion. Closure is considered a tax revenue gathering tool, meant as an incentive for firms to "pay up" where they might not otherwise. As such, our analysis suggests that it is quite ineffective in the long-run. Even under extensive tax evasion, the absence of closure can boost long-term tax revenues substantially. At first glance, (ignoring the per-audit cost) tax penalties do not appear to be an effective deterrent when compared to audits.

We proceeded to lift the assumption of risk-neutrality on behalf of the firm, making the problem of computing the optimal policy impossible to solve via DP. In this new setting, we initially removed the closure option, thus formulating a classical optimal portfolio selection problem. The idea behind this simplified model was to evaluate the effect of the firm's "current" decision concerning tax evasion for five years into the future. The optimal portfolio approach was effective precisely because of this five-year statutory limit on auditing. When risk aversion was thus included, the results obtained for a risk-neutral firm still hold: under the current set of tax parameters, tax evasion remains persistent. This becomes especially evident considering the high risk aversion coefficient ($\lambda = 6$), that seems to be "required" in order to begin limiting tax evasion. Moreover, tax penalties remain ineffective as a deterrent when compared to more frequent audits.

The portfolio selection problem could easily accommodate the non-linear reward function that results from risk aversion, but not the closure option. To do that, we returned to our initial Markov-based model and used Approximate Dynamic Programming (ADP) to compute the optimal firm policy. In our case, ADP was applied by i) discretizing the firm's state and decision spaces, and ii) using an on-line policy evaluation method (specifically, the SARSA(0) algorithm) together with a function approximator (a feed-forward neural network). The neural network was trained to "store" the firm's optimal long-term revenues, given a starting state and decision. The SARSA(0) algorithm was used to efficiently "learn" the optimal firm decisions through simulations of the firm's state evolution. The ADP approach was validated by initially assuming risk neutrality and comparing the results thus obtained (optimal policy and long-term firm revenues) to those previously computed via DP. We subsequently introduced a constant relative risk aversion utility function for the firm, and demonstrated that we can indeed compute the optimal policy and revenues in the "full" model which includes both non-linearity in the reward function and the option of closure. We also estimated the risk aversion coefficient of the average Greek firm to be approximately $\lambda = 2$, by assuming that the Greek "hidden" economy is approximately 40% as suggested by the literature.

Based on our results, we can argue that a typical firm in Greece currently has no incentive to behave honestly. The prevailing combination of tax parameters (tax penalty, audit probability and closure cost) does not prevent the firm from engaging in tax evasion. This phenomenon is exacerbated when the closure option is available, thus allowing the firm an alternative "route" to the usual tax payment. Closure significantly reduces the government's long-term revenues for both risk-neutral and risk-averse firms. The incorporation of closure into the Greek taxation system has other side effects as well: it limits the effectiveness of tax penalties and audits as tax evasion deterrents. Finally, based on our numerical experiments, we can argue that the Greek government should invest in a stronger audit mechanism instead of higher tax penalties, based on the fact that audits appear to be a more effective deterrent and that increased tax penalties may have other socioeconomic effects (i.e., effects on the firms' welfare).

Our model is not a complete description of the Greek taxation system, and not all tax-paying entities are firms. Also, some parameter values were indirectly estimated because the Greek government does not publish relevant data (i.e., on audit and closure rates). As a result, this thesis cannot be regarded as a complete analysis of tax evasion but instead as a new quantitative alternative to the existing literature which historically was been based on macroeconomic or game-theoretic models.

Opportunities for further work include the use of a variation of the basic algorithm used here, namely SARSA(λ), to improve the performance of ADP in the case of a risk-averse firm. To improve the precision and convergence rate of the "learning process", one could also explore the use of more sophisticated NNs. In addition, there are alternative function approximators to consider, such as Coarse, Tile, and Kanerva Coding.

With respect to the descriptive power of our model, it would be interesting to incorporate other forms of taxation (such as VAT) or alternative forms of a tax amnesty. The model could also be adapted to a multi-class setting where different taxation schemes are applied for different taxpayer groups. One potential use of this approach could be the classification of different taxpayers' decision patterns together with the development of appropriate metrics that express *how close* are the decision patterns of two different taxpayers, with an eye towards developing more sophisticated taxation schemes that would result in a more fair system with no incentives for tax evasion. Finally, the model could be augmented to include "hidden" Markov states corresponding to the government's decision process regarding tax parameters such as audits, penalties and tax rates. In such a setting, the firm would have to make implicit estimates of the parameters' values by "observing the environment" through a "learning process" such as the one in the ADP approach.

Appendix A

Markov transition matrices

$$
M_{no} = \begin{bmatrix}
0.0025 & 0.0025 & 0.0025 & 0.0025 & 0.0025 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.0025 & 0.0025 & 0.0025 & 0.0025 & 0.0025 & 0.0025 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0025 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0025 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.04 & 0.04 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.9975 & 0.9975 & 0.9975 & 0.9975 & 0.9975 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.9975 & 0.9975 & 0.9975 & 0.9975 & 0.9975 & 0.9975 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9975 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9975 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.96 & 0.96
\end{bmatrix} \tag{A.1}
$$

Table A.1: Transition probabilities M_{no}: closure is not available

$$
M_a = \begin{bmatrix}
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{bmatrix} \tag{A.2}
$$

Table A.2: Transition probabilities M_a: closure is available and the firm decides to use it.

$$M_d = \begin{bmatrix}
0.0075 & 0.0075 & 0.0075 & 0.0075 & 0.0075 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.0075 & 0.0075 & 0.0075 & 0.0075 & 0.0075 & 0.0075 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0075 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.0075 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.12 & 0.12 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0.9925 & 0.9925 & 0.9925 & 0.9925 & 0.9925 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0.9925 & 0.9925 & 0.9925 & 0.9925 & 0.9925 & 0.9925 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9925 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.9925 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.88 & 0.88
\end{bmatrix} \tag{A.3}$$

Table A.3: Transition probabilities M_d: closure is always available and the firm declines it.

Appendix B

Algorithm

```
%%%%%%%%%%%%%%%%% FILE mdp_tax.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%This file handles the case of a stochastically available option.
%It computes the firm's maximum expected rewards and optimal decisions
%with respect to tax evasion and use of the closure option, when the latter
%is available.

clear all

global Jk
%Firm reward matrix (eq. (11)). Jk will be filled in by row using
%value iteration. Columns of Jk will correspond to firm states (see
%discussion in Section 2.3).

global Sk %Similar to Jk but for Govenment revenues.

global PERIODIC
%BOOLEAN set to indicate whether the option is offered
%periodically. This file is designed to work with
%PERIODIC=0 ONLY.
PERIODIC=0; %This file is tailored to NON-periodic option.

global v %Probability of closure being available each year.
v=0.2; %Set closure probability (v=1 means closure always available)

global gam
gam=1/(1+0.03); %Gamma - discount factor

global DOGOV
DOGOV=0;    %Set to 1 to compute government revenues explicitly. Useful if
%one wishes to take into account complications such as the
%government being unable to collect a portion of the tax owed
%(see discussion in Sec. 3.5)

%Tax parameters (Sec. 2.3)
global R %Firm annual revenue
R=100;
global r
r=0.24; %tax rate
global b
b=0.24; %0.24; %xxx tax penalty rate
```

```
global l
l=0.023;% closure cost as a fraction of firm's annual revenue

Nstates=15*2*2^5;
%number of discrete states on which the firm's reward will be computed
eps=0.01; %convergence threshold
Jk=zeros(1,Nstates);
Sk=Jk;
uopt=Jk; %holds optimal decision with respect to tax evasion at each state
aopt=Jk; %holds optimal decision with respect to closure usage

%%% Begin Value Iteration

converged=0; %Boolean flag
k=1; %initialize the 1st row of Jk
for i=1:Nstates %fill in rewards for all states
    %The state (see eq. 4) is a 7-tuple (s, c, h) where h is 5-dimensional.
    %Determine which 7-tuple the counter i corresponds to.
    s=floor((i-1)/(2*2^5))+1; %Markov state
    c=floor(mod(i-1,2*2^5)/(2^5))+1; %Closure availability
    hs=mod(i-1,2^5); %now form history vector, h
    hs=dec2bin(hs,5);
    h=zeros(5,1);
    for j=1:5
        h(j)=str2num(hs(j));
    end
    %compute rewards/gov.revenues with no history of tax evasion by the
    %firm
    Jk(i)=g(s,c,h,0);
    Sk(i)=ggov(s,c,h,0);
end
while (converged==0)
    k=k+1 %fill in a new row of Jk, Sk
    tempFirm=zeros(1,Nstates);
    tempGov=tempFirm;
    %compute firm & government rewards, optimal decisions for all states
    for j=1:Nstates
        [tempFirm(j), tempGov(j), uopt(j), aopt(j)]=expreward(j,k);
    end
    %Append new rows to Jk, Sk
    Jk=[Jk; tempFirm];
    Sk=[Sk; tempGov];
    %display firm reward corresponding to the state x_k=(1,2,[0 0 0 0 0]')
    %i.e., starting from an audit, no option available, clean history
    %Jk(k,1)
    Jk(k,33)
    uopt(33)
    %Sk(k,1)
    %plot(uopt); drawnow;
    [evade, useoption]=checkfeasible(uopt,aopt);
evade
useoption
    %check for convergence
    maxdiff=max(abs(Jk(k,:)-Jk(k-1,:)))
    if (maxdiff)<eps
        converged=1;
    end
end

fprintf(1,'firm rewards=%f\n',Jk(k,33));
fprintf(1,'gov.t. rewards=%f\n',100*1/(1-gam)-Jk(k,33));
```

```
%%%%%%%%%%%%%%%%%% end of FILE mdp_tax.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%% FILE mdpperiodic_tax.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear all
global Jk
%Firm reward matrix (eq. (11)). Jk will be filled in by row using
%value iteration. Columns of Jk will correspond to firm states (see
%discussion in Section 2.3).

global Sk %Similar to Jk but for Govenment revenues.

global PERIODIC
%BOOLEAN set to indicate whether the option is offered
%periodically. This file is designed to run with
%PERIODIC=1 option ONLY.
PERIODIC=1;

global CLOSURE_PERIOD
CLOSURE_PERIOD=5; %period for closure option

%global v %Probability of closure being available each year.
%v=1; %Set closure probability (v=1 means closure always available)

global gam
gam=1/(1+0.03); %Gamma - discount factor

global DOGOV
DOGOV=0;    %Set to 1 to compute government revenues explicitly. Useful if
%one wishes to take into account complications such as the
%government being unable to collect a portion of the tax owed
%(see discussion in Sec. 3.5)

%Tax parameters (Sec. 2.3)
global R %Firm annual revenue
R=100;
global r
r=0.24; %tax rate
global b
b=0.24; %tax penalty rate
global l
l=0.023;% closure cost as a fraction of firm's annual revenue

Nstates=15*2*2^5;
eps=0.01;

Jk=zeros(1,Nstates);
Sk=Jk;
uopt=Jk;
aopt=Jk;

converged=0;
k=1; %initialize the first row of Jk

for i=1:Nstates %fill in rewards for all states
    %determine which state x_k=(s,c,h) j corresponds to
    s=floor((i-1)/(2*2^5))+1;
    c=floor(mod(i-1,2*2^5)/(2^5))+1;
    %hs=mod(i-1,2^5);
    %hs=dec2bin(hs,5);
```

```
    %h=zeros(5,1);
    %for j=1:5
    %    h(j)=str2num(hs(j));
    %end
    Jk(i)=g(s,c,[0 0 0 0 0]',0);
    Sk(i)=ggov(s,c,[0 0 0 0 0]',0);

end

%initialize the remaining time steps within the 1st period
for q=2:CLOSURE_PERIOD
    k=k+1;
    Jk(q,:)=Jk(1,:);
    Sk(q,:)=Sk(1,:);
end

%Apply value iteration
while (converged==0)

    %for a 5-year period, fill 5 rows of Jk
    for q=1:CLOSURE_PERIOD
        k=k+1 %add a row
        tempFirm=zeros(1,Nstates);
        tempGov=tempFirm;

        for j=1:Nstates %fill in rewards for all states
            [tempFirm(j), tempGov(j), tempu(j), tempa(j)]=expreward(j,k);
        end

        Jk=[Jk; tempFirm];
        Sk=[Sk; tempGov];
        uopt=[uopt; tempu];
        aopt=[aopt; tempa];

    end

    %compare rows k to k-5.
    maxdiff=max(max(abs(Jk(k-CLOSURE_PERIOD+1:k,:)-Jk(k-2*CLOSURE_PERIOD+1:k-CLOSURE_PERIOD,:))))

    %display firm reward from state x_k=(1,1,[0 0 0 0 0]');
    Jk(k,33)
    %Sk(k,33)
    if maxdiff<eps
        converged=1;
    end
end

%display firm and gov't rewards when the firm begins from state
%x_k=(11,2,[0 0 0 0 1]'), i.e., the firm hid all its profits since
%its last closure, 1 year ago, and has no closure option at its disposal
%for the past 4 years.

fprintf(1,'firm rewards=%f\n',Jk(k,nexti(11,2,[0 0 0 0 1]')));
fprintf(1,'gov.t. rewards=%f\n',100*1/(1-gam)-Jk(k,nexti(11,2,[0 0 0 0 1]')));

%save testPer Jk Sk

%%%%%%%%%%%%%%%%% end of FILE mdpperiodic_tax.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%% FILE expreward.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

```
%Function expreward(i,k)
%INPUTS: index i of a column in the reward matrix Jk
%        k: stage of value iteration
%OUTPUTS: er: firm's expected reward
%         egovr: gov't expected revenue
%         u1,u2: firm's optimal decisions (tax evation, closure usage).
function [er, egovr, u1, u2]=expreward(i,k)

global gam
global DOGOV


%The firm's state is a 7-tuple (see eq. 3).
%Determine which 7-tuple i corresponds to
[s,c,h]=findstate(i);

%determine expected reward under decision u

%reward function is independent of c - closure state, so we can just
%evaluate it before hand and use it whether c=1 or c=2
g0=g(s,c,h,0);
g1=g(s,c,h,1);

if (c==2) %no option available

    %compute expected value with u1=0 or u1=1
    %CASE 1: no tax evasion, i.e., u1=0; u2=1;
    [firmrev govrev]=expval(s,c,h,k,0,1);
    er0=g0+gam*firmrev; %add using discounting
    if (DOGOV) %work out government revenue explicitly
        egr0=ggov(s,c,h,0)+gam*govrev;
    else
        egr0=0;
    end

    %CASE 2: Max. tax evation, i.e., u1=1; u2=1;
    [firmrev govrev]=expval(s,c,h,k,1,1);
    er1=g1+gam*firmrev;
    if (DOGOV)
        egr1=ggov(s,c,h,1)+gam*govrev;
    else
        egr1=0;
    end

    %find maximum, record optimal decision
    values=[er0,er1];
    gvalues=[egr0, egr1];
    [er, u1]=max(values);
    egovr=gvalues(u1);
    u1=u1-1; %convert index (1 or 2) to 0 or 1.
    u2=0;    %there is no option to consider

else
    if (c==1) %c=1, option available

        %compute exp. reward with u1=0 or u1=1, while TAKING THE OPTION
        %u1=0;
        %u2=1;
        [firmrev govrev]=expval(s,c,h,k,0,1);
        er01=g0+gam*firmrev;
        if (DOGOV)
```

```
            egr01=ggov(s,c,h,0)+gam*govrev;
        else
            egr01=0;
        end

        %u1=1; u2=1;
        [firmrev govrev]=expval(s,c,h,k,1,1);
        er11=g1+gam*firmrev;
        if (DOGOV)
            egr11=ggov(s,c,h,1)+gam*govrev;
        else
            egr11=0;
        end


        %compute exp val with 0 or 1, while REJECTING THE OPTION
        %u1=0; u2=2;
        [firmrev govrev]=expval(s,c,h,k,0,2);
        er02=g0+gam*firmrev;
        if (DOGOV)
            egr02=ggov(s,c,h,0)+gam*govrev;
        else
            egr02=0;
        end

        %u1=1; u2=2;
        [firmrev govrev]=expval(s,c,h,k,1,2);
        er12=g1+gam*firmrev;
        if (DOGOV)
            egr12=ggov(s,c,h,1)+gam*govrev;
        else
            egr12=0;
        end

        %pick the best value, record decisions for u1, u2
        values=[er01 er11 er02 er12];
        govvalues=[egr01 egr11 egr02 egr12];
        [er, idx]=max(values);
        egovr=govvalues(idx);
        u1=mod(idx-1,2);
        u2=floor((idx-1)/2)+1;
    else
        fprintf(1,'Oops...\n');
        pause;
    end
end

end


%The firm's state is a 7-tuple (see eq. 3).
%This function takes as input an integer i corresponding
%to a column of Jk, uopt, or aopt, and determines which 7-tuple
%corresponding to that column
function [s c h]=findstate(i)

s=floor((i-1)/(2*2^5))+1;
c=floor(mod(i-1,2*2^5)/(2^5))+1;
%c=mod(s1,(2^5))+1
hs=mod(i-1,2^5);
hs=dec2bin(hs,5);
```

```
h=zeros(5,1);
for j=1:5
    h(j)=str2num(hs(j));
end

%%%%%%%%%%%%%%%%% end of FILE expreward.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%% FILE expval.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%Function expval(s,c,h,k,u1,u2). Used in Value Iteration to computes the
%expected reward firm and government rewards from stage k onwards.
%
%INPUTS: s: firm's Markov state
%        c: closure state (see Sec. 2.1)
%        h: firm's tax evasion history vector
%        k: Value iteration stage
%       u1: firm's current decision with respect to tax evasion (u1 is the
%           fraction of profit the firm chooses to conceal).
%       u2: firm's current decision to use (or not) the closure option.
%
%OUTPUT: efirmrev: expected firm revenue while at the state defined
%                  by (s,c,h), making decision u1, u2.
function [efirmrev egovrev]=expval(s,c,h,k,u1,u2)

%see main file
global Jk
global Sk
global PERIODIC
global DOGOV

%initialize firm and gov't revenue values to 0.
efirmrev=0;
egovrev=0;

%See Sec. 2.1
H=[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];

e5=[0 0 0 0 1]';

%Based on the firm's decision, the next history vector will be...
nexth=H*h+e5*u1;

i=s;

for t=1:2 %iterate on closure being available next year, or not
    for j=1:15
        %the next possible sstate will be x_k=(j,t,H*h+e_5*u_1)
        if (PERIODIC==0)
            efirmrev=efirmrev+Pr(c,j,i,u2)*Pre(t)*Jk(k-1,nexti(j,t,nexth));
            if (DOGOV)
                egovrev=egovrev+Pr(c,j,i,u2)*Pre(t)*Sk(k-1,nexti(j,t,nexth));
            end
        else
            efirmrev=efirmrev+Pr(c,j,i,u2)*PrePer(t,k)*Jk(k-1,nexti(j,t,nexth));
            if (DOGOV)
                egovrev=egovrev+Pr(c,j,i,u2)*PrePer(t,k)*Sk(k-1,nexti(j,t,nexth));
            end
```

```
        end
    end
end

end

%%%%%%%%%%%%%%%%%% end of FILE expval.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%% FILE g.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%Function g(s,c,h,u1) firm reward (see eq. 9).
%
%INPUTS: s: Firm's Markov state (see Sec. 2.1)
%        c: closure state - NOT USED HERE
%        h: firm's tax evasion history vector
%       u1: firm's current decision with respect to tax evasion (u1 is the
%           fraction of profit the firm chooses to conceal).
%OUTPUT: rew: firm's reward while at state x=(s,c,h'), making decision u1.
function rew=g(s,c,h,u1)

%Tax parameters (see Section 2.3)
global R %Firm annual revenue - set in main file
global r %tax rate
global b %tax penalty rate
global l %closure cost as a fraction of firm's annual revenue

%R=100;
%r=0.24;
%b=0.24;
%l=0.023;

rew=0;
%reward computation. See eq. (9) for formula and explanation
if (s>=11 && s <=15) %we are in a N1...N5 state (no audit/ no option)
    rew=1-r+r*u1;
else
    if (s>=6 && s<=10) %We are in an O1,...,O5 (option) state
        rew=1-r+r*u1-l*(s-5);
    else
        if (s>=1 && s<=5) %We are in a V1...V5 (audit) state
            sh=0;
            sih=0;
            for i=1:s
                sh=sh+h(length(h)-i+1);
                sih=sih+i*h(length(h)-i+1);
            end
            %penalty if audited. Includes 3/5 discount term for prompt pmt
            rew=1-r+r*u1-r*sh-(3/5)*b*r*sih;
        else
            fprintf(1,'Oops...\n');
            pause;
        end
    end
end

rew=rew*R;

end

%1-r*(1-u)=1-r+ru
```

```
%%%%%%%%%%%%%%%%% end of FILE g.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%% FILE ggov.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%Function ggov(s,c,h,u1)
%
%INPUTS: s: firm's Markov state (see eq. (4)).
%        c: option state
%        h: firm's tax evasion history vector
%        u1: firm's decision with respect to tax evasion (see Sec. 2.1).
%OUTPUT: rew=government revenues when the firm is in state x_k=(s,c,h) and
%        conceals u1 fraction of its profit.
%Note: Government revenue can usually be computed in the main file simply
%from knowledge of the firm's optimal reward. This function is provided for
%experimentation with various options, such as the government being unable
%to collect a portion of the revenue it is owed.

function rew=ggov(s,c,h,u1)

%Tax parameters (see Section 2.3)
global R %Firm annual revenue - set in main file
global r %tax rate
global b %tax penalty rate
global l %closure cost as a fraction of firm's annual revenue

%Set to <1 if some penalties/back taxes go uncollected
fraction_collected=1;

%h
rew=0;
if (s>=11 && s <=15)
    rew=r*(1-u1);
else
    if (s>=6 && s<=10)
        rew=r*(1-u1)+l*(s-5);
    else
        if (s>=1 && s<=5)
            sh=0;
            sih=0;
            for i=1:s
                sh=sh+h(i);
                sih=sih+i*h(i);
            end
            rew=r*(1-u1)+fraction_collected*(r*sh+r*(3/5)*b*sih);
        else
            fprintf(1,'Oops...\n');
            pause;
        end
    end
end

rew=rew*R;

end

%1-r*(1-u)=1-r+ru

%%%%%%%%%%%%%%%%% end of FILE ggov.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%% FILE Pr.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Pr(q,i,j,m): see eq. (8).
% INPUTS: q=1 (option available) or 2 (no option).
%          i: starting state
%           j: end state
%           m: 1 (firm takes option) or 2 (firm declines option).
%OUTPUT:  probability of transitioning from state j to state i given the
%availability of the closure option (q) and firm's decision to use the
%option (m)
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
function p=Pr(q,i,j,m)

%Transition matrix without option
Mno=[
0.0025  0.0025  0.0025  0.0025  0.0025  0       0       0       0       0       0       0      0        0      0
0       0       0       0       0       0.0025  0.0025  0.0025  0.0025  0.0025  0.0025  0       0        0      0
0       0       0       0       0       0       0       0       0       0       0       0.0025  0        0      0
0       0       0       0       0       0       0       0       0       0       0       0       0.0025   0      0
0       0       0       0       0       0       0       0       0       0       0       0       0        0.04   0.04
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
0.9975  0.9975  0.9975  0.9975  0.9975  0       0       0       0       0       0       0    0        0      0
0       0       0       0       0       0.9975  0.9975  0.9975  0.9975  0.9975  0.9975  0       0        0      0
0       0       0       0       0       0       0       0       0       0       0       0.9975  0        0      0
0       0       0       0       0       0       0       0       0       0       0       0       0.9975   0      0
0       0       0       0       0       0       0       0       0       0       0       0       0        0.96   0.96
];

%Transition matrix with option available & firm accepts
Ma=[
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
1       1       1       1       1       1       1       1       1       1       1       0       0       0   0
0       0       0       0       0       0       0       0       0       0       0       1       0       0   0
0       0       0       0       0       0       0       0       0       0       0       0       1       0   0
0       0       0       0       0       0       0       0       0       0       0       0       0       1   0
0       0       0       0       0       0       0       0       0       0       0       0       0       0   1
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
];

%Transition matrix with option available & firm declines
Md=[0.0075  0.0075  0.0075  0.0075  0.0075  0       0       0       0       0       0       0    0        0      0
     0       0       0       0       0       0.0075  0.0075  0.0075  0.0075  0.0075  0.0075  0       0        0      0
     0       0       0       0       0       0       0       0       0       0       0       0.0075  0        0      0
     0       0       0       0       0       0       0       0       0       0       0       0       0.0075   0      0
     0       0       0       0       0       0       0       0       0       0       0       0       0        0.12   0.12
     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
```

```
0.9925      0.9925  0.9925  0.9925  0.9925  0       0       0       0       0       0       0    0        0      0
    0       0       0       0       0       0.9925  0.9925  0.9925  0.9925  0.9925  0.9925  0       0        0      0
    0       0       0       0       0       0       0       0       0       0       0       0.9925  0        0      0
    0       0       0       0       0       0       0       0       0       0       0       0       0.9925   0      0
    0       0       0       0       0       0       0       0       0       0       0       0       0        0.88   0.88
    ];

if (q==2) %closure not available
    p=Mno(i,j);
else
    if (q==1) && (m==1) %take option
        p=Ma(i,j);
    else
        if (q==1) && (m==2) %decline
            p=Md(i,j);
        else
            fprintf(1,'Ooops...\n');
            pause;
        end
    end
end

%%%%%%%%%%%%%%%%%% end of FILE Pr.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%% FILE pre.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% Pre(t,k) corresponds to Pr(e_k=i) in eq. (6)
% INPUTS: t=1 if closure is available, 0 otherwise
% OUTPUT: probability that closure is given
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
function p=pre(t)

%c is the probability of closure - defined in the main mdp.m file
global v

%uncomment this statement for stochastic closure
if (t==1) %closure available
    p=v;
else %no closure
    p=1-v;
end

%%%%%%%%%%%%%%%%%% end of FILE pre.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%% FILE PrePer.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%
% PrePer(t,k) corresponds to Pr(e_k=i) in eq. (6) in the case
% where the option is given periodically, with period CLOSURE_PERIOD
% INPUTS: t=1 if closure is available, 0 otherwise
% OUTPUT: probability that closure is given
%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
function p=PrePer(t,k)

global CLOSURE_PERIOD %set in mdpperiodic_tax.m
%CLOSURE_PERIOD=5; %closure period in years
```

```
if mod(k,CLOSURE_PERIOD)==1 %we are in a year where closure is offered
    if (t==1)
        p=1;
    else
        p=0;
    end
else %not in a year where closure is offered
    if (t==1)
        p=0;
    else
        p=1;
    end
end

%%%%%%%%%%%%%%%%% end of FILE PrePer.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

%%%%%%%%%%%%%%%%% FILE checkfeasible.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

function [evade, useoption]=checkfeasible(uopt,aopt)
global v

H=[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];

e5=[0 0 0 0 1]';

feasible=0*uopt;

converged=0;

feasible(1)=1; %seed initial state

while (~converged)

    savedfeasible=feasible;

    for i=1:length(feasible)

        if (feasible(i)==1) %find next feasible states

            %find state and controls
            [s,c,h]=column2state(i);
            u1=uopt(i);
            a1=aopt(i);

            %Based on the firm's decision, the next history vector will be...
            nexth=H*h+e5*u1;

            %find which are the next possible states

            vstate=zeros(15,1);
            for j=1:15
                vstate(j)=(Pr(c,j,s,a1)~=0);
            end

            %mark states
            for j=1:15
```

```
                if vstate(j)==1
                    if (v>0)
                        feasible(nexti(j,1,nexth))=1;
                    end
                    if (v<1)
                        feasible(nexti(j,2,nexth))=1;
                    end
                end
            end

        end
    end

    if (feasible==savedfeasible)
        converged=1;
    end
end

evade=sum(uopt.*feasible);
useoption=sum((aopt==1).*feasible);

%check if uopt=1 in all feasible states
%figure(1);
%plot(uopt.*feasible);

%figure(2);
%plot((aopt==1).*feasible);

%figure(3);
%plot(feasible);
%drawnow;

%%%%%%%%%%%%%%%%% end of FILE checkfeasible.m %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

Appendix C

First Passage Probabilities

First passage probabilities $f_{1j}^{(n)}$

			n		
	0.0025	0.0025	0.0025	0.0025	0.0396
	0.0025	0.0025	0.0025	0.0397	0.0380
j	0.0025	0.0025	0.0398	0.0381	0.0365
	0.0025	0.0399	0.0382	0.0352	0.0324
	0.0400	0.0369	0.0340	0.0314	0.0290

Table C.1: First passage probabilities f_{1j} from each Markov state j to the state 1 (audit)

		$1 - \sum_{n=1}^{5} f_{1j}^{(n)}$		
0.9504	0.9148	0.8806	0.8518	0.8287

Table C.2: Probabilities that an audit will not occur within five years given that the firm is in Markov state j

Bibliography

[1] Ministry of Finance, 2004, Law N.3259/2004 (POL.1034/2005) (in Greek).

[2] ——, 2008, Law N.3697/2008 (POL.1130/2008) (in Greek).

[3] J. Andreoni, B. Erard, and J. Feinstein, "Tax compliance," *Journal of Economic Literature*, vol. 36, no. 2, pp. 818–860, Jun 1998.

[4] A. Sandmo, "Optimal taxation: An introduction to the literature," *Journal of Public Economics*, vol. 6, pp. 37–54, 1976.

[5] R. G. Lipsey and K. Lancaster, "The general theory of second best," *Review of Economic Studies*, vol. 24, pp. 11–32, 1957.

[6] M. G. Allingham and A. Sandmo, "Income tax evasion: A theoretical analysis," *Journal of Public Economics*, vol. 1, no. 6, pp. 988–1001, Jun. 1972.

[7] G. S. Becker, "Crime and punishment: An economic approach," *Journal of Political Economy*, vol. 76, no. 2, pp. 169–217, 1968.

[8] J. Andreoni, "Reasonable doubt and the optimal magnitude of fines: Should the penalty fit the crime?" *The RAND Journal of Economics*, vol. 22, no. 3, pp. 385–395, 1991b.

[9] S. Shavell, "The optimal use of nonmonetary sanctions as a deterrent," *American Economic Review*, vol. 77, no. 4, pp. 584–592, 1987.

[10] D. Mookherjee and I. Png, "Optimal auditing, insurance, and redistribution," *The Quarterly Journal of Economics*, vol. 104, no. 2, pp. 399–415, May 1989.

[11] D. Mookherjee and I. P. L. Png, "Monitoring vis-a-vis investigation in enforcement of law," *American Economic Review*, vol. 82, no. 3, pp. 556–565, 1992.

[12] S. Yitzhaki, "A note on income tax evasion: A theoretical analysis," *Journal of Public Economics*, vol. 3, no. 2, pp. 201–202, 1974.

[13] J. C. Baldry, "Tax evasion and labour supply," *Economic Letters*, vol. 3, no. 1, pp. 53–56, 1979.

[14] J. H. Pencavel, "A note on income tax evasion, labour supply and nonlinear tax schedules," *Journal of Public Economics*, vol. 12, no. 1, pp. 115–124, 1979.

[15] F. A. Cowell, "Taxation and labour supply with risky activities," *Economica*, vol. 48, no. 192, pp. 365–379, Nov. 1981.

[16] A. Sandmo, "Income tax evasion, labour supply, and the equity-evasion trade-off," *Journal of Public Economics*, vol. 16, no. 3, pp. 265–288, 1981.

[17] J. P. P. Gordon, "Individual morality and reputation costs as deterrents to tax evasion," *European Economic Review*, vol. 33, no. 4, pp. 797–805, 1989.

[18] B. S. Frey and L. P. Feld, "Deterrence and morale in taxation: An empirical analysis," CESifo, Tech. Rep. 760, 2002. [Online]. Available: http://www.ifo-dresden.de/

[19] J. Hindriks, M. Keen, and A. Muthoo, "Corruption, extortion and evasion," *Journal of Public Economics*, vol. 74, no. 3, pp. 395–430, 1999.

[20] K. Basu, S. Bhattacharya, and A. Mishra, "Notes on bribery and the control of corruption," *Journal of Public Economics*, vol. 48, no. 3, pp. 349 – 359, 1992.

[21] T. Besley and J. McLaren, "Taxes and bribery: The role of wage incentives," *The Economic Journal*, vol. 103, no. 416, pp. 119–141, 1993.

[22] P. Chander and L. Wilde, "Corruption in tax administration," *Journal of Public Economics*, vol. 49, no. 3, pp. 333 – 349, 1992.

[23] D. Mookherjee and I. P. L. Png, "Corruptible law enforcers: How should they be compensated?" *The Economic Journal*, vol. 105, pp. 145–159, 1995.

[24] E. M. Engel and J. R. Hines Jr, "Understanding tax evasion dynamics," National Bureau of Economic Research, Tech. Rep. 6903, 1999. [Online]. Available: http://www.nber.org/papers/w6903

[25] B. Dalamagas, "A dynamic approach to tax evasion," *Public Finance Review*, vol. 39, no. 2, pp. 309–326, 2011.

[26] N. Goumagias, D. Hristu-Varsakelis, and A. Saraidaris, "A decision support model for tax revenue collection in Greece," *Decision Support Systems*, vol. 53, no. 2, pp. 76–96, 2012.

[27] N. Goumagias and D. Hristu-Varsakelis, "Tax evasion by risk-averse firms in Greece: A discrete markov-based optimization model," *Optimization*, vol. 62, no. 8, pp. 1153–1167, 2013.

[28] J. F. Reinganum and L. L. Wilde, "Equilibrium verification and reporting policies in a model of tax compliance," *International Economic Review*, vol. 27, no. 3, pp. 739–760, Oct. 1986.

[29] I. Sanchez and J. Sobel, "Hierarchical design and enforcement of income tax policies," *Journal of Public Economics*, vol. 50, no. 3, pp. 345 – 369, 1993.

[30] K. M. Eisenhardt, "Agency theory: An assessment and review," *Academy of Management Review*, vol. 14, no. 1, pp. 57–74, 1989.

[31] R. Rees, "The theory of principal and agent part I," *Bulletin of Economic Research*, vol. 37, no. 1, pp. 3–26, 1985.

[32] ——, "The theory of principal and agent part II," *Bulletin of Economic Research*, vol. 37, no. 2, pp. 75–97, 1985.

[33] D. M. Kreps and R. Wilson, "Sequential equilibria," *Econometrica*, vol. 50, no. 4, pp. 863–894, Jul. 1982.

[34] M. Spence, *Market signalling.* Cambridge, MA: Harvard University Press, 1974.

[35] S. Scotchmer, "Equity in tax enforcement," *Harvard discussion paper*, 1986, Harvard Institute of Economic Research.

[36] J. A. Mirrlees, "An exploration in the theory of optimum income taxation," *The Review of Economic Studies*, vol. 38, no. 2, pp. 175–208, 1971.

[37] M. J. Graetz, J. F. Reinganum, and L. L. Wilde, "The tax compliance game: Toward an interactive theory of tax enforcement," *Journal of Law, Economics, & Organization*, vol. 2, no. 1, pp. 1–32, 1986.

[38] S. Baiman and J. S. Demski, "Economically optimal performance evaluation and control systems," *Journal of Accounting Research*, no. 18, pp. 184–220, 1980, supplement.

[39] R. A. Dye, "Optimal monitoring policies in agencies," *The RAND Journal of Economics*, vol. 17, no. 3, pp. 339–350, 1986.

[40] R. C. Sansing, "Information acquisition in a tax compliance game," *The Accounting Review*, vol. 68, no. 4, pp. 874–884, Oct. 1993.

[41] P. J. Beck and W.-O. Jung, "Taxpayers' reporting decisions and auditing under information asymmetry," *Accounting Review*, vol. 64, no. 3, pp. 468–487, Jul. 1989.

[42] H. Cremer, M. Marchand, and P. Pestieau, "Evading, auditing and taxing," *Journal of Public Economics*, vol. 43, no. 1, pp. 67–92, Feb. 1990.

[43] N. D. Melumad and D. Mookherjee, "Delegation as commitment: The case of income tax audits," *The RAND Journal of Economics*, vol. 20, no. 2, pp. 139–163, 1989.

[44] J. Greenberg, "Avoiding tax avoidance: A (repeated) game-theoretic approach," *Journal of Economic Theory*, vol. 32, no. 1, pp. 1–13, 1984.

[45] B. Erard and J. S. Feinstein, "Honesty and evasion in the tax compliance game," *The RAND Journal of Economics*, vol. 25, no. 1, pp. 1–19, 1994a.

[46] M. J. Graetz and R. Bursick, "Conscience, significant others, and rational choice: Extending the deterrence model," *Journal of Law and Economic Organizations*, vol. 2, no. 1, pp. 1–32, 1986.

[47] L. Harris and Associates, "Tax payer opinion survey," US Internal Revenue Service, Washington D.C., Tech. Rep., Jan. 1987.

[48] J. Martinez-Vazquez and M. Rider, "Multiple modes of tax evasion: Theory and evidence," *National Tax Journal*, vol. 58, no. 1, pp. 51–76, 2005.

[49] S. C. Rhoades, "The impact of multiple components reporting on tax compliance and audit strategies," *The Accounting Review*, vol. 74, no. 1, pp. 63–85, Jan. 1999.

[50] C. Adams, *For Good and Evil: The Impact of Taxes on the Course of Civilization.* Lanham MD: Madison Books, 1993.

[51] H. M. Groves, "Empirical studies of income-tax compliance," *National Tax Journal*, vol. 11, no. 4, pp. 291–301, Dec. 1958.

[52] H. H. Hinrichs, "Unreporting of capital gains on tax returns or how to succeed in gainsmanship without actually paying taxes," *National Tax Journal*, vol. 17, pp. 158–163, Jun. 1964.

[53] C. T. Clotfelter, "Tax evasion and tax rates: An analysis of individual returns," *The Review of Economics and Statistics*, vol. 65, no. 3, pp. 363–373, Aug. 1983.

[54] J. Tobin, "Estimation of relationships for limited dependent variables," *Econometrica*, vol. 26, no. 1, pp. 24–36, Jan. 1958.

[55] J. S. Feinstein, "An econometric analysis of income tax evasion and its detection," *The RAND Journal of Economics*, vol. 22, no. 1, pp. 14–35, 1991.

[56] ——, "The safety regulation of U.S. nuclear power plants: Violations, inspections and abnormal occurrences," *Journal of Political Economy*, vol. 97, no. 1, pp. 115–154, Feb. 1989.

[57] B. Erard and J. S. Feinstein, "An econometric model of the tax compliance game," Mimeo, Stanford University Graduate School of Business, Tech. Rep., 1990.

[58] N. Friedland, S. Maital, and A. Rutenberg, "A simulation study of income tax evasion," *Journal of Public Economics*, vol. 10, no. 1, pp. 107–116, Aug. 1978.

[59] J. Alm, B. R. Jackson, and M. McKee, "Estimating the determinants of tax payer compliance with experimental data," *National Tax Journal*, vol. 45, no. 1, pp. 107–114, Mar. 1992b.

[60] J. Slemrod, "An empirical test of tax evasion," *The Review of Economics and Statistics*, vol. 67, no. 2, pp. 232–238, May. 1985.

[61] C. W. Christian and S. Gupta, "New evidence on secondary evasion," *The Journal of the American Taxation Association*, vol. 15, no. 1, pp. 72–93, 1993.

[62] P. J. Beck, J. S. Davis, and W.-O. Jung, "Experimental evidence on taxpayer reporting under uncertainty," *Accounting Review*, vol. 66, no. 3, pp. 535–558, Jul. 1991.

[63] J. A. Dubin and L. L. Wilde, "An empirical analysis of federal income tax auditing and compliance," *National Tax Journal*, vol. 41, no. 1, pp. 61–74, Mar. 1988.

[64] A. D. Witte and D. F. Woodbury, "The effect of tax laws and tax administration on tax compliance: The case of the U.S. individual income tax," *National Tax Journal*, vol. 38, no. 1, pp. 1–13, Mar. 1985.

[65] K. J. Beron, H. V. Tauchen, and A. D. Witte, "The effect of audits and socioeconomic variables on compliance," in *Why people pay taxes: Tax compliance and enforcement*, J. Slemrod ed. Ann Arbor: University of Michigan Press, 1992, pp. 67–89.

[66] W. W. Pommerehne and B. S. Frey, "The effects of tax administration on tax morale," *Unpublished Manuscript, University of Saarbrucken*, 1992. [Online]. Available: http://www.opengrey.eu/item/display/10068/222947

[67] S. Klepper and D. Nagin, "The anatomy of tax evasion," *Journal of Law, Economics, & Organization*, vol. 1, no. 5, pp. 1–24, 1989.

[68] S. B. Long and J. A. Swingden, "The role of third-party information reporting in our self-assessing income tax system," in *Research conference report: How do we affect taxpayer behavior*. Washington D.C.: Internal Revenue Service Document, 1990, no. 7302, pp. 3–91.

[69] H. S. J. Robben, "Decision frame and opportunity as determinants of tax cheating: An international experimental study," *Journal of Economic Psychology*, vol. 11, no. 3, pp. 341–364, 1990.

[70] B. Erard, "The influence of tax audits on reporting behavior," in *Why people pay taxes: Tax compliance and enforcement*, J. Slemrod ed. Ann Arbor: University of Michigan Press, 1992, pp. 95–114.

[71] P. Webley, "Audit probabilities and tax evasion in a business simulation," *Economic Letters*, vol. 25, no. 3, pp. 267–270, 1987.

[72] M. W. Spicer and R. E. Hero, "Tax evasion and heuristics: A research note," *Journal of Public Economics*, vol. 26, no. 2, pp. 263–267, Mar. 1985.

[73] O. H. Erekson and D. H. Sullivan, “A cross-section analysis of IRS auditing,” *National Tax Journal*, vol. 41, no. 2, pp. 175–189, Jun. 1988.

[74] J. Alm, R. Bahl, and M. N. Murray, “Tax structure and tax compliance,” *The Review of Economics and Statistics*, vol. 72, no. 4, pp. 603–613, Nov. 1990.

[75] J. A. Dubin, M. J. Graetz, and L. L. Wilde, “The effect of audit rates on the federal individual income tax, 1977-1986,” *National Tax Journal*, vol. 43, no. 4, pp. 395–409, Dec. 1990.

[76] J. Alm, R. Bahl, and M. N. Murray, “Audit selection and income tax underreporting in the tax compliance game,” *Journal of Development Economics*, vol. 42, no. 1, pp. 1–33, 1993.

[77] J. Slemrod, “Cheating ourselves: The economics of tax evasion,” *The Journal of Economic Perspectives*, vol. 21, no. 1, pp. 25–48, 2007.

[78] C. W. Christian, “Voluntary compliance with the individual income tax: Results from 1988 TCMP study,” *The IRS Research Bulletin*, vol. 1500, pp. 35–42, 1994.

[79] M. Lubell and J. T. Scholz, “Cooperation, reciprocity, and the collective-action heuristic,” *American Journal of Political Science*, vol. 45, no. 1, pp. 160–178, Jan. 2001.

[80] S. Klepper, M. Mark, and D. Nagin, “Expert intermediaries and legal compliance: The case of tax preparers,” *Journal of Law and Economics*, vol. 34, no. 1, pp. 205–229, Apr. 1991.

[81] B. Erard, “Taxation with representation: An analysis of the role of tax practitioners in tax compliance,” *Journal of Public Economics*, vol. 52, no. 2, pp. 163–197, Sep. 1993.

[82] J. Alm, B. Erard, and J. S. Feinstein, “The relationship between federal and tax audits,” in *Empirical Foundations of Household Taxation.* University of Chicago Press, 1996, pp. 235–277.

[83] B. Erard, “Self-selection with measurements errors: A microeconometric analysis of the decision to seek tax assistance and its implications for tax compliance,” *Journal of Econometrics*, vol. 81, no. 2, pp. 319–356, Dec. 1997.

[84] M. Hanlon and S. Heitzman, “A review of tax research,” *Journal of Accounting and Economics*, vol. 50, no. 2-3, pp. 127–178, Dec. 2010.

[85] B. Erard and J. S. Feinstein, “The role of moral sentiments and audit perceptions in tax compliance,” *Public Finance=Finances publiques*, vol. 49, pp. 70–89, 1994b, supplement.

[86] M. W. Spicer and S. B. Lundstedt, "Understanding tax evasion," *Public Finance=Finances publiques*, vol. 31, no. 2, pp. 295–305, 1980.

[87] P. Webley, H. Elffers, and D. Hessing, *Tax evasion: An experimental approach.* Cambridge: Cambridge University Press, 1991.

[88] W. W. Pommerehne, A. Hart, and B. S. Frey, "Tax morale, tax evasion and the choice of policy instruments in different political systems," *Public Finance=Finances publiques*, vol. 49, pp. 52–69, 1994, supplement.

[89] K. Krause, "Tax complexity: Problem or opportunity?" *Public Finance Review*, vol. 28, no. 5, pp. 395–414, Sep. 2000.

[90] F. Cowel, *Cheating the Government: The Economics of Evasion.* Cambridge, MA: MIT press, 1990.

[91] M. Daunton, "Trusting leviathan: British fiscal administration from the napoleonic wars to the second world war," *In Trust and Governance*, vol. 57, pp. 102–134, 2004, New York: Russel Sage Foundation.

[92] J. L. Mikesell, "Amnesties for state tax evaders: The nature of and response to recent programs," *National Taxation Journal*, vol. 39, no. 4, pp. 507–525, Dec. 1986.

[93] A. Das-Gupta and D. Mookherjee, *Tax Amnesties in India: An Empirical Evaluation.* Boston University, Institute for Economic Development, 1995, discussion Paper no. 53.

[94] R. C. Fisher, J. H. Goddeeris, and J. C. Young, "Participation in tax amnesties: The individual income tax," *National Tax Journal*, vol. 42, no. 1, pp. 15–27, Mar. 1989.

[95] J. Slemrod, "The economics of corporate tax selfishness," *National Tax Journal*, vol. 57, no. 4, pp. 877–899, Dec. 2004.

[96] M. C. Jensen and W. H. Meckling, "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, vol. 3, no. 4, pp. 305–360, Oct. 1976.

[97] K. Chen and C. Y. C. Chu, "Internal control versus external manipulation: A model of corporate income tax evasion," *The RAND Journal of Economics*, vol. 36, no. 1, pp. 151–164, 2005.

[98] K. J. Crocker and J. Slemrod, "Corporate tax evasion with agency costs," *Journal of Public Economics*, vol. 89, no. 9-10, pp. 1593–1610, Sep. 2005.

[99] M. A. Desai, A. Dyck, and L. Zingales, "Theft and taxes," *Journal of Financial Economics*, vol. 84, no. 3, pp. 591–623, Jun. 2007.

[100] D. P. Bertsekas, *Dynamic Programming and Optimal Control*, 3rd ed. Belmont MA: Athena Scientific, 2005.

[101] J. E. Smith and K. F. McCardle, "Structural properties of stochastic dynamic programs," *Operations Research*, vol. 50, no. 5, pp. 796–809, Oct. 2002.

[102] K. Kanellopoulos, I. Kousoulakos, and B. Rapanos, "The underground economy in Greece: What official data show," *Greek Economic Review*, vol. 14, pp. 215–236, 1992.

[103] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*, 4th ed. McGraw Hill Higher Education, 2002.

[104] R. C. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.

[105] R. Feldbrugge, "Using reinforcement learning to make optimal use of available power and improving overall speed of a solar-powered boat," Master's thesis, University of Groningen. Faculty of Mathematics & Natural Sciences, 2010. [Online]. Available: http://www.ai.rug.nl/ mwiering/