On the Design of Incentive Mechanisms in the Presence of Externalities

Miguel A. Sanchez Villalba
London School of Economics and Political Science

Dissertation submitted for the degree
DOCTOR OF PHILOSOPHY
in the Field of Economics

London, August 2007

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science (University of London) is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

Date Signature

Abstract

In a world of ever increasing interrelations among people, firms and countries, externalities become more and more significant as time passes by and, consequently, incentive schemes that overlook them will fail to fulfil their objectives.

This dissertation analyses how those schemes are affected by the presence of externalities by focusing on some special cases. Depending on the type of externality involved, the cases can be labelled as either "endogenous" (if it is created by the designer's choice of scheme) or "exogenous" (otherwise).

The analysis of the latter case finds that delegation of contracting rights improves the efficiency of a multi-agent organisation because it closes the gap between society's and agents' marginal benefits. The analysis of endogenous externalities, on the other hand, shows that sometimes the designer's optimal action is to create an externality between agents and to take full advantage of the new interactions thus generated.

These findings indicate that when externalities are present incentive schemes can be radically affected and, moreover, that the mechanism designer may have incentives to create externalities between agents in order to advance her goals. These effects are illustrated using leading examples and experimental data.

Acknowledgements

There is a long list of people who in one way or another contributed to this achievement.

First of all, my family, Dad, Mum, Jose, Juan, Maria and Lucas, whose help and support was the engine that kept me going. Then my supervisor, Frank Cowell, who was priceless as a source of both knowledge and encouragement. Third, my co-author Rafael Hortala-Vallve, who shared the ups and downs of this adventure.

Also a large group of people who showed an interest in my work and gave me valuable feedback: my advisers Jonathan Leape, Bernardo Guimaraes and Georg Weizsacker, my fellow students Oliver Denk and Giovanni Ko, and a number of people including Madhav Aney, Munia Bandyopadhay, Ralph Bayer, Jordi Blanes, Mariano Selvaggi, Rocco Macchiavelo and many others who attended my seminars.

Last, but definitely not least, my friends here in London: those from STICERD (Sue, Leila, Pedro, Carmen, etc.), my football buddies (Eduardo, Marco, Moacir, Josep, Rui, Alex, Carlos, etc.), my flatmates (Boris, Laszlo and Paul), and many others (Veronica, Cecilia, Dragomir, Fernando, Ana, Coco, Maria, Nuria, Nicolas, etc.). All of them helped me to overcome the bad times and to enjoy the good ones.

To all of you, my eternal gratitude.

Thank you.

Contents

List of Figures

So much to do, so little time.

Chapter 1

Introduction

The objective of this thesis is to analyse the impact of externalities on the design of optimal incentive schemes. This is a relevant topic because in an increasingly "globalised" world the importance of the associated external effects is also on the rise. This increased sensitivity to other people's actions means that actors that seemed to be entirely unrelated to one another in the past are now firmly connected, an example being the late 90s financial crisis that hit countries as disparate and far apart as Russia, Indonesia and Argentina.

Unsurprisingly, mechanisms designed to operate optimally in those once self-contained units will become inefficient if kept in place in the interconnected scenario, so it is worth exploring how those mechanisms should be adjusted to achieve their goals under the new circumstances.

Such exploration is, however, a vast enterprise, and so the present dissertation will instead analyse a few special cases in-depth. These cases can be classified, depending on the type of externality involved, into two categories, namely, exogenous and endogenous.

An externality of the first type is one whose existence is *not* the result of the designer's choice of scheme, i.e., one that links actors' payoffs via technological or preference-related channels. Its analysis (jointly undertaken with Rafael Hortala-Vallve and presented in part I of the thesis) focuses on the choice of the optimal contracting structure in a multi-agent organisation and uses the British railway system as a leading example. The main result is that in a hidden action scenario in which jointly-produced output is the only contractible variable, the associated positive externality that leads to inefficiently low effort can be mitigated by delegating contracting rights. Intuitively, agents with contracting rights realise the effect of their actions on other agents' wellbeing, thus decreasing the gap between their private- and the social- marginal benefit and helping to partially internalise the externality.

The second part of the thesis investigates the case of endogenous externalities, that is, those created by the mechanism designer as part of the optimal scheme. These externalities are entirely different from the standard, "exogenous" type because, unlike the latter, in the absence of the designer's actions the externality would not exist. Their existence is, however, fundamental to the mechanism devised by the principal and, as such, plays a crucial role in the provision of incentives.

The study of the endogenous case is divided into two complementary modules, namely, the theoretical analysis (undertaken in chapter 4) and the empirical one (in chapter 5).

Using the design of anti-evasion policy as a leading example, the theoretical analysis investigates the problem of choosing an auditing policy when a group of similar taxpayers are affected by common income shocks and are imperfectly informed about the "type" of tax agency they face (i.e., how tough on evasion the agency is). The analysis finds that the optimal strategy consists of following a *"contingent rule"*, namely, auditing a given taxpayer with a probability that is a (weakly) increasing function of her fellow taxpayers' declarations. Intuitively, since taxpayers are very similar to each other, other people's declarations are informative about the likelihood of a particular one being an evader. This policy endogenously creates a negative externality between taxpayers: the expected utility of any of them is negatively related to her probability of being audited, which is a (weakly) increasing function of every other taxpayer's declarations. This gives rise to the presence of strategic complementarities between declarations and, consequently, to a coordination game between taxpayers. The associated problems of multiple equilibria are avoided by the fact that taxpayers do not know exactly the "type" of agency they face (soft/tough on evaders), but only get noisy private signals about it. This heterogeneity in information sets ensures, by the tenets of the global game technique, a unique equilibrium.

The predictions of the "Tax Evasion as a Global Game" (TEGG) model of chapter 4 are tested in chapter 5 using data from a computerised experiment where participants interacted with each other in situations that resembled the tax compliance game. The findings suggest strong support for the superiority of the TEGG model's "contingent rule" over the "cut-off rule" usually prescribed by the literature (a comparison between the TEGG rule and other rules is part of my future research agenda). Also, the estimated coefficients have the signs predicted by the global game's comparative statics, but the data seem to reject the idea that people use higher-order beliefs when making decisions (though they may learn/adapt and make the same decisions as if using them).

The main conclusion to draw is that the presence of externalities significantly affects the design of incentive schemes. Some times they could work against the designer (like in part I, leading to inefficiently low levels of effort), some other times they can work for her (like in part II, increasing the chances of taxpayers' mistakes by creating a more complex game). Also, externalities can lead to the radical modification of incentive schemes (delegation being preferred over centralisation in part I) or become an integral and crucial part of them (modifying altogether the nature of the game in the process, as in part II). Furthermore, while in part I the externality is exogenously generated by the non-contractibility of output, in part II it is the designer itself that, by following the "contingent" policy, creates an otherwise non-existent externality.

Chapter 2

Literature review

The role of externalities in economics has been recognised for a long time, to the point that their absence is required for the Pareto-optimality of a competitive equilibrium to hold. Their presence, on the other hand, robs Adam Smith's "Invisible Hand" of its ability to generate efficient outcomes, and so it generates the need for policies that correct or at least mitigate their effects. From the quintessential example of the tragedy of the commons (Hardin (1968)) to the idea of geographical "clustering" (Krugman (2005)) and from tax competition (Devereux and Pearson (1990)) to network effects (Katz and Shapiro (1985)) and self-control (Thaler and Shefrin (1981)), externalities have been thoroughly analysed by economists –usually with the goal of reducing the ensuing inefficiency.

The first part of the thesis follows this path and intends to design a mechanism that minimises the detrimental effect of the positive externality arising in an organisation in which contracts can only be contingent on jointly-produced output: each agent chooses her own effort based on her private marginal benefit and private marginal cost, without realising the positive effect her effort has on the utility of every other agent, and thus exerting too little effort compared to society's optimum. The solution suggested here consists of the delegation of contracting rights from the Principal to one of the agents, who then subcontracts with the second one. This way, the intermediate agent becomes residual claimant, thus noticing the effect her effort has on the second agent and increasing her effort accordingly (as well as the overall level of efficiency).

The seminal reference on moral hazard in teams is Holmstrom (1982). Most of the subsequent literature focused on the problem of collusion among the agents (see Tirole (1986)), an issue that does not play a role here because –given that output is the same for everyone– there are no "lucky" agents that can compensate "unlucky" ones. In fact, the only possibility of collusion among agents consists in forming a kind of "cartel" and making their effort decisions jointly, thus internalising the externality they impose on each other but also increasing the Principal's profits. The literature has also investigated the role of hidden information (like Melumad et al. (1995), who study the effects of different productivities), the design of monitoring schemes (e.g., Faure-Grimaud et al. (2003) or Baliga and Sjolstrom (1998)) and the allocation of different tasks to different agents (as in Prendergast (1995)). Though certainly all these topics are relevant, they have been explored frequently and in detail, and so will not be part of the present analysis, which focuses instead on the allocation of contracting rights. Closer to the topic of interest are the papers by Itoh (Itoh (1991) and Itoh (1994)), that consider the allocation of tasks when agents have incentives to help other agents and the desirability of implementing relative performance schemes. These factors, however, do not affect the results in settings where output is produced jointly, as is the case here. The closest references to our study are those of Macho-Stadler and Pérez-Castrillo (1998) and Macho-Stadler and Jelovac (2002). They compare different contracting structures (using the Spanish health sector as their leading example) but the key aspect of their work is the timing of events rather than delegation and the internalisation of the externality.

The second part of the thesis deals with a less frequent scenario. In this case, the externality is not an inherent part of the economy under consideration but the result of the deliberate choice of the mechanism designer. The leading example used as illustration is that of a tax agency that has to decide its auditing strategy while knowing that the taxpayers' incomes are subject to common shocks. In such a setting, it is found that the agency's optimal strategy requires the probability of auditing a taxpayer that declares low income to be a (weakly) increasing function of the declarations of other taxpayers. This way, the agency creates a negative externality between agents: the higher the declarations of other taxpayers, the higher the probability of being caught if I evade, and so the higher my incentives to comply. This, in turn, means that taxpayers' declarations are strategic complements and that the agency's policy forces taxpayers to play a coordination game between them, a game that –as the externality that generates it– would not exist if the agency did not create it.

Unlike the previous case, the literature on this kind of "endogenous" externality is rather scarce. Papers like Itoh (1991) do consider how incentive schemes could amplify or temper the effects of externalities, but the externalities are not created by the designer. A similar problem affects the analysis of "club goods" (Cornes and Sandler (1996)). Morgan (2000)'s innovative analysis of lotteries as means to finance the provision of public goods is one of the few studies that can be included in the literature on endogenous externalities. In the particular area of tax compliance considered here, the closest reference is Basseto and Phelan (2004), who analyse how the optimal tax system can lead to "tax riots", though they ignore the issue of the optimal auditing strategy that plays a critical role in the present analysis. Several studies, though, do consider the effects of externalities, but they are not endogenous. This is the case of, among others, Benjamini and Maital (1985), who introduce psychological costs and find that they lead to "epidemics" of compliance or evasion. Others like Fortin et al. (2004) and Myles and Naylor (1996) rely on social norms and model utility as an increasing function of conformity. This is also the case of Kim (2005), who uses the same equilibrium selection technique used here, namely, the global game approach (Carlsson and van Damme (1993), Morris and Shin (2002b)). For the empirical part the literature is limited as well, in spite of the large number of experiments framed as tax compliance problems that span from the seminal one by Baldry (1986) who compares tax evasion to gambling, to papers that study the connection between tax evasion and voting (Feld and Tyran (2002)). But the closest reference is, undoubtedly, Alm and McKee (2004), where tax compliance is analysed as a coordination game. The present study goes one step further and introduces uncertainty about the agency's "type" and models tax evasion as a global game.

Part I

Exogenous Externality

Chapter 3

Delegation of contracting rights

joint with Rafael Hortala-Vallve

3.1 Introduction

When a Virgin train derailed near Graygigg (U.K.) on Friday 23 February 2007 and the sorrow associated with the casualties (including one death) settled down, old arguments regarding the organisation of the British railway system soon resurfaced.

The agency in charge of the track (Network Rail) was found guilty of negligence and several people (including Virgin chairman Sir Richard Branson) demanded a greater say in track maintenance by train companies. Britain's biggest train company First Group and the Conservative party went further and proposed that train operators should do their own maintenance, thus reverting to pre-privatisation arrangements.

Following the 1993 Railways Act, the state owned British Rail was privatised. This implied the separation between maintenance (undertaken by Railtrack) and train operation in an attempt to improve the efficiency of the system. Nevertheless the industry's safety record suffered (5 accidents, 59 deaths) and so in 2002 the government created Network Rail (NR) to replace Railtrack. The latter's reliance on sub-contracted personnel was blamed for the crashes, so since its inception the new agency has only used its own staff for maintenance tasks.

The newly created Network Rail is monitored by the Office of Rail Regulation (ORR) but operates as a commercial business. Half of its income comes from the government and the rest is raised by access charges paid by the Train Operating Companies (TOC). The three actors (ORR, NR and TOC) can therefore be seen as the building blocks of a team production problem in which output is the quality of train services and the inputs are the efforts of the agents regarding the provision of the two basic determinants of quality, namely, track maintenance (NR) and train operation (TOC).[1] These efforts are not verifiable, and so contracts can only be contingent on output, i.e., on observable proxies for quality such as customer satisfaction or punctuality.[2] The combination of joint production and contracts being contingent only on output, therefore, creates a positive externality between the agents: the effort exerted by one of them increases output, which in turn increases not only its own payoff but that of the other party's as well.[3] Agents, however, ignore this interaction when they make their individual effort decisions, with the end result being inefficiently low levels of effort being exerted. We show that the principal (ORR) can mitigate this problem by simply *delegating* contracting rights to one of the agents. In this way, the latter becomes residual claimant and realises the positive externality that her effort has on the other party. This means that the difference between private and social marginal benefits is reduced and that overall efficiency increases as a result.

In the absence of contractual restrictions the Revelation Principle states that a delegated structure cannot improve upon a centralised one. In our setting, the Principal could *centralise the delegated structure* by offering one of the agents a contract contingent on the subsequent contracts this agent writes with her subordinates. However, this practice is not common in reality: a Manager/Subcontractor is unlikely to accept a contract that is contingent on her own actions as this will strip her of all freedom of choice. Indeed, no such contracts control the relationships between the component parties of the railway system, and so it is under the assumption that those contracts are ruled out that we find that delegation is optimal.

The seminal reference on moral hazard in teams is Holmstrom (1982). Most of the contract

[1] Clearly this is an extremely simplified model of the railway industry that ignores many (and important) industry-related elements. However, our goal is not the suggestion of policies applicable to this particular industry: we only intend to use this very stylised model of the British railway system to illustrate the problem of choosing the contracting structure of an organisation.

[2] Alternatively, as is common in the moral hazard literature, efforts can be considered to be verifiable and verification to be costly, a reasonable assumption in this case in which monitoring a vast network is needed.

[3] If tracks are not properly maintained, trains need to slow down and punctuality may suffer. Conversely, if train companies do not train their drivers or maintain their trains properly, they cannot take advantage of well maintained tracks. *Efforts* could be interpreted as maintenance or training activities but one could also think of *efforts* as investments as long as they are difficult (or too costly) to verify in a court of law and hence no contracts can be written upon them.

theory literature on teams analyses collusion and ways in which the principal can avoid it (see Tirole (1986) and subsequent literature). In our model, however, collusion can only benefit the Principal because it implies the agents make decisions as a unit, so that they internalise the positive externality they exert on each other and increase their effort accordingly. This may yield a positive surplus to the agents if the principal does not anticipate such behaviour but the Principal's profits would still increase relative to the situation where no collusion occurs. The literature has also analysed repeatedly the role of hidden information and the effects of monitoring among agents (see Baliga and Sjolstrom (1998) and Faure-Grimaud et al. (2003)), and so we do not consider these issues here.

In a similar vein to our work, Felli and Hortala-Vallve (2007) show how delegation can costlessly prevent collusion between a Supervisor and an Agent. Itoh (1991) and Itoh (1994) focus instead on the best way to allocate tasks among agents and whether the principal benefits from offering relative performance schemes. Finally, the closest references to our study are those of Macho-Stadler and Pérez-Castrillo (1998) and Macho-Stadler and Jelovac (2002). They compare different contracting structures in the health sector using a binary effort model but the key aspect of their work is the timing of events rather than delegation.

3.2 The Model

The model presented below is based on Holmstrom and Milgrom (1987).

A Principal hires two identical agents ($i = 1, 2$) to undertake the production of a joint output, x. Agents are assumed identical in order to concentrate solely on the effects of different contracting structures. Output is assumed to be normally distributed with mean μ and variance $\sigma^2 > 0$. The expected output of the project increases with the effort exerted by the agents: $\mu(e_1, e_2) = e_1 + e_2$.

Efforts are *non-verifiable* and contracts ($w_i(x)$, $i = 1, 2$) can only be contingent on realised output. Moreover, we restrict contracts to be linear in output, i.e. $w_i(x) = a_i + b_i x$, $i = 1, 2$.[4] Contract offers are assumed to be public (i.e. observed by all parties). We assume the Principal can credibly commit to her proposed policy, thus avoiding the

[4] Whenever contracts are not constrained to be linear and arbitrarily large punishments are allowed, first best can be achieved. However, this scenario is not realistic within our setting of passenger rail services.

The question of limited liability is explored in appendix A.1.

The possibility of writing contracts contingent on profits is analysed in appendix A.2.

issue of renegotiation. Hereafter, capital letters denote aggregate variables: $E := e_1 + e_2$, $W(x) := w_1(x) + w_2(x)$...

The Principal is risk neutral and seeks to maximise the expected output minus total wages. Agents are assumed to be risk averse with *constant absolute risk aversion* (CARA) utility and index of risk aversion $r > 0$. The disutility of effort is $\frac{1}{2}e_i^2$. The setting just described allows us to rewrite agents' expected utility functions in terms of their certainty equivalent: $a_i + b_i \cdot E - \frac{\gamma}{2}b_i^2 - \frac{1}{2}e_i^2$, where $\gamma := r\sigma^2 > 0$.[5]

3.2.1 First Best

When efforts are contractible the Principal maximises expected profits and ensures that the agents' *Participation Constraints* are satisfied, i.e., that the expected utility they derive from the contract is not lower than their reservation utility $\underline{U}$ (which, without loss of generality, can be normalised to 0). The Principal's program is therefore:

$$\begin{gathered}\max_{\{w_1(x), e_1, w_2(x), e_2\}} E_x\{\pi(x)\} \\ \text{s.t.}\ \{PC_i : E_x\{U(w_i(x), e_i)\} \geq \underline{U},\ i = 1,2\end{gathered} \tag{3.1}$$

The solution to the problem is such that the levels of effort exerted by the agents are $e_1^* = e_2^* = 1$, the wages are $w_i^*(x) = \frac{1}{2}$, $i = 1,2$ and the Principal's expected profit is $E\pi^* = 1$. As expected, risk averse agents face no risk and (since they are identical to each other) are treated identically.

3.2.2 Centralised Second Best

The second best situation requires providing output-based incentives and insuring the agents against the resultant risk. This means that when the Principal optimises, she needs not only to ensure that both agents' participation constraints are satisfied, but

[5] The neutrality of the principal is not fundamental for the analysis. If she were risk averse, her objective function (assuming CARA utility) would differ from that of a risk averse person only in terms of a constant that reflects the disutility associated to the risk borne by her. Her choice of contracts, however, will not be affected.

also to take into account that they will react optimally to the contract they are offered (*Incentive Constraints*). Her program, therefore, reads as follows:

$$\max_{\{w_1(x),e_1,w_2(x),e_2\}} E_x\left\{\pi\left(x\right)\right\}$$
$$\text{s.t.} \begin{cases} PC_i : E_x\left\{U\left(w_i\left(x\right),e_i\right)\right\} \geq \underline{U},\ i=1,2 \\ IC_i : e_i \in \arg\max_{\dot{e}_i} E_x\left\{U\left(w_i\left(x\right),\dot{e}_i\right)\right\},\ i=1,2 \end{cases} \tag{3.2}$$

The second best efforts and contracts are $e_1^{**}=e_2^{**}=\frac{1}{1+\gamma}$ and $w_i^{**}(x)=a^{**}+\frac{1}{1+\gamma}\cdot x, i=1,2$ respectively, where a^{**} makes the participation constraints binding. The associated expected profit is $E\pi^{**}=\frac{1}{1+\gamma}$. Henceforth we will call this situation the *centralised second best (CSB)*.

It is worth noting, however, that this solution to the 2-agent problem is equivalent to the solution of two independent 1-agent problems. That is, the contracts offered by the Principal do not take into account the existence of the positive externality between the agents that increases *all* agents' expected wages when any of them increases her effort. Due to contractual restrictions, the Principal is unable to induce agents to internalise this externality when contracting with them in a centralised way. As a matter of fact, our simple model tells us that the best possible arrangement involves integrating all the activities so that effort decisions are made jointly.[6] However, this is not possible in the case of the British railway industry because unmodelled aspects of it prevent such arrangements being reached, especially the EU Directive 91/440 that requires all EU member states to separate '...the management of railway operation and infrastructure from the provision of railway transport services, separation of accounts being compulsory and organisational or institutional separation being optional'. This means that a realistic model of the current situation has to reflect this constraint, and this is precisely what we do in the following section.

3.2.3 Delegated Second Best

The Principal can improve on the Centralised Second Best by changing the contracting structure to establish a hierarchy between the agents. The rationale for the improvement

[6]Such scenario is explored in appendix A.3. A comparison between this "Cartel" case and the Centralised- and Delegated- Second Best cases is also presented there.

derives from the fact that the agent higher in the hierarchy recognises the positive externality her effort has on the other agent and is willing to exert more effort at no extra cost.

Under the delegated structure, the Principal contracts with one agent, who then subcontracts with the remaining one. The timing of the game is as follows: the Principal offers a contract $W(x) = A + Bx$ to Agent 1, who subsequently decides her own level of effort together with the contract she offers to Agent 2, $w_2(x) = a + bx$; then Agent 2 decides her level of effort. If their participation constraints are met, agents exert effort. Finally, output is realised and payments are made.

The Principal's program, therefore, reads as follows:

$$\max_{\{W(x),e_1,w_2(x),e_2\}} E_x\{\pi(x)\} \tag{3.3}$$

$$\text{s.t.}\begin{cases} PC_1 : E_x\{U(W(x) - w_2(x), e_1)\} \geq \underline{U} \\ IC_1 : \begin{cases} \max\limits_{\{\dot{e}_1, w_2(x), e_2\}} E_x\{U(W(x) - w_2(x), \dot{e}_1)\} \\ \text{s.t.}\begin{cases} PC_2 : E_x\{U(w_2(x), e_2)\} \geq \underline{U} \\ IC_2 : e_2 \in \arg\max\limits_{\dot{e}_2} E_x\{U(w_2(x), \dot{e}_2)\} \end{cases} \end{cases} \end{cases}$$

Agent 1's program has a unique solution: $e_1 = B$, $e_2 = b = \frac{1+\gamma}{1+2\gamma}B$ and a is such that the participation constraint of Agent 2 is binding. Since Agent 2's participation constraint enters the Lagrangean for Agent 1's programme directly, it is clear that Agent 1 now internalises the effect of her effort on Agent 2's payoff:in

$$\mathcal{L} = (A - a) + (B - b)\cdot E - \frac{\gamma}{2}(B - b)^2 - \frac{1}{2}e_1^2 + \lambda\left(a + b\cdot E - \frac{\gamma}{2}b^2 - \frac{1}{2}e_2^2\right) + \varpi(e_2 - b)$$

Its derivative with respect to e_1 should be equal to zero at the optimum,

$$\frac{\partial\mathcal{L}}{\partial e_1} = (B - b) - e_1 + \lambda\cdot b = 0$$

The above condition differs from the incentive constraint under the centralised structure because it includes the term $(\lambda \cdot b)$. This term captures the positive externality that Agent 1 has on Agent 2: more effort by Agent 1 increases the expected output which in turn relaxes the participation constraint of Agent 2.

The unique solution has all participation constraints binding and the following effort levels:

$$\begin{cases} e_1^d = B^d = \frac{2+3\gamma}{(1+\gamma)(2+\gamma)} \\ e_2^d = b^d = \frac{2+3\gamma}{(1+2\gamma)(2+\gamma)}. \end{cases}$$

The associated expected profit is

$$E\pi^d = \frac{1}{2}\frac{(2+3\gamma)^2}{(1+2\gamma)(2+3\gamma+\gamma^2)} \tag{3.4}$$

3.3 Results

Proposition 3.1 *The delegated structure always yields a higher expected profit to the Principal than the centralised one. The relative gains are higher the higher the risk of the project and/or the higher the index of risk aversion of the agents.*

The proof is immediate from comparing $E\pi^{**} = \frac{1}{1+\gamma}$ and $E\pi^d = \frac{1}{2}\frac{(2+3\gamma)^2}{(1+2\gamma)(2+3\gamma+\gamma^2)}$ and showing that $E\pi^d/E\pi^{**}$ is increasing in γ.[7]

Before moving to the next proposition, remember that the disutility of risk borne by Agent i is equal to $\frac{r}{2}Var(w_i(x))$. Thus, Agent 2's disutility of risk is $\frac{\gamma}{2}\left(b^d\right)^2$, and Agent 1's is $\frac{\gamma}{2}\left(B^d - b^d\right)^2$.

Proposition 3.2 *Under the delegated structure, Agent 1 exerts more effort than Agent 2. Moreover, both agents exert more effort than under the second best structure, i.e. $e_1^d > e_2^d > e_1^{**} = e_2^{**}\ \forall\gamma > 0$. In terms of risk, Agent 1 bears less risk than Agent 2 ($b^d > B^d - b^d$) and the latter bears more risk under the delegated structure than under the centralised one ($b^d > b^{**}$).*

The effort results lead to the conclusion that expected output is greater under delegation than under centralisation, and the prediction that managers (agents higher in the

[7]Note that the system's total expected welfare (or total expected surplus, $ES \equiv E\pi + EU_1 + EU_2$) is equal to the principal's expected profits, and so this result indicates that the Delegated structure is more efficient than the Centralised one. This is a consequence of having assumed linear contracts and ignored limited liability constraints, thus ensuring that participation constraints are always binding in equilibrium, that is, $EU_1 = 0$ and $EU_2 = 0$.

When limited liability constraints are taken into consideration the equivalence between expected profits and total surplus breaks down because agents can get a strictly positive expected payoff, $EU_i > 0$.

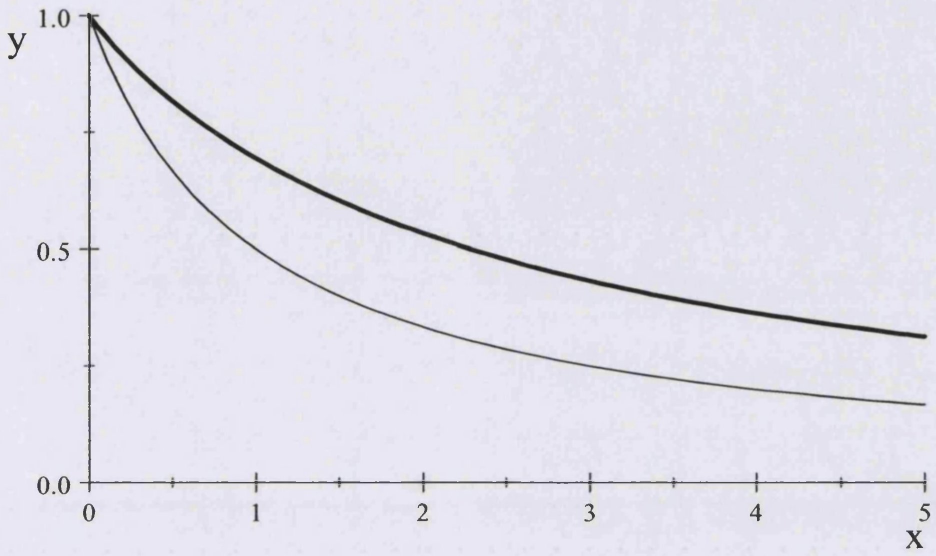

Figure 3.1: Expected Profits.
Horizontal axis: γ. Vertical axis: Thick line: $E\pi^d$; thin line: $E\pi^{**}$.

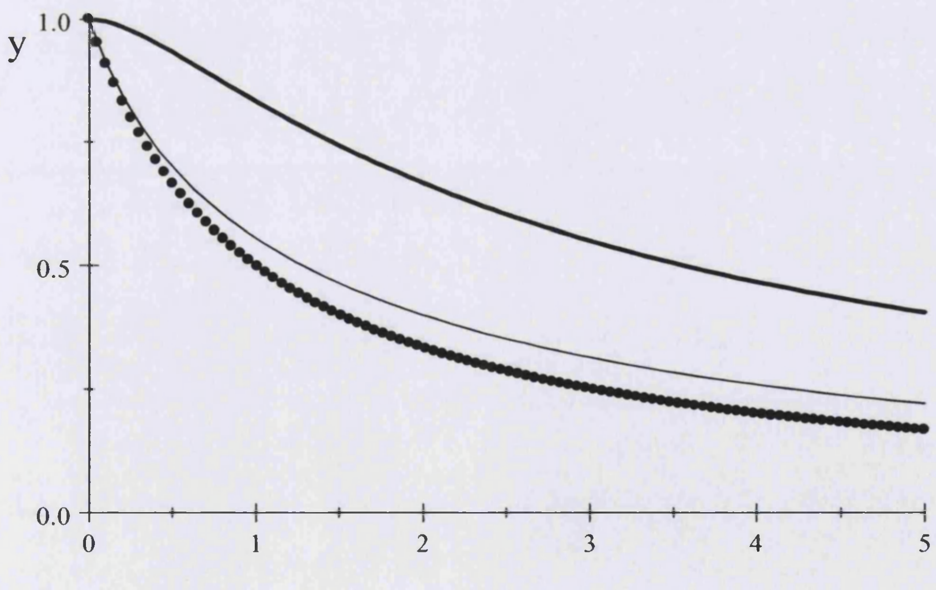

Figure 3.2: Efforts.
Horiz. axis: γ. Vert. axis: Thick solid line: e_1^d; thin solid line: e_2^d; dotted line: e^{**}.

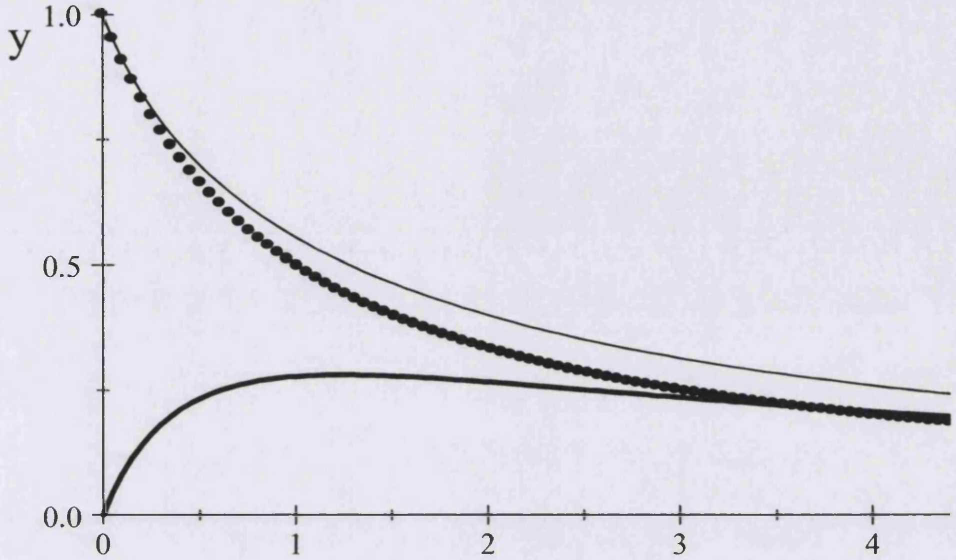

Figure 3.3: Risk.
Horiz. axis: γ. Vert. axis: Thick solid line: $B^d - b^d$; thin solid line: b^d; dotted line: b^{**}.

hierarchy) exert more effort than subordinates is in line with findings in Prendergast (1995).

From Agent 1's program we find that $\frac{\partial b^d}{\partial B^d} = \dfrac{1+\gamma}{1+2\gamma} \in \left[\frac{1}{2}, 1\right]$; i.e. whenever the Principal induces a higher effort from Agent 1, the latter also provides more incentives to Agent 2. From the Principal's perspective this generates a (second-order) *trickle down effect* that multiplies the initial (first-order) effect of an increase in B^h by increasing also b^h. This effect is decreasing in γ.

Notice that the slopes of the wage contracts, besides providing information on the optimal levels of effort, also indicate the risk borne by the agents. Moreover, while Agent 1's effort depends on her *gross* wage ($e_1^d = B^d$), her disutility of risk $\frac{\gamma}{2}\left(B^d - b^d\right)^2$ depends on her *net* wage instead. As a consequence, Agent 1 is able to transfer most of her risk to Agent 2, though at the expense of exerting more effort than her subordinate (see figure 3.4). This result also shows how the fundamental trade-off in moral hazard situations, that of incentives versus risk, is lessened by the delegation of contracting rights: for Agent 1, more incentives (greater B^d) does not mean as much extra risk $\left(B^d - b^d\right)$ as in the centralised case since, as mentioned above, she will transfer some of the risk to Agent 2 by increasing b^d.

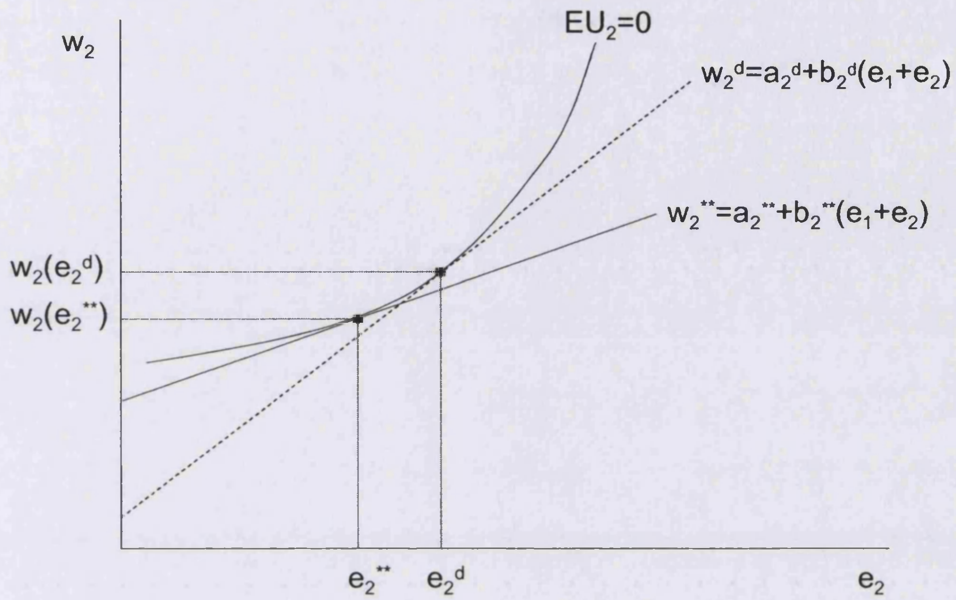

Figure 3.4: Agent 2's incentive scheme. Second Best and Delegated Second Best.

At first sight, the fact that the agent who is higher in the hierarchy bears less risk than her subordinates may seem puzzling. However, it is well known that outsourcing, subcontracting or decentralisation are ways to pass risk to those at lower levels of the hierarchy. Indeed, their tendency to generate precarious, deregulated working conditions is one of the main criticisms levelled at practices such as subcontracting in the construction industry, outsourcing in manufacturing and the privatisation of public services.[8]

3.4 Discussion

In spite of the convexity of the disutility of effort and the fact that agents are homogeneous and risk averse, the delegated structure results in an asymmetric distribution of risk and effort between agents. However, the associated internalisation of the externality overcomes the inefficiency generated by the unequal treatment of agents and overall efficiency is higher than under the centralised structure.

Two modelling assumptions are needed for our results to hold. First, each agent's effort choice depends only on the power of her own incentive scheme (slope of her contract)

[8] The analysis of the relative importance of risk and effort in an agent's disutility under each contracting structure is undertaken in appendix A.4.

and is independent of the effort exerted by the other agent (i.e. there are no strategic complementarities or substitutabilities). Consequently, any agent affects the other only by modifying the latter's participation constraint. Second, under the delegated structure Agent 1 fully internalises her externality on Agent 2, hence her effort does not depend on the distribution of the wages among the agents but only on the aggregate wage bill $(W(x))$.

When we consider a more general setting the analysis becomes ambiguous precisely because these assumptions no longer hold. For instance, strategic complementarities reinforce the pre-eminence of the delegated structure over the centralised one, while strategic substitutabilities work against the incentives of Agent 1 to exert more effort under the delegated structure.

The benefits of the delegated structure stem from the fact that Agent 1 internalises the externality. This happens when she takes advantage of the interaction of her two choice-variables (e_1 and $w_2(x)$); i.e. the individual at the top of the hierarchy not only proposes a contract to her subordinates but also decides her plans for the future (her effort decision).[9] It would be boundedly rational for her not to take into account the interaction between both decisions. Indeed, in the British railway example, we can expect that if the Office of Rail Regulation (Principal) contracts with, for example, Network Rail (Agent 1), then the latter will choose simultaneously its investment in track maintenance (e_1) and the access charges ($w_2(x)$) that the Train Operation Companies (Agent 2) have to pay in order to use the tracks.

Finally, it is important to note that although Agent 1 decides on her own effort before Agent 2 makes her decision, Agent 2 does not need to observe Agent 1's choice.[10] Though our model certainly works in such circumstances, it does not require them: it also works under the less demanding condition that agents are able to determine the equilibrium of the game, a common assumption in the literature.

Given the nature of both agents' activities our model seems to suggest that ORR should delegate contracting rights to NR, which should in turn contract with the TOC (note that our model and results naturally extend to the situation where there are many TOCs). The rationale for this suggestion is that, since the same piece of railtrack may be used by several TOCs, the alternative delegated structure (i.e., that in which TOCs are given joint control over it) will imply the multiplication of fixed costs and/or costly coordination between companies. These problems would be eliminated if control is given to a unique

[9] This is the reason why the Principal cannot replicate the Delegated Second Best outcome by simply offering the agents the DSB contracts. See appendix A.5 for a formal proof.

[10] As was noted above, this last scenario introduces the possibility (which is out of the scope of the present paper) of message games where Agent 2 could disclose the information she oberves to the principal.

maintenance company (NR), which will also be allowed to choose the fees that the TOCs should pay in order to use the track.

Finally, we remain silent on the question of the public or private ownership of the organisations involved. Rather, we have proposed a contractual framework that, while remaining within the regulatory framework created by EU directive 91/440, offers a way of improving passenger railway services.

3.5 Extensions

The number of extensions that could be pursued taking the present model as a starting point is large, so we concentrate on three of them, namely, the cases of costly contracting and of heterogeneity of agents due to differences in risk aversion and productivity. The first one is relevant because it can shed light on the relative frequencies of delegated and centralised structures in the real world; the latter two because they affect the allocation of agents to different tiers of a hierarchy.

3.5.1 Costly contracting

Writing a contract is an activity whose cost can be divided into fixed and variable components. The first one is usually larger than the second one, since legal fees and administrative costs are significantly higher than printing another copy of the contract. For all practical reasons, therefore, one can assume that variable costs are zero and so only fixed costs matter. Since the Principal always contracts with at least one agent, this means that her expected profits become

$$E\pi' := E\pi - \kappa \tag{3.5}$$

where $\kappa > 0$ is the contracting cost. Agent 1's objective function is unaffected when she does not have contracting rights (case CSB) but becomes

$$Eu_1' := Eu_1 - \kappa \tag{3.6}$$

when she does have them (case DSB). Agent 2 never subcontracts anyone, so her objective function is unaffected by the presence or absence of contracting costs.

Under the centralised structure (CSB), the costly-contracting solution is identical to the costless one, with the only difference that the Principal's expected profits decrease by the amount corresponding to the contracting cost κ. Under delegation (DSB), however, both the Principal and Agent 1 contract (with Agent 1 and Agent 2, respectively) and so the social cost of contracting is doubled. The Principal foresees that subcontracting will occur and, in order to satisfy Agent 1's participation constraint, she will have to compensate the latter for the contracting cost she will bear and expected profits will be reduced (when compared to the costless expected profits) by the *social* cost of contracting, 2κ. Figure 3.5 compares the Principal's expected profits in cases CSB and DSB when contracting is costly and shows the areas in the parameter space (γ, κ) where each case dominates the other (DSB dominates CSB below the curve and the opposite is true above the curve).

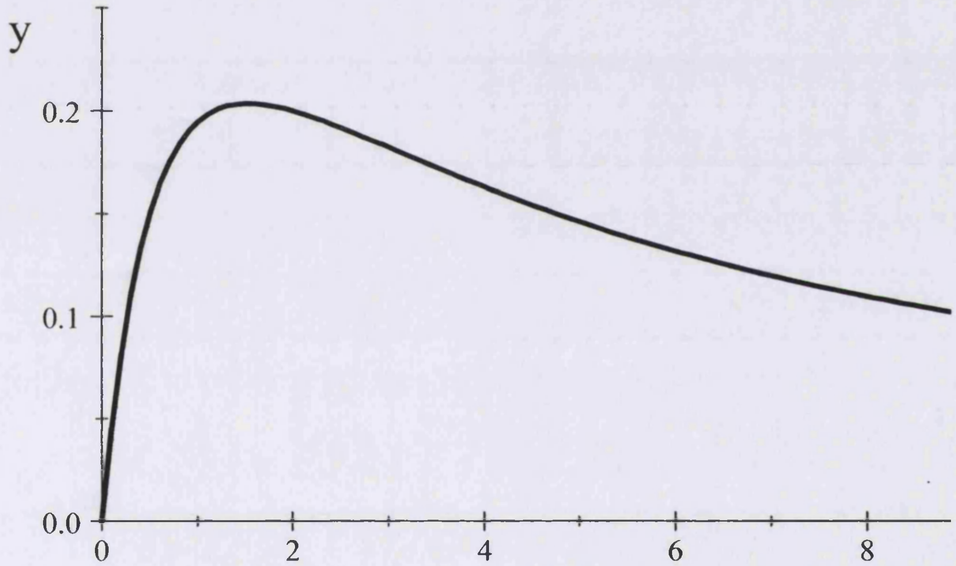

Figure 3.5: Costly contracting. CSB *v* DSB.
Horizontal. axis: γ. Vertical. axis: κ.

As expected, DSB dominates CSB when contracting costs are low (in particular, the costless contracting case of figure 3.1 corresponds to the horizontal axis in figure 3.5), but the dominance relationship is non-monotonic in γ: for low values of γ (when agents are only slightly risk averse) CSB dominates DSB because of the duplication of contracting costs; however, as γ increases, the gains from (partially) internalising the externality more than compensate the additional contracting cost and the inequality is reversed; finally, as γ gets sufficiently high, the marginal gain from internalising the externality decreases and the duplication of costs recovers its prominent role, so that CSB becomes dominant again.

This result, therefore, sheds light on the question of the frequency of centralised and delegated contracting structures in the real world. In particular, it can be seen that delegation will work when contracting costs are low and the agents are moderately risk averse (alternatively, the project is moderately risky). Also, because of the rapid escalation of contracting costs, it shows that multi-tiered hierarchies/subcontracting chains are unlikely to be observed in reality unless the externalities that are internalised are significant.

3.5.2 Heterogeneous agents

In the basic model of section 3.2, the assumption of homogeneous agents was chosen to avoid other factors that may have created incentives to treat them differently (in particular, to choose one agent over the other to play the role of Agent 1) and to highlight the beneficial effects of delegation even when there are no differences among agents. Indeed, it is only because of the efficiency gains from internalising the externality via delegation that the principal deviates from the (apparently obvious) policy of treating identical agents identically. And it is because the same efficiency gains that the pre-eminence of delegation over centralisation is maintained when agents are heterogeneous. But heterogeneity brings forward an important question regarding the contracting structure, namely, who should be contracted by the principal? Two cases are analysed here: in one of them agents differ in their degrees of risk aversion, in the other in their productivities.

3.5.2.1 Risk aversion

The heterogeneity in terms of risk aversion is reflected by the values that γ takes for different agents: from the definition of γ on page 22, the higher the degree of risk aversion r, the greater is γ. Without loss of generality one agent (call her A) is assumed more risk averse than the other (named B). Formally,

$$\gamma_A > \gamma_B > 0 \tag{3.7}$$

In a similar way expected profits can be labelled depending on the contracting structure: $E\pi^d_{AB}$ if the principal contract with agent A and $E\pi^d_{BA}$ if she contracts with agent B. It is straightforward to show that

$$E\pi_{AB}^{d} - E\pi_{BA}^{d} = -\frac{(\gamma_B - \gamma_A)\left((3\gamma_B + 2)\gamma_A^2 + (\gamma_B + 2)(\gamma_B + 1)(3\gamma_A + 2)\right)}{2(1 + \gamma_A + \gamma_B)(2 + \gamma_B)(1 + \gamma_A)(2 + \gamma_A)(1 + \gamma_B)} > 0 \quad (3.8)$$

which implies that the principal should contract with the most risk averse agent (A in this case).

This is consistent with the finding (with homogeneous agents) that Agent 1 faces a lower risk than Agent 2 (proposition 3.2). Intuitively, a very risk averse Agent 1 will transfer most of her risk to Agent 2 by choosing b very close to B, and this in turn implies that Agent 2's effort ($e_2 = b$) will be very close to Agent 1's effort ($e_1 = B$), thus increasing (expected) output and profits.

3.5.2.2 Productivity

The second source of heterogeneity to be explored is related to the differences in productivity among agents. These can be modelled by assuming that different agents have different marginal disutilities of effort $\frac{\partial\phi(e_i)}{\partial e_i}$, where $\phi(e_i) := \frac{1}{2}\phi_i e_i^2$, $\phi_i > 0$. Without loss of generality one can normalise the marginal disutility of Agent 1 to 1 ($\phi_1 = 1$) and of Agent 2 to ϕ ($\phi_2 = \phi$). The principal's expected profits then become

$$E\pi^d = \frac{1}{2}\frac{(\phi(1+\gamma) + 1 + 2\gamma)^2}{(1+2\gamma)\phi(1 + 2\gamma + \phi(1 + \gamma + \gamma^2))} \quad (3.9)$$

which is a decreasing function of ϕ, so that the principal should contract with the most productive agent (A if $\phi < 1$ and B if $\phi > 1$).

This result is in line with that of proposition 3.2 that states that Agent 1 exerts more effort than Agent 2. Intuitively, since Agent 1's disutility of effort is lower than that of Agent 2 (for a given level of effort), the principal needs to pay less to the former than to the latter in order to satisfy their participation constraints, thus increasing expected profits.

3.6 Conclusion

We analysed how team production should be organised when contracts can be written contingent only on joint output. We find that the positive externality associated with this setup can be partially internalised by the delegation of contracting rights, which decreases the misalignment between the principal's and the agents' interests and thus increases the overall efficiency of the organisation. The superiority of delegation over centralisation shows, therefore, that the mere addition of the externality may have a radical effect on the design of incentive schemes.

Regarding the British passenger rail services, the present study seems to suggest the convenience of allowing Network Rail both to choose how to allocate its investment in track maintenance and to determine the access charges that train companies should pay for using the network. However, these conclusions are the result of analysing a very simplified model of the industry and hence have to be tempered by the fact that many crucial features that would affect the outcome (e.g., price regulation) were left out in order to concentrate on the organisational problem.

Part II

Endogenous Externality

Chapter 4

Anti-Evasion Policy: Theory

4.1 Introduction

It is common practice for tax agencies worldwide to use observable characteristics of taxpayers to partition the population into fairly homogeneous categories in order to better estimate their incomes: all other things being equal, those who declare well below the estimate are likely to be evaders and are audited, while those who declare about or above it are likely to be compliant taxpayers and are not inspected. But this "cut-off" auditing policy (Reinganum and Wilde (1985)) can lead to systematic mistargeting in the presence of common shocks: in good years the category would be under-audited (bars and pubs in a heat-wave); in bad years it would be over-audited (chicken-breeders in an avian-flu outbreak).

The present chapter focuses on the problem a tax agency faces when deciding its auditing policy within each audit category in such scenario. To avoid systematic mistargeting, the government needs *contemporaneous* data correlated with the common shock. I examine the possibility of using the profile of declarations of the taxpayers in a category as a signal of the shock experienced by them and show that, for a government facing a low-income declarer, the optimal auditing strategy is (weakly) increasing in the other taxpayers' declarations. Intuitively, the higher these declarations, the more likely the shock was a positive one, and so the more likely that someone who declares low income is an evader. Precisely this type of reasoning is presumed to be behind the method used by the IRS's "Discriminant Index Function" (DIF) to determine which taxpayers to audit.[1]

[1]On page 301, Alm and McKee (2004) say: "(...) a taxpayer's probability of audit is based not only

This policy introduces a *negative externality* among taxpayers, one that *would not exist otherwise*: if someone increases her declaration, everyone else's probability of detection is increased. This changes the nature of the evasion problem by creating a coordination game among agents: each one of them has incentives to evade if most other people evade as well, and prefers to comply if most of the rest are compliant. The resulting multiplicity of equilibria and its associated policy design problems are avoided by the presence of an information asymmetry in favour of the tax agency. A government's innate "toughness" with respect to evasion is a parameter that is its private information, enters its objective function and affects its optimal policy: *ceteris paribus*, tougher agencies will audit more intensively than softer ones. Since this parameter is an agency's private information, taxpayers need to estimate it in order to decide how much income to declare and they do it based on the information available to them, namely, their incomes and their signals. Each taxpayer's previous experiences, conversations with friends and colleagues and interpretation of media news constitute noisy signals of the government's type and are taxpayers' private information. The heterogeneity of signals makes different taxpayers perceive their situations as different from other taxpayers', and yet every one of them follows the same income declaration strategy. This leads to the survival of only one equilibrium in which (usually) some people evade and others comply, a result that is empirically supported and yet unlikely to be predicted by other tax evasion models.

Previous research on the area (started by Allingham and Sandmo (1972) and surveyed by Cowell (1990) and Andreoni et al. (1998)) did analyse the effect of asymmetric information in the tax compliance game. Some only considered the presence of "strategic uncertainty" (i.e., the uncertainty that taxpayers face in coordination games about which equilibrium will be selected), usually generated by psychological and/or social externalities (Benjamini and Maital (1985), Fortin et al. (2004), etc.). Others restricted their attention to the "fundamental uncertainty" faced by the taxpayers with respect to the type of agency (Scotchmer and Slemrod (1989), Stella (1991), etc.). The present study, on the other hand, considers both types of uncertainty and thus models the situation as a global game (Carlsson and van Damme (1993), Morris and Shin (2002b)).

The closest references to the present article are Alm and McKee (2004), Basseto and Phelan (2004) and Kim (2005). The first one is a laboratory experiment where the (*ad hoc*) auditing policy is contingent on the distribution of income declarations, while the second and third ones use the global game technique to determine the optimal tax system and the auditing policy, respectively. This paper presents a theoretical analysis in which –unlike the laboratory experiment– the agency's optimal strategy is derived instead of

upon his or her reporting choices, but also upon these choices relative to other taxpayers in the cohort. In short, there is a taxpayer-taxpayer game that determines each individual's chances of audit selection."

assumed. The other two studies employ the same technique that I use here, but while Basseto and Phelan (2004) are concerned with the optimal tax system as designed by a government, this article focuses only on the targeting aspect of *one* of the agencies of the government. Finally, Kim (2005) generates the strategic interaction among taxpayers by adding a "stigma cost" to their utility functions, whereas in my case it is the result of a cunning tax agency that sets its auditing policy to maximise its objective function.

4.2 Model

The model focuses on the interaction between the tax agency (also referred to as "the government") and the taxpayers (or "agents") *within a given category.* For simplicity, I will use "population of taxpayers" and "common shocks" to indicate the members of the category and the shocks faced by them, and *not* those of the whole population (i.e., the set which is the union of all the categories), unless indicated otherwise.

4.2.1 Timing

The timing of the game is as follows:

1. Actors (tax agency and taxpayers) receive their pieces of private information (the agency its "type" λ, the taxpayers their incomes y and signals s).
2. Taxpayers submit their income declarations d and pay taxes accordingly.
3. Finally, the agency observes the vector of declarations $\mathbf{d}$ and undertakes audits and collects fines (if any).

4.2.2 Objective functions

TAXPAYERS

Taxpayers are uniformly distributed on the $[0,1]$ segment and are assumed to be risk-neutral, so that their utility is a linear function of their disposable income:

$$u(d_i, a_i, y_i) = y_i - td_i - a_i \cdot \Phi(d_i, y_i) \qquad \forall i \in [0,1] \tag{4.1}$$

where $y_i \in \{0,1\}$ is agent i's gross (taxable) income, $t \in (0,1)$ is the income tax rate,[2] $d_i \in \{0,1\}$ is agent i's income declaration, $a_i \in \{0,1\}$ is an indicator function defined as

$$a_i = \begin{cases} 1 & \text{if agent } i \text{ is audited} \\ 0 & \text{if agent } i \text{ is not audited} \end{cases} \tag{4.4}$$

and $\Phi(d_i, y_i)$ is the fine agent i should pay if audited, defined as

$$\Phi(d_i, y_i) = \begin{cases} f \cdot (y_i - d_i) & \text{if } d_i < y_i \\ 0 & \text{otherwise} \end{cases} \tag{4.5}$$

where $f := (1+\varsigma)\, t$ and $\varsigma \in (0,1)$ is the surcharge rate that has to be paid by a caught evader on every dollar of evaded taxes).[3]

TAX AGENCY

Narrowly defined, a tax agency's objective is to raise revenue. More generally, its problem consists of determining which citizens should be audited and which ones should not.

An agency, therefore, chooses its auditing strategy in order to minimise its targeting errors.[4]

These errors can be of two types: Negligence and Zeal. A negligence mistake occurs when a "profitable audit" is not undertaken. A zeal error takes place when an "unprofitable audit" is carried out.

[2]The tax system can be easily transformed into a progressive one by using the following change of variables

$$y := Y - B \tag{4.2}$$

where $Y \in \{B, B+1\}$ is a taxpayer's gross income, B is the exemption level and y is taxable income. The simplest progressive tax system is therefore

$$T = \begin{cases} t(D-B) & \text{if } \; D \geq B \\ 0 & \text{if } \; D < B \end{cases} \tag{4.3}$$

where $D \in \{B, B+1\}$ is the taxpayer's declaration. Thus, in bad years everyone in the class is exempt and in good ones everyone is liable to pay taxes.

[3]The IRS applies rates between 20% (misconduct) and 75% (fraud) (Andreoni et al. (1998)), so $\varsigma \in (0,1)$ covers the relevant range. It is assumed that $(1+\varsigma)\, t < 1$, such that the fine if caught evading does not exhaust a high-income person's income.

[4]The analysis also holds if the the agency's objective function is based on expected net revenue. See appendix B.1.

An audit is defined as "profitable" if the fine obtained if undertaken more than compensates for the cost of carrying it out (formally, if $\Phi(d_i, y_i) > c$, where $\Phi(d_i, y_i)$ is the fine –as defined in equation 4.5– and $c \in (\varsigma t, (1+\varsigma)t)$ is the cost of the audit). It is assumed that an audit that discovers an evader is always profitable, while an audit that targets a compliant taxpayer is always unprofitable. Formally, if $\alpha_i = 1$ means that auditing agent i is profitable, then

$$\alpha_i := \begin{cases} 1 & \text{if } y_i = 1 \text{ and } d_i = 0 \\ 0 & \text{otherwise} \end{cases} \tag{4.6}$$

Hence, a negligence error (N_i) occurs when the audit is profitable ($\alpha_i = 1$) and it is not undertaken ($a_i = 0$). On the other hand, a zeal error (Z_i) occurs when the audit is not profitable ($\alpha_i = 0$) and yet it is undertaken ($a_i = 1$). Formally,

$$N_i := \begin{cases} 1 & \text{if } \alpha_i = 1 \text{ and } a_i = 0 \\ 0 & \text{otherwise} \end{cases} \qquad Z_i := \begin{cases} 1 & \text{if } \alpha_i = 0 \text{ and } a_i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{4.7}$$

For the rest of the article, and due to the fact that they make the problem more tractable, I will use –without loss of generality– the following two error functions:

$$N_i := (1-a_i)(1-d_i)\,y_i \qquad Z_i := a_i\left[1-(1-d_i)\,y_i\right] \tag{4.8}$$

Different agencies can, however, value each kind of error differently. If λ is defined as the weight attached to negligence errors, the loss inflicted by agent i on an agency of type λ can be expressed as

$$L_i := \lambda N_i + (1-\lambda)\,Z_i \tag{4.9}$$

Aggregating over all taxpayers, the loss function becomes

$$L(\mathbf{a}, \mathbf{d}, y) := \int_0^1 \left[\lambda N_i + (1-\lambda)\,Z_i\right] di \tag{4.10}$$

since i (that indexes taxpayer i) is uniformly distributed on the $[0,1]$ segment.

4.2.3 Strategy spaces

A taxpayer's strategy consists of choosing an income declaration $d \in \{0,1\}$. A tax agency's strategy consists of choosing a vector of auditing decisions $\mathbf{a}$, such that its ith argument,

$a_i \in \{0, 1\}$, indicates the auditing decision regarding taxpayer i.

4.2.4 Information sets

At the node where an actor A makes her decision, her information set, I_A, consists of the union of two sets: one that is common to all actors, I^c, and one that includes the actor's private information, I^p_A. Formally,

$$I_A = I^c \cup I^p_A \tag{4.11}$$

where $A \in \{TA, i\}$ stands for *actor*, TA for *tax agency* and i for *taxpayer* i, and superindices c and p identify the *common* and *private* sets, respectively.

COMMON SET I^c

The common set I^c includes the exogeneous parameters of the problem (like the tax rate t and the surcharge rate ς) and the parameters of the probability distributions of the private information variables (income y_i, type of agency λ and signal s_i).

Incomes are assumed perfectly correlated to reflect the fact that common shocks affect similar agents in similar ways:[5]

$$y_i = y \ \forall i \in [0, 1] \tag{4.12}$$

"Good years" ($y = 1$) occur with probability $\gamma \in (0, 1)$ and "bad years" with probability $1 - \gamma$.[6]

The agency's type λ is a non-manipulable characteristic of the agency that affects the government's auditing policy. It is uniformly distributed on the $[\varepsilon, 1 - \varepsilon]$ segment $\left(0 < \varepsilon < \frac{1}{2}\right)$ and is independent of the income shock.

Taxpayers' signals s_i convey information about the government's type λ and are, on average, correct. They reflect the information about the agency's type that taxpayers get from all available sources: media news, previous experiences, conversations with colleagues and friends, etc. Formally,

$$s_i := \lambda + \varepsilon_i \qquad \forall i \in [0, 1] \tag{4.13}$$

[5]The imperfect correlation case is analysed in appendix B.2.

[6]The case in which income can take more than two values is considered in appendix B.3.

where ε_i is the error term, which is assumed to be white noise ($E(\varepsilon_i) = 0\ \forall i$), uniformly distributed on the $[-\varepsilon, \varepsilon]$ segment, and independent of income y_i, other taxpayers' errors $\varepsilon_{j\neq i}$ and the government's type λ.

TAXPAYER i'S PRIVATE SET I_i^p

Each taxpayer knows the realisation of her private information variables, namely, her income y_i and her private signal s_i. Furthermore, since all taxpayers know that the income distribution is degenerate, they know that every taxpayer has the same income y ($y = 0$ if $y_i = 0$ and $y = 1$ if $y_i = 1$).

TAX AGENCY'S PRIVATE SET I_{TA}^p

The agency knows the realisation of her private information variable, its type λ. Also, given the timing of the game, it observes the vector of income declarations $\mathbf{d}$, each argument $d_i \in \{0, 1\}$ being the declaration of a taxpayer. Given the dichotomous nature of the declarations, the vector of income declarations can be summarised by a sufficient statistics, namely, the average declaration $D \in [0, 1]$, which will be used henceforth.

4.2.5 Schematic representation of the game

Given the elements presented so far, the game can be represented as in tables 4.1 (for bad years) and 4.2 (for good years).

$y = 0$ $(1-\gamma)$		Agency's strategy: Audit $(a = 1)$	Agency's strategy: Do not audit $(a = 0)$
Taxpayer's strategy	Declare low $(d = 0)$	$1 - \lambda$ 0	0 0
	Declare high $(d = 1)$	$1 - \lambda$ $-t$	0 $-t$

Note: In each cell, Bottom-left element: taxpayer's utility; Top-right element: agency's loss.

Table 4.1: Compliance game in bad years (y=0).

$y=1$ (γ)		Agency's strategy	
		Audit ($a=1$)	Do not audit ($a=0$)
Taxpayer's strategy	Declare low ($d=0$)	0 $1-f$	λ 1
	Declare high ($d=1$)	$1-\lambda$ $1-t$	0 $1-t$

Note: In each cell, Bottom-left element: taxpayer's utility; Top-right element: agency's loss.

Table 4.2: Compliance game in good years (y=1).

Taxpayers observe their income before making their declarations, so that they know which of the two games is being played. The tax agency, on the other hand, does not know the true value of y, and therefore does not know which of the two games is being played.

4.2.6 Equilibrium concept

A perfect Bayesian equilibrium must specify actors' posterior beliefs, taxpayers' income declaration strategies, the agency's auditing strategy and the average declaration in the category. Formally,

$$\lambda \mid s \sim U[s-\varepsilon, s+\varepsilon] \tag{4.14}$$

$$s \mid \lambda \sim U[\lambda-\varepsilon, \lambda+\varepsilon] \tag{4.15}$$

$$\Pr(y=1 \mid D) = \begin{cases} 1 & \text{if } D>0 \\ \gamma & \text{if } D=0 \end{cases} \tag{4.16}$$

$$d_i^*(s,y) \in \underset{d_i\in\{0,1\}}{\arg\max} E\{u(d_i,a_i,y_i) \mid I_i\} \;\; \forall i \in [0,1] \tag{4.17}$$

$$\mathbf{a}^*(D,\lambda) \in \underset{a_i\in\{0,1\}\;\forall i\in[0,1]}{\arg\min} E\{L(\mathbf{a},\mathbf{d},y) \mid I_{TA}\} \tag{4.18}$$

$$D(\mathbf{d},\lambda) = \int_0^1 d_i^*(s,y)\, di \tag{4.19}$$

The first three lines indicate that all actors have Bayesian beliefs. In particular, taxpayers know that the posterior distribution of the type of the agency λ (conditional on the taxpayer's signal s) is uniform with support $[s-\varepsilon, s+\varepsilon]$ (equation 4.14). The tax agency, on the other hand, knows that the posterior distribution of signals (conditional on the agency's type λ) is uniform with support $[\lambda-\varepsilon, \lambda+\varepsilon]$ (equation 4.15). The agency also knows that the posterior distribution of income (conditional on the observed average declaration D) is such that equation 4.16 holds.

The following equations indicate that actors choose their actions optimally: taxpayers choose their declarations d in order to maximise their expected utility, conditional on the available information I_i (equation 4.17) and the tax agency chooses its auditing strategy $\mathbf{a}$ so as to minimise the expected losses due to targeting mistakes, conditional on its available information I_{TA} (equation 4.18). Finally, equation 4.19 aggregates the taxpayers' decisions to give the average declaration.

4.3 Solving the model

4.3.1 Preliminaries

In bad years ($y = 0$) taxpayers know that the game being played is the one depicted in table 4.1. The taxpayer's optimal strategy is therefore:

Proposition 4.1 *In bad years, every taxpayer declares low income. Formally, for all* $i \in [0,1]$

$$d_i^* (s, 0) = 0 \tag{4.20}$$

Proof. From direct inspection of table 4.1. Declaring low income strictly dominates declaring high income: the payoff is 0 in the first case and $-t$ in the second one, irrespective of the other taxpayers' declarations and the agency's auditing decision. ■

It is therefore common knowledge that the game shown in table 4.1 simplifies to

$y = 0$ $(1 - \gamma)$		Agency's strategy	
		Audit $(a = 1)$	Do not audit $(a = 0)$
Taxpayer's strategy	Declare low ($d = 0$)	$1 - \lambda$ 0	0 0

Note: In each cell, Bottom-left element: taxpayer's utility; Top-right element: agency's loss.

Table 4.3: Compliance game in bad years (y=0).

Thus, the compliance game consists of the ones shown in tables 4.2 (for good years) and 4.3 (for bad ones). Given the timing presented in section 4.2.1, the game is solved by backwards induction. Hence, I will analyse first the second stage and will solve the tax agency poblem

4.3.2 Tax agency problem

The agency's problem consists of choosing the audit vector $\mathbf{a}$ so as to minimise the expected losses, conditional on its information set I_{TA}. The expected loss function is therefore:

$$E\{L(\mathbf{a}, D, y) \mid I_{TA}\} = \Pr(y = 0 \mid I_{TA}) \cdot L(\mathbf{a}, D, 0) + \Pr(y = 1 \mid I_{TA}) \cdot L(\mathbf{a}, D, 1) \quad (4.21)$$

where $L(\mathbf{a}, D, 0)$ is the loss when income is low ($y = 0$) and $L(\mathbf{a}, D, 1)$ is the loss when income is high ($y = 1$). Define the probability of high income conditional on observing average declaration D as

$$\phi(D) := \Pr(y = 1 \mid I_{TA}) \quad (4.22)$$

Note, also, that D represents the proportion of the population that declares high income. Hence, without loss of generality, we can assume that taxpayers on the $[0, 1 - D]$ segment declare low income and those on the $(1 - D, 1]$ segment declare high income. The loss function (equation 4.10) can therefore be re-written as

$$L(\mathbf{a}, D, y) = \lambda y \int_0^{1-D} (1 - a_i)\, di + (1 - \lambda)(1 - y) \int_0^{1-D} a_i di + (1 - \lambda) \int_{1-D}^{1} a_i di \quad (4.23)$$

where the first two terms correspond to the expected loss generated by those who declare low income and the last one corresponds to the loss generated by those who declare high income, and where the assumption of perfectly correlated incomes ($y_i = y \; \forall i \in [0, 1]$) allows us to take y out of the integral.

Furthermore, from the perspective of the agency, taxpayers who declare low income are indistinguishable from each other, and the same is true for those who declare high income. This means that the government will treat every person in each group in an identical way. This implies, therefore, that the agency has only two policy variables: a_0 and a_1, which are the audit decisions for taxpayers who declare low and high income, respectively. The loss function then becomes

$$L(\mathbf{a}, D, y) = \lambda y (1 - a_0)(1 - D) + (1 - \lambda)(1 - y) a_0 (1 - D) + (1 - \lambda) a_1 D \quad (4.24)$$

which takes the value

$$L(\mathbf{a}, D, 0) = (1-\lambda)\, a_0 (1-D) + (1-\lambda)\, a_1 D \tag{4.25}$$

when the income shock is negative ($y = 0$), and the value

$$L(\mathbf{a}, D, 1) = \lambda (1-a_0)(1-D) + (1-\lambda)\, a_1 D \tag{4.26}$$

Thus, from equations 4.22, 4.25 and 4.26, the expected loss function of equation 4.21 becomes

$$E_{TA}(L) = \{[1-\phi(D)](1-\lambda)\, a_0 + \phi(D)\, \lambda (1-a_0)\}(1-D) + (1-\lambda)\, a_1 D \tag{4.27}$$

where the subindex TA indicates that the expectation is conditional on the information set of the tax agency and the arguments of the loss function were omitted for simplicity.

The agency therefore minimises the expected loss, as indicated in equation 4.27. The results are summarised in the following proposition.

Proposition 4.2 *For every taxpayer, a λ-type agency's optimal auditing strategy is as follows:*

- *if a taxpayer declares high income ($d_i = 1$), do not audit her ($a_1 = 0$);*
- *if every taxpayer declares low income ($D = 0$) and the agency is "soft" (λ sufficiently low), do not audit anyone ($a_0 = 0$);*
- *if every taxpayer declares low income ($D = 0$) and the agency is "tough" (λ sufficiently high), audit everyone ($a_0 = 1$);*
- *if some taxpayers declare low income and others declare high income ($D > 0$), audit everyone who declares low income ($a_0 = 1$).*

Formally, for every taxpayer $i \in [0,1]$,

$$a_i^*(d_i, D, \lambda) = \begin{cases} 0 & \text{if } d_i = 1 \\ 0 & \text{if } d_i = 0,\ D = 0, \text{ and } \lambda < \tilde{\lambda} \\ \in [0,1] & \text{if } d_i = 0,\ D = 0, \text{ and } \lambda = \tilde{\lambda} \\ 1 & \text{if } d_i = 0,\ D = 0, \text{ and } \lambda > \tilde{\lambda} \\ 1 & \text{if } d_i = 0, \text{ and } D > 0 \end{cases} \tag{4.28}$$

where $\tilde{\lambda} := 1 - \gamma$ *and* $\gamma \in (0, 1)$ *is the probability of a good year.*

Proof. In appendix B.4, page 122. ■

Intuitively, the proposition says that an agency's optimal auditing decision with respect to a given taxpayer i depends on the taxpayer's decision d_i, the declarations of all other taxpayers (summarised by the average declaration D) and the agency's type λ. When at least one person declares high income (and so $D > 0$), the government knows for sure –thanks to the perfect correlation assumption– that the shock was a positive one (it was a "good year"), and so the optimal strategy consists of auditing everyone who declares low income ($a_i^*(0, D > 0, \lambda) = 1$, since they are evaders) and not auditing anyone who declares high income ($a_i^*(1, D, \lambda) = 0$, since only "rich" taxpayers ever declare high income, and so their declarations are truthful). When everyone declares low income (so $D = 0$), the government cannot tell whether it faces a population of "poor" compliant taxpayers or one of "rich" evaders. The optimal policy therefore depends on how tough the government is (i.e., how high λ is) and how likely it is for the taxpayers to face a good year (i.e., the value of γ). If the agency is rather tough (λ is rather high), the optimal policy consists of auditing everyone (and the same is true if the probability of a good year, γ, is high). Otherwise (if the agency is rather soft or a bad year is very likely), it is better for the agency to audit no one.

These results are summarised in the following proposition:

Proposition 4.3 *For every taxpayer, a λ-type agency's optimal auditing strategy is:* (1) *(weakly) increasing in the agency's type* λ, *and* (2) *(weakly) increasing in the probability of a good year* γ. *Formally,*

$$(1) \quad \frac{\partial a_i^*(d_i, D, \lambda)}{\partial \lambda} \geqslant 0 \qquad\qquad (2) \quad \frac{\partial a_i^*(d_i, D, \lambda)}{\partial \gamma} \geqslant 0 \qquad\qquad (4.29)$$

Proof. By direct inspection of equation 4.28. ■

Further characterizing the agency's optimal strategy, the next result describes how it depends on the taxpayer's own declaration as well as on every other taxpayer's declarations:

Proposition 4.4 *For every taxpayer, a λ-type agency's optimal auditing strategy is:* (1) *(weakly) increasing in every other taxpayers' declaration* $d_{j \neq i}$, *and* (2) *(weakly) decreasing in the taxpayer's own declaration* d_i. *Formally,*

$$(1) \quad \frac{\partial a_i^*(d_i, D, \lambda)}{\partial d_{j \neq i}} \geqslant 0 \qquad\qquad (2) \quad \frac{\partial a_i^*(d_i, D, \lambda)}{\partial d_i} \leqslant 0 \qquad\qquad (4.30)$$

Proof. By direct inspection of equation 4.28. ∎

Intuitively, this means that the agency audits individuals who declare high income with a lower probability than those who declare low income (as is standard in tax evasion models). The novelty of the present study is in the result of equation 4.30.1, which shows that a loss-minimising agency would use the information conveyed by the vector of income declarations (or the average declaration, which in this case is a sufficient statistics) when deciding its optimal policy. In particular, the declarations of other taxpayers provide *contemporaneous* information about the likelihood of a given income shock, improving the targeting proficiency of the agency that can thus perfectly distinguish between truthful and untruthful declarations when the average declaration is different from 0.

4.3.3 Taxpayer problem

Once the second stage game is solved, we can turn to the first stage and solve the taxpayer problem.

As shown in section 4.3.1, in bad years all taxpayers declare low income. Hence, here I will focus on the case when income is high.

In good years, each taxpayer i chooses her income declaration d_i so as to maximise her expected utility, conditional on her information set I_i. Her expected utility function is:

$$E\{u(d_i, a_i, y = 1) \mid I_i\} = \Pr(a_i = 1 \mid I_i) \cdot u(d_i, 1, 1) + \Pr(a_i = 0 \mid I_i) \cdot u(d_i, 0, 1) \quad (4.31)$$

Noting that the expected value of the audit decision simplifies to the probability of an audit:

$$E_i(a_i) = \Pr(a_i = 1 \mid I_i) \quad (4.32)$$

and using equation 4.1, the expected utility function of equation 4.31 becomes

$$E_i(u) = 1 - td_i - f_i \cdot E_i(a_i) \quad (4.33)$$

where the subindex i indicates that the expectation is conditional on the information set of taxpayer i and the arguments of the utility function were omitted for simplicity.

If the taxpayer evades $d_i = 0$, her expected utility equals gross income minus the expected fine:

$$E_i(u(evasion)) = 1 - f \cdot E_i(a_i(d_i = 0)) \quad (4.34)$$

If the taxpayer complies, she gets utility

$$u(compliance) = 1 - t \tag{4.35}$$

with certainty.

The taxpayer's optimal decision $d^*(y_i, E_i(a_i(d_i = 0)))$ depends on the comparison between the two as follows

$$d^*(1, E_i(a_i(d_i = 0))) = \begin{cases} 0 & \text{if } E_i(a_i(d_i = 0)) < P \\ \in [0,1] & \text{if } E_i(a_i(d_i = 0)) = P \\ 1 & \text{if } E_i(a_i(d_i = 0)) > P \end{cases} \tag{4.36}$$

where $P := \frac{1}{(1+\varsigma)}$ is the probability of detection that eliminates evasion.

Intuitively, in good years taxpayers evade only if their subjective belief about the probability of being audited is not too high. This implies that an agent's declaration is (weakly) increasing in her expectation over the probability of detection.

Combining the results for bad and good years (proposition 4.1 and equation 4.36), the solution to the taxpayer problem is

$$d^*(y_i, E_i(a_i(d_i = 0))) = \begin{cases} 0 & \text{if } y_i = 0 \\ 0 & \text{if } y_i = 1 \text{ and } E_i(a_i(d_i = 0)) < P \\ \in [0,1] & \text{if } y_i = 1 \text{ and } E_i(a_i(d_i = 0)) = P \\ 1 & \text{if } y_i = 1 \text{ and } E_i(a_i(d_i = 0)) > P \end{cases} \tag{4.37}$$

from which it is clear that an agent's declaration is (weakly) increasing in her gross income.

The latter results are summarised in the following proposition:

Proposition 4.5 *A taxpayer's optimal declaration strategy is:* (1) *(weakly) increasing in her (subjective) expectation over the probability of detection* $E_i(a_i(d_i = 0))$, *and* (2) *(weakly) increasing in her gross income* y_i. *Formally,*

$$(1) \quad \frac{\partial d^*(y_i, E_i(a_i(d_i=0)))}{\partial E_i(a_i(d_i=0))} \geqslant 0 \qquad (2) \quad \frac{\partial d^*(y_i, E_i(a_i(d_i=0)))}{\partial y_i} \geqslant 0 \tag{4.38}$$

Proof. By direct inspection of equation 4.37. ■

Equation 4.30.1 and the first part of proposition 4.38 make taxpayer i's optimal declaration strategy a (weakly) increasing function of the other taxpayers' declarations:

Proposition 4.6 *Taxpayers' declarations are (weakly) strategic complements. Formally, for every $j \neq i$,*

$$\frac{\partial d_i^*(y_i, s_i)}{\partial d_{j\neq i}} \geqslant 0 \tag{4.39}$$

Proof. Directly from propositions 4.5 and 4.4. ■

This proposition opens a channel through which a higher signal leads to a higher declaration: a high signal means that other taxpayers are also likely to receive high signals –and to declare high income too– which increases the expected probability of detection and makes compliance relatively more attractive (i.e., provides incentives to (weakly) increase the amount of income declared).

Even more importantly, this result transforms the nature of the tax evasion problem, because it creates a *coordination game* among the taxpayers *on top of* the cat-and-mouse game that each one of them plays against the agency and that is usually the only one considered by the literature. The strategic complementarity between taxpayers' declarations, however, is not an inherent characteristic of the game, but rather one that is created by the agency in its attempt to minimise its targeting errors. Indeed, it is the fact that the auditing strategy is an increasing function of other taxpayers' declarations (Proposition 4.4) that *creates a negative externality* between taxpayers (proposition 4.6). That is, a cunning agency, willing to minimise its targeting-related losses, designs its optimal auditing strategy by introducing some strategic uncertainty (i.e., by creating a coordination game between taxpayers) that improves its ability to distinguish compliant from non-compliant agents and thus decreases the occurrence of targeting mistakes.

The taxpayer's optimal declaration strategy can be further characterised in terms of private signals, as shown in the next proposition:

Proposition 4.7 *In good years, a taxpayer's optimal declaration strategy:* (1) *is the same for all taxpayers, and* (2) *is (weakly) increasing in her private signal s_i. Formally,*

$$(1) \quad d^*(1, s_i) = \begin{cases} 0 & \text{if } s_i < \hat{s} \\ \in [0,1] & \text{if } s_i = \hat{s} \\ 1 & \text{if } s_i > \hat{s} \end{cases} \qquad (2) \quad \frac{\partial d^*(1,s_i)}{\partial s_i} \geqslant 0 \tag{4.40}$$

where $\hat{s} := \tilde{\lambda} + \varepsilon(2P-1)$, $\tilde{\lambda} := 1-\gamma$ *and* $P := \frac{1}{(1+\varsigma)}$.

Proof. The first part is the result of $\hat{s}$ being a constant that is independent of the identity of the taxpayer whose strategy is being studied. The determination of $\hat{s}$ is shown in appendix B.4, page 123. For the second part, by direct inspection of equation 4.40.1. ■

The intuition is straightforward: the higher the signal received ($s_i := \lambda + \varepsilon_i$ from equation 4.13), the higher is the taxpayer's (subjective) expectation over the government's type λ, meaning that the agent believes that, very likely, she faces a tough agency and, consequently, a high probability of detection. This decreases the (subjective) expected return of evasion and makes compliance more attractive, which leads the taxpayer to (weakly) increase her income declaration.

The first part of the proposition highlights the fact that, though having different private signals, all taxpayers agree on the "switching point" below which one should evade and above which one should comply. Note also that, as expected, each "type" of taxpayer (defining agent i's "type" as its private information pair (y_i, s_i)) has a unique optimal strategy: taxpayers with low income ($y_i = 0$) ignore their signals and always declare low income; taxpayers with high income ($y_i = 1$) do take into account the signals they receive and declare income as shown in equation 4.40.1.

4.4 Equilibrium

A priori, the generation of a coordination game among taxpayers does not look as a good idea for the agency because this kind of games present multiple equilibria, which make policy design a complicated matter. Nevertheless, this difficulty is overcome thanks to the presence of a second source of uncertainty (called "fundamental uncertainty") that allows for the tax evasion problem to be modelled as a "global game" (Carlsson and van Damme (1993), Morris and Shin (2002b)).[7]

This equilibrium-selection technique eliminates all but one equilibria owing to the introduction of some heterogeneity in taxpayers' information sets in the form of the noisy private signals they receive and that convey information about the government's private information parameter λ (the source of the "fundamental uncertainty"). Thus, taxpayers do not observe the true coordination game (as they would do if signals were 100% accurate), but slightly different versions of it. This is the case since taxpayers with different

[7]In other applications (bank runs, currency crises, etc), this technique has been criticised because of not taking into account the coordinating power of markets and prices (Atkeson (2000)). This criticism is greatly mitigated in the case of tax evasion, since there is no "insurance market against an audit" to aggregate information about the government's type (the "fundamental", in global games jargon).

signals would work out different estimates of the agency's type λ and the average declaration D, and so of their probabilities of detection. The optimal declaration strategy, however, is one and the same for every "type" of taxpayer (propositions 4.1 and 4.7). The rationale for this result goes along the lines described in the paragraph immediately after the proof of proposition 4.6: my own signal gives me information about the possible signals that other taxpayers may have received and, more importantly, about the signals that they *cannot* have received, thus allowing me to discard some strategies that they cannot have followed. The application of this process iteratively by *every* taxpayer leads to the elimination of all strictly dominated strategies and leaves only one optimal strategy to be followed by every taxpayer (Morris and Shin (2002a)), namely, the ones in propositions 4.1 and 4.7.

As a consequence, once the private information variables (the agency's type λ and taxpayers' incomes and signals $(\mathbf{y}, \mathbf{s})$) are realised, the equilibrium will be *unique*.

However, depending on the value of λ, the equilibrium can present different features, as illustrated by the following proposition:[8]

Proposition 4.8 *In bad years ($y_i = 0\ \forall i \in [0,1]$), the average declaration is zero ($D = 0$), as is the level of evasion ($\kappa^* = 0$). In good years ($y_i = 1\ \forall i \in [0,1]$), the corresponding values are as follows:*

	Full evasion ($\lambda < \hat{s} - \varepsilon$)	*Partial evasion* ($\hat{s} - \varepsilon < \lambda < \hat{s} + \varepsilon$)	*Full compliance* ($\hat{s} + \varepsilon < \lambda$)
Average declaration D	0	$\frac{\lambda+\varepsilon-\hat{s}}{2\varepsilon}$	1
Level of evasion κ^*	1	$1 - \frac{\lambda+\varepsilon-\hat{s}}{2\varepsilon}$	0

(4.41)

Proof. In appendix B.4, page 127. ■

This shows that, as expected, evasion is lower the tougher the government is.

Equilibrium strategies for each actor are shown in the following proposition:

Proposition 4.9 *The unique equilibrium of the tax evasion game looks like one of the following cases:* (1) Full evasion *($\lambda < \hat{s} - \varepsilon$): in good years, every taxpayer evades and nobody is audited,* (2) Partial evasion *($\hat{s} - \varepsilon < \lambda < \hat{s} + \varepsilon$): in good years, taxpayers with*

[8]Since in bad years taxpayers declare low income in every scenario, the three cases are characterised (and labelled) according to the actions taken by taxpayers in good years.

low signals ($s_i < \hat{s}$) evade and are audited with certainty while those with high signals ($s_i > \hat{s}$) comply and are not audited, and (3) Full compliance *($\hat{s} + \varepsilon < \lambda$): in good years, every taxpayer complies and everyone who declares low income is audited. In bad years, taxpayers always declare truthfully. Formally,*

	Full evasion	*Partial evasion*	*Full compliance*
$d^*(0, s_i)$	0	0	0
$d^*(1, s_i)$	0	$\begin{cases} 0 & \textit{if } s_i < \hat{s} \\ \in [0,1] & \textit{if } s_i = \hat{s} \\ 1 & \textit{if } s_i > \hat{s} \end{cases}$	1
$a_i^*(d_i, D, \lambda)$	0	$\begin{cases} 0 & \textit{if } d_i = 1 \\ 1 & \textit{if } d_i = 0 \end{cases}$	$\begin{cases} 0 & \textit{if } d_i = 1 \\ 1 & \textit{if } d_i = 0 \end{cases}$

(4.42)

Proof. Follows directly from the optimal strategies of the players (propositions 4.1 and 4.7 for the taxpayers, propostion 4.2 for the agency) and the characterisation of the equilibrium in terms of the average declaration (proposition 4.8). ■

The full evasion case occurs when the agency is so soft ($\lambda < \hat{s} - \varepsilon$) that all taxpayers know it will audit nobody who declares low income, and so everyone evades. The opposite occurs in the full compliance case, in which the agency is so tough ($\hat{s} + \varepsilon < \lambda$) that all taxpayers know it will audit everyone who declares low income, and so everyone complies. The partial evasion case occurs when the government is not too soft nor too tough ($\hat{s} - \varepsilon < \lambda < \hat{s} + \varepsilon$) and so taxpayers would like to do as most taxpayers do (*strategic complementarity*). They follow the optimal strategy described in proposition 4.7, which means that the vector of declarations will be different from zero. The agency, observing this, would know for sure that true income is high and so will audit everyone who declares 0 and nobody that declares 1.

Building on these results, one can further characterise the three cases:

Proposition 4.10 *The payoffs of the players in the three possible scenarios are as fol-*

lows:

<table>
<tr><th></th><th>Full evasion</th><th>Partial evasion</th><th>Full compliance</th></tr>
<tr><td>Taxpayer/ Bad year</td><td>0</td><td>0</td><td>0</td></tr>
<tr><td>Taxpayer/ Good year</td><td>1</td><td>$\begin{cases} 1-t & \textit{if } d_i = 1 \\ 1-(1+\varsigma)t & \textit{if } d_i = 0 \end{cases}$</td><td>$1-t$</td></tr>
<tr><td>Tax Agency</td><td>$\gamma\lambda$</td><td>$\begin{cases} 0 & \textit{if } \lambda < \tilde{\lambda} \\ (1-\gamma)(1-\lambda) & \textit{if } \lambda > \tilde{\lambda} \end{cases}$</td><td>$(1-\gamma)(1-\lambda)$</td></tr>
</table>

(4.43)

Proof. Follows directly from the definition of the payoff functions of the players (equations 4.1 and 4.10), their optimal strategies (propositions 4.1, 4.2 and 4.7) and the characterisation of the equilibrium in terms of the average declaration (proposition 4.8). ■

In bad years a taxpayer's payoff is a direct consequence of her declaring truthfully her low income and getting no punishment or reward for doing so, regardless of the value of λ. The other two actors' payoffs, on the other hand, are different depending on the case under consideration. In good years, with full evasion, every taxpayer evades and, since no one is audited, they keep their gross incomes. In turn, since the agency audits no one, it suffers an expected loss of $\gamma\lambda$ because with probability γ the year is a good one and so everyone is an evader who is not caught (negligence errors) and with probability $1-\gamma$ the year is a bad one, everyone complies and nobody is audited (no zeal errors). Analogously, with full compliance, all taxpayers comply and so their disposable income is simply their gross income minus their voluntarily paid taxes, $1-t$. The expected loss of the agency is now $(1-\gamma)(1-\lambda)$ because with probability γ the year is a good one, everyone complies and nobody is audited (no negligence errors) and with probability $1-\gamma$ the year is a bad one and everyone complies but is audited anyway (zeal errors). The most interesting scenario is, however, the partial evasion one. Here, a soft agency ($\hat{s}-\varepsilon < \lambda < \tilde{\lambda}$) makes no targeting error whatsoever, thus reaching the best outcome it could aspire to. The rationale behind this result is that in good years some taxpayers will evade (those with low signals) while others will comply (those with high signals) and so the agency can perfectly distinguish evaders from compliant taxpayers, which implies that evaders are always caught (their payoffs are equal to gross income minus fine, $1-f$) while compliant taxpayers are never targeted (they get payoffs equal to gross income minus taxes $1-t$). In bad years, everyone declares zero and nobody is audited, so no mistake is made. A tough agency ($\tilde{\lambda} < \lambda < \hat{s}+\varepsilon$) will also catch every evader in good years, but will audit everyone in bad years, thus leading to the same expected loss than the "Full Evasion" case.

A surprising corollary can thus be stated: *"The relationship between the level of evasion and the agency's payoff is not monotonic."* Indeed, as we move from the right to the left column in equation 4.43 (i.e., as evasion increases), the agency's welfare first increases and then decreases. This means that the government is better off when it can create a *coordination game* among agents but, especially, when it in turn *makes taxpayers take different actions* (some evade, others comply), thus getting valuable information about the true income of the population and increasing its targeting accuracy.

To conclude the characterisation of the equilibrium, it is important to analyse how more accurate signals affect the level of evasion and the agency's payoff:

Proposition 4.11 *More precise information (formally, a lower ε) leads to:* (1) *(weakly) less compliance if the agency is soft ((weakly) more if it is tough), and* (2) *a (weakly) higher expected loss. Formally,*

$$(1) \quad \frac{\partial D^*}{\partial \varepsilon} \begin{cases} \geqslant 0 & \text{if } \lambda < \tilde{\lambda} \\ \leqslant 0 & \text{if } \lambda > \tilde{\lambda} \end{cases} \qquad\qquad (2) \quad \frac{\partial EL^*}{\partial \varepsilon} \leqslant 0 \qquad (4.44)$$

where $\tilde{\lambda} := 1 - \gamma$.

Proof. In appendix B.4, page 127. ■

The first part of the proposition highlights the fact that the impact of better information on the level of evasion depends on the type of the agency. This is at odds with previous studies, which usually find that better information leads to more evasion, through the argument that it decreases the risk borne by taxpayers who, assumed to be risk averse, have therefore more incentives to evade.

Though compelling, this argument cannot be applied to the present case because here agents are assumed risk neutral. Yet, what matters is that the relationship between compliance and accuracy of information is not *intrinsically* (weakly) increasing or decreasing, but rather one whose shape depends on the type of the government. Intuitively, when an agency is soft (λ is low) it dislikes targeting compliant taxpayers and so would audit with a very low probability. For signals of a given precision $\varepsilon > 0$, agents will estimate the probability of detection and decide their income declarations accordingly. If signals became more precise (if ε decreased), agents would be more aware of the fact that the agency is soft (in the extreme case, when $\varepsilon = 0$, they would know it with certainty), and so would expect a lower probability of detection, which in turn makes evasion relatively more attractive and leads to lower compliance. An analogous story can be used when the agency is tough (λ is high).

The second part of the proposition, on the other hand, shows that more accurate information is *never* good for the tax agency. Though previous studies also found this relationship, they relied on the above mentioned risk aversion of taxpayers and on the monotonic relationship between the level of evasion and the tax agency's payoffs (debunked by proposition 4.10).

The channel used here, on the other hand, hinges on the new feature introduced by the agency's policy: the coordination game played by taxpayers. From the agency's perspective, and using proposition 4.10, the coordination scenarios ("Full Evasion" and "Full Compliance" cases) are (weakly) dominated by the coordination failure one ("Partial Evasion" case). Since more precise information decreases the likelihood of the latter scenario (because the probability of $\lambda \in (\hat{s} - \varepsilon, \hat{s} + \varepsilon)$ decreases), then agencies prefer low-precision signals over very accurate ones. Note, however, that this benefit is only available to soft agencies ($\hat{s} - \varepsilon < \lambda < \tilde{\lambda}$), because it increases an agency's ability to distinguish evaders from compliant taxpayers in good years and thus decreases the number of negligence errors it makes. On the other hand, tough agencies ($\tilde{\lambda} < \lambda < \hat{s} + \varepsilon$) cannot take advantage of it as the situation in bad years (which is the origin of such agencies' zeal errors) is unaffected by a change in the informativeness of signals.

Alternatively, a lower ε can be interpreted as an increase in the degree of aggregation of information (or information-sharing). That is, if taxpayers shared their signals, the effect would be equivalent to an increase in their precision, since the group's average signal is expected to be closer to the true value of λ than the individual ones. In the limit, if all signals were shared, taxpayers would know the government's type with certainty –this is exactly the same result as if all signals were perfectly accurate (i.e., if $\varepsilon \to 0$).

4.5 Discussion

As every other model, the one developed here is built around some simplifying assumptions that make it more tractable and elegant, but also more restrictive and unrealistic.

Indeed, it could be argued that tax agencies do not follow a "bang-bang" policy such that either everyone is audited or nobody is, but rather one where a fraction of the population is audited while the rest is not. The first approach is a direct consequence of the "*ex-post* horizontal equity" condition, while the second one would fit a situation that satisfies the condition of "horizontal equity in expectation". The former is a stronger version of the latter, but also leads to situations where those who declare equal amounts are *effectively* treated equally, a desirable feature of an optimal auditing policy in my view. However, if the second approach were used, the results would not be significantly different from the ones presented in the text, the only "major" difference being that a tough agency would not audit everyone, but rather just a fraction of the population sufficiently large as to eliminate all incentives to evade (with the added benefit that the enforcement costs will be lower due to the smaller number of audits undertaken).

Also unlikely to be found in the real world is the dichotomous character of income assumed here. When more than two levels of income are allowed, the auditing decision with respect to a given individual depends on the *relative position* of the taxpayer's declaration compared to the rest of the population's: if it is among the highest ones, then the taxpayer's probability of detection is still (weakly) increasing in the agency's type and, under mild assumptions, (weakly) decreasing in the amount declared; if it is not, the agency knows the taxpayer is lying and audits her with certainty. When only two levels of income are considered, this policy collapses to the one presented earlier in this chapter.[9]

Along similar lines, it is clear that the assumption of perfect correlation among the taxpayers' incomes is an implausible one. However, it is just intended to capture the fact that usually taxpayers that belong to the same category are homogeneous in most aspects, including income. Relaxing it will not change the (qualitative) results, as long as the common shocks are maintained as the *main source of income variability.* This ensures that there is still a significant degree of correlation among incomes and, therefore, that other taxpayers' declarations convey *useful information* about the common shock that affects the category. Also very important for the analysis is the fact that incomes within a class are more homogeneous than the signals received by its members, such that the differences among them are mainly due to disparate perceptions of the government's type. Thus, the assumption of perfect uniformity allows us to observe the effect of the fundamental uncertainty unadulterated by the presence of income heterogeneity, and so the analysis is greatly simplified.

Finally, the importance of the partitioning of the taxpayer population into fairly homogeneous categories is highlighted by the fact that the above mentioned "relatively high correlation" condition is achieved when the category consists of agents that are very similar to each other in terms of their "observables" (age, profession, gender, etc.), since in this case their idiosyncratic shocks will be relatively small compared to the category-wide ones.[10] However, since the partitioning problem is an issue this paper is not concerned

[9] Also, irrespective of the levels of income allowed, if they are bounded above (i.e., $y_i \leq y_{\max}\ \forall i \in [0,1]$), the agency would never audit those who declare $y_{\max}$. In the more realistic case of unbounded domain, the probability of detection simply decreases as the declaration increases, as is standard in the literature.

[10] These "observables" refer to variables that are exogenous to (or costly to manipulate by) the agents, and so do *not* include taxpayers' current declarations.

with, the only related matter worth discussing here is the type of classes that favours the present model. And since the latter clearly relies on some degree of uniformity within the class, its predictions are more likely to fit the data from classes with a large number of rather homogeneous people (e.g., unskilled manufacture workers or non-executive public servants) than the ones from small and/or heterogeneous classes.

4.6 Conclusion

The question of a tax agency's optimal auditing strategy in the presence of common income shocks is relevant because it is not unusual for such shocks to be the main source of income variability for a group of fairly homogeneous taxpayers. Under these circumstances an agency's best policy consists of auditing those who declare low income with a probability that is (weakly) increasing in the declarations of the other taxpayers in the category. Intuitively, the higher these declarations, the more likely the shock was a positive one, and hence the more likely that someone who declares low income is an evader.

Implementing this policy does not require new information to be gathered by the agency, just using the available information better. Yet, it changes the nature of the problem for the taxpayers: on top of the standard cat-and-mouse game each one of them plays against the agency, they also play a coordination game against each other, a game in which a negative *externality* between them is *created*, a game taxpayers would *not* play if the policy were not contingent on the vector of declarations.

The heterogeneity in private signals eliminates the policy design difficulties that the multiplicity of equilibria appears to generate and paves the way for modelling the problem as a global game which not only is more realistic, but also predicts a unique equilibrium which is consistent with empirical evidence.

Chapter 5

Anti-Evasion Policy: Experiment

5.1 Introduction

Common income shocks that affect similar agents in similar ways are well documented: airlines' sales plummeted after 9/11, chicken breeders faced low demand after the avian flu outbreak, and emergent markets have difficulties attracting investors every time the U. S. Federal Reserve increases interest rates. As these examples show, often the common shocks are the main source of income variability, with a *common/idiosyncratic* ratio well above 1.

It is therefore not surprising that a tax agency that ignores common income shocks when deciding its auditing policy will act suboptimally. But this is exactly what happens if they follow the most popular policy prescribed by the literature: the "cut-off rule" (Reinganum and Wilde (1988)) which states that the agency should not audit any firm that declares about or above a certain fixed cut-off income level, while auditing those who declare below it with a sufficiently high probability. Combined with common income shocks, this policy leads to systematic mistargeting: the agency audits "too much" in bad years and "too little" in good ones.

In this environment the optimal policy is a contingent rule in which the agency audits every firm with a probability that is a non-decreasing function of every other taxpayer's declarations. This is because other firms' declarations give the agency information about the realisation of the shock and so the probability of a given taxpayer being an evader is (weakly) higher the higher are her fellow taxpayers' declarations.

The purpose of the present analysis is therefore to test the *TEGG* ("Tax Evasion as a Global Game") model presented in chapter 4. This is a relevant task because it will help determining which of the alternative rules (contingent or cut-off) is superior to the other and, indirectly, whether the data is consistent with the modelling of tax evasion as a global game and its associated predictions.

However, real-world data on tax evasion is not readily available. Those who engage in tax evasion are not willing to declare it and tax agencies are reluctant to provide the information because of the confidentiality of tax returns: even if the datapoints are not labelled, in many cases it is quite easy to identify which individual firm they belong to, thus revealing sensitive information about taxpayers.

For this reason, the analysis will use the second-best available dataset, namely, the one collected in a computerised experiment in which participants interacted with each other in situations that resembled the scenario described by the TEGG model. This methodology has the obvious disadvantage of making difficult the extrapolation of results from the sample to the population, but allows the experimenter a greater control over the variables under study and is, as mentioned before, the only available one anyway.

The econometric analysis finds that the agency is better off when using the contingent rule than when using the cut-off one, and so that the key prediction of the TEGG model is strongly supported: the associated, *artificially created negative externality* does increase the payoff of the agency. The data also support the hypothesis that people make decisions consistent with higher-order beliefs (which play an important role in ensuring the uniqueness of the global game equilibrium) and that the comparative statics follow the ones predicted by the global game method.

Although there are many laboratory experiments framed as tax compliance problems (see e.g. Baldry (1986), that compares tax evasion and gambling, Alm et al. (1992), that investigates the effects of institutional uncertainty, and Alm and McKee (2004), which analyses tax evasion as a coordination game) none has investigated tax evasion as a global game (see figure 5.1). The latter requires not only the strategic uncertainty generated by the coordination game but also the "fundamental uncertainty" created by the incompleteness of information regarding the payoff functions. Tests of the global game technique seem to support it in terms of predictive power (Cabrales et al. (2002)) and/or comparative statics (Heinemann et al. (2004b)), but are less supportive of the participants' use of higher-order beliefs when making decisions. The latter result is also the conclusion of other studies, like Stahl and Wilson (1994) and Bosch-Domenech et al. (2002).

		Strategic uncertainty	
		No	Yes
Fundamental uncertainty	No	Baldry (86)	Alm & McKee (04)
	Yes	Alm *et al.* (92)	This study

Figure 5.1: Position in the tax evasion experimental literature.

5.2 Experiment design

The experiment took place on the 28 of November 2006, at the ELSE computer laboratory (UCL, London). The pool of participants was recruited by ELSE from their database of about 1,000 people (most of them UCL students). 200 of them were chosen randomly and invited to take part and the first 100 who accepted the offer were allocated to sessions according to their time preferences.[1] No person was allowed to participate in more than one session.

The day of the experiment 76 people took part in four treatments (labelled GC, GE, LC and LE), each involving a 60-to-90-minute long session. The treatments were defined according to the policy used (contingent v cut-off, or "global" (G) v "lottery" (L)) and the predicted optimal strategy of the participants (which for this experiment, as will be shown later, reduces to determining the optimal choice when signal b is received: to evade (E) or to comply (C)).[2] This way the experimental setup can be visualised as in table

[1] Five "reserve" people were invited to each session and 7 of them had to be turned down because the target number (20 per session) was reached or because the treatment required an even number of participants (treatments GC and GE). Each one of them was paid the £5 show-up fee before being dismissed.

[2] Tax evasion has often been compared to a gamble in which the taxpayer "wins" (i.e., gets away with evasion) with probability w, and "loses" (i.e., is caught and has to pay a fine on top of the unpaid taxes) with probability $1 - w$.

The cut-off rule is equivalent to a standard lottery (and hence the name of the treatment) because it *fixes* the chances of winning (say $w = 1 - p$) and losing ($1 - w = p$). Evasion can therefore be seen as equivalent to buying $(1 - p)\,N$ out of a pool of N raffle tickets, each one of them equally likely to be the winner.

In the Global treatments, on the other hand, those probabilities are *not* fixed, because they are affected by what other people do. In particular, since other people's compliance has a negative impact on my payoff, the fact that other people comply is equivalent to having the total number of tickets increased to, say, $N' > N$, so that my probability of winning w' (in spite of my holding the same number of tickets as

5.1.

		Participant's strategy	
		Comply (C)	Evade (E)
Auditing rule	Contingent (G)	GC	GE
	Cut-off (L)	LC	LE

Note: Participant's strategy refers to the optimal strategy of a participant when receiving hint b.

Table 5.1: Treatments.

Participants were lined up outside the lab according to their arrival time. At the designated time they entered and freely chose where to sit. They were not allowed to communicate for the entirety of the session and could not see other people's screens.

Each session consisted of 6 stages, namely, instructions, short quiz, trial rounds, experimental rounds, questionnaire and payment. The instructions were read aloud by the instructor and, in order to ensure their correct understanding, the participants were asked to complete a "short quiz" (shown in appendix C.1; correct answers and the rationale for them were provided by the instructor after a few minutes). For the same reason, participants then played two "trial" (practice) rounds whose outcomes did not affect their earnings. After each of these first three stages the instructor answered subjects' questions in private. The experimental rounds (20 per session) were then played, and after that, subjects completed a questionnaire with information regarding personal data and the decision-making process they followed. Finally, each participant was paid an amount of money consisting of a fixed show-up fee (£5) and a variable component equal to the earnings accumulated over the 20 experimental rounds.[3] Table 5.2 shows the exchange rate used to translate experimental currency into money, as well as other payment-related summary statistics.[4]

before, $(1-p)\,N$) is now comparatively lower: $w' = \frac{(1-p)N}{N'} < \frac{(1-p)N}{N} = w$.

[3]In other experimental studies (Heinemann et al. (2004a) among them) participants were paid according to the result of one randomly-chosen round. The rationale for this is that it avoids hedging, something that is not a problem here: the maximum payment a person can receive in any given round is £0.50 or £0.90 (depending on the treatment), with expected values in the £0.30-£0.35 range.

[4]In order to minimise delays and computational hassle, every person's payment was rounded *up* to the closest multiple of £0.20. Participants were not told about this arrangement until after they completed their questionnaires in order to avoid strategic play with respect to this peripheral matter.

Treatment	Participants	£ per 1000 points	Min/Avg/Max Payment		
GC	18	0.50	10.80	11.52	11.80
GE	18	0.90	7.40	9.30	9.80
LC	20	0.50	11.60	11.65	11.80
LE	20	0.90	9.80	11.20	11.60
All	76		7.40	10.95	11.80

Note: £ per 1000 points is the exchange rate at which 1000 "experimental points" where transformed into pounds.

Table 5.2: Treatments. Participants and Money.

Each experimental round consisted of two stages: the "Choice" one, where participants had to make a decision that would affect their payoffs, and the "Feedback" one, where they got information about their outcomes for the round.

		Column player	
		Y	Z
Row player	Y	$x(Y,Y,q)$	$x(Y,Z,q)$
	Z	$x(Z,Y,q)$	$x(Z,Z,q)$

Note: Only Row player's payoffs (x) are shown. Payoff's components are Row's action, Column's action and the realisation of the random variable q. Column's payoffs are symmetrical.

Table 5.3: Stage game.

In the "choice" stage a one-shot game was played where the subjects had to choose one of two possible actions (Y or Z) interpreted as *Evasion* and *Compliance*, respectively (the game's normal form for the 2-person case is shown in table 5.3). A participant i's payoff was determined by her own decision, $d_i \in \mathcal{D} := \{Y, Z\}$, the decisions of the other $n-1$ people in her category, $\mathbf{d}_{-i} := (d_1, ..., d_{i-1}, d_{i+1}, ..., d_n)$, $\mathbf{d}_{-i} \in \mathcal{D}^{n-1}$, and the realisation of a random variable, $q \in \mathcal{Q} := \{A, B, C\}$. Formally,

$$x_i := x(d_i, \mathbf{d}_{-i}, q) \tag{5.1}$$

Different choices have different effects on payoffs, and so, while the payoff from option Y is uncertain (reflecting the uncertainty about being audited), that of option Z is a known, fixed quantity: formally, for every $\mathbf{d}_{-i}, \mathbf{d}'_{-i} \in \mathcal{D}^{n-1}$, $q, q' \in \mathcal{Q}$,

$$x(Z) := x(Z, \mathbf{d}_{-i}, q) = x(Z, \mathbf{d}'_{-i}, q') \tag{5.2}$$

The random variable q can take values A, B and C with probabilities $p(A) = 0.20$, $p(B) = 0.60$ and $p(C) = 0.20$, respectively. It represents the different possible "types" of agency (A : soft on evasion, B : medium, C : tough) and corresponds to the "λ" mentioned in the chapter 4 (page 38). It affects evasion payoffs (i.e., Y-payoffs) negatively: the tougher the agency, the more likely the evader will be audited and the lower her payoff. Formally, for every $\mathbf{d}_{-i} \in \mathcal{D}^{n-1}$,

$$x(Y, \mathbf{d}_{-i}, A) > x(Y, \mathbf{d}_{-i}, B) > x(Y, \mathbf{d}_{-i}, C) \tag{5.3}$$

At the time of making a decision participants do not know the value of q, but each one of them gets a private signal $s \in \mathcal{S} := \{a, b, c\}$ (called "hint" in the experiment) that is related to the realised value of q as shown in table 5.4 (and in the Instructions sample in appendix C.1). The instructions highlighted the fact that different people could get different hints but q was the same for everyone. No other probabilities were provided explicitly, though the instructions did supply the information required for their computation, namely, the prior probability distribution of q, $p(q)$, and the conditional probability, $f(q|s)$.[5]

If hint =...	...then q =...	...with probability $f(q\|s)$ =...
a	A	1.000
	A	0.125
b	B	0.750
	C	0.125
c	C	1.000

Table 5.4: Hints and q.

The participant's submission of her decision (Y or Z) ended the "Choice" stage and gave way to the "Feedback" one, in which the person was informed about the realised value of q, the signal she received, her choice and her payoff for the round. At no stage was a subject given any information about the signals or choices of any other participant.[6]

By clicking on the "Continue" button, participants exited the "Feedback" stage and moved on to the next round (if any was left). Rounds were identical to each other in terms of their structure (Choice and Feedback stages) and rules (payoff computations, prior and conditional probability distribution of q), but may have differed in the *realised* values of

[5] A "Choice stage" screenshot (labelled "Choice screen" in the experiment) can be seen in the instructions sample in appendix C.1, figure C.1. The programme used was z-Tree (Fischbacher (2007)).

[6] A "Feedback stage" screenshot (labelled "Results screen" in the experiment) can be seen in the instructions sample in appendix C.1, figure C.2.

the random variables (q and s). Participants were told explicitly about this and informed that each round was independent from every other one.

5.2.1 Treatments

As shown in table 5.1, treatments were defined according to the type of game created by the agency's policy (Global or Lottery) and the predicted optimal strategy of the participants (Evade or Comply).

The difference between Global and Lottery treatments is related to the effect of other subjects' choices on the payoffs of individual participants. In the Lottery treatments the rule implemented by the agency is of the cut-off type, and so what other people do does not affect taxpayer i's payoff. Formally, for every $q \in \mathcal{Q}$,[7]

$$x(Y, Y, q) = x(Y, Z, q) \tag{5.4}$$

In Global treatments, on the other hand, the auditing policy followed is the contingent one, implying that other people's declarations have a negative impact on taxpayer i's payoff via the increased probability of detection (as in proposition 4.6 in chapter 4). Formally, for every $q \in \mathcal{Q}$,

$$x(Y, Y, q) > x(Y, Z, q) \tag{5.5}$$

It is worth mentioning here that the Lottery treatment can be interpreted as a special (limit) case of the Global treatment in which the effect of other people's decisions on a certain participant's payoff is arbitrarily small. Consequently, and without loss of generality, henceforth the analysis will be restricted to the Global case, with the occasional reference to the Lottery one provided only when relevant.

For the experiment, participants in the Global treatments were divided into 9 groups of 2 people each, the matching protocol being random (equi-probable) within rounds and independent across them. The experimental setup reproduced the three typical scenarios described by the global game literature:

† The two **extreme** cases in which the "fundamentals" are "so bad"/"so good" that there is a strictly dominant strategy which is chosen by everyone. In the experiment the fundamental is q, the agency's "toughness", and so strict dominance requires that

[7] I restrict my attention to the 2-person case, which will be the relevant one throughout the paper. The extension to the n-person case is straightforward.

everyone should evade when the agency is very soft ($q = A$) and that everyone should comply when it is very tough ($q = C$). Formally, for every $d' \in \mathcal{D}$,

$$x(Y, d', A) > x(Z) \quad (5.6)$$

$$x(Y, d', C) < x(Z) \quad (5.7)$$

† The **intermediate** one in which the "fundamentals" are not so bad but no so good either. In this case a coordination game is created and, consequently, no strategy dominates all others: which one is optimal depends on what other people do. In the experiment, this corresponds to the scenario in which the agency is "medium" ($q = B$): if the other person in my group evades, it is optimal for me to evade as well; if the other person complies, I am better off if I comply too.[8]

Formally,

$$x(Y, Y, B) > x(Z) > x(Y, Z, B) \quad (5.8)$$

Turning now to the other dimension that defines treatments, the difference between the Evasion and Compliance cases is due to their different predictions regarding what a participant's optimal strategy should be. Thus, distinguishing E from C treatments demands the solving of the taxpayer problem, namely, that of choosing between Evasion (Y) and Compliance (Z) using all the information available (s) in order to maximise expected utility. In this setup, therefore, a taxpayer's strategy $\boldsymbol{\sigma}$ is a vector of decisions, one for each possible signal $s \in \mathcal{S}$. Formally, $\boldsymbol{\sigma} := (\sigma(a), \sigma(b), \sigma(c))$, where $\sigma : \mathcal{S} \to \mathcal{D}$ is a function that maps signals into decisions.[9] Therefore, finding the solution requires comparing the (certain) utility of compliance, $u(Z)$, and the expected utility from evasion,

$$Eu(Y, \mathbf{k}'(\mathbf{s}') | s) := \sum_{q \in \mathcal{Q}} f(q|s) \sum_{s' \in \mathcal{S}} \Pr(s'|q) \cdot \left\{ k'(s') u(Y, Y, q) + [1 - k'(s')] u(Y, Z, q) \right\} \quad (5.9)$$

where $u(Y, d', q) := u(x(Y, d', q))$ is the utility from receiving payoff $x(Y, d', q)$; $s' \in \mathcal{S}$ and $d' \in \mathcal{D}$ are respectively the signal and decision of the other member of the group; $\Pr(s'|q) \in [0, 1]$ is the conditional probability of the other member getting a signal s' given that the agency's type is q; $\mathbf{k}'(\mathbf{s}') := (k'(a), k'(b), k'(c))$; and $k'(s') \in [0, 1]$ is the taxpayer's belief about the probability of evasion of a group-mate that receives a signal s'.

[8] Clearly, this condition does not apply to the Lottery case.

[9] Actually, it maps signals into *probability distributions* over decisions, if one allows for mixed strategies. However, this possibility was explicitly ruled out here because its inclusion would not have provided any extra, significant insight as to justify the complexity-associated problems it would have entailed.

This comparison depends crucially on the beliefs of the optimiser with respect to the actions of the other member of the group, $\mathbf{k}'(\mathbf{s}')$, and, therefore, on the ability and sophistication of subjects at forming them, a matter that is directly related to the concepts of common knowledge and higher-order beliefs (HOBs, Carlsson and van Damme (1993)). These HOBs refer to the levels of reasoning involved in reaching a conclusion and are neatly connected to the Iterated Deletion of Strictly Dominated Strategies (IDSDS) method: for each iteration, the order of beliefs increases one level. Furthermore, HOBs are the key factor behind the uniqueness of the global game equilibrium: In the first iteration, $i = 1$, my private signal gives me information about the set of strategies (out of the original set, Σ^0) that are strictly dominated (*SDed*) by others and will therefore *never* be played. In the second iteration, $i = 2$, the set of those strategies that survived the previous round of deletions is the new feasible set, Σ^1. Via an analogous mechanism, a new group of *SDed* strategies will be discarded and after that a new iteration $i = 3$ with feasible set Σ^2 will begin. The theory of global games proves that in the limit, after an arbitrarily large number of iterations, the feasible set Σ^∞ has only one element, $\boldsymbol{\sigma}^*$. In other words, the equilibrium is unique (Morris and Shin (2002b)).

In the present experiment, only 2 iterations are needed to find the unique solution to the taxpayer problem.[10] Depending on the number of iterations used (1 or 2), a taxpayer is then classified as "Rudimentary" or "Advanced", respectively. Their behaviour is summarised in the following two propositions.

Proposition 5.1 (RDom) ***Rudimentary Dominance:*** *According to Rudimentary taxpayers (RTPs): 1. if $s = a$ (signal is low), Evasion strictly dominates (SDs) Compliance: $Y \succ_R Z$; 2. if $s = b$ (signal is medium), no strategy SDs the other: $Y \not\succ_R Z$ and $Z \not\succ_R Y$; and 3. if $s = c$ (signal is high), Compliance SDs Evasion: $Z \succ_R Y$.*

Proposition 5.2 (ADom) ***Advanced Dominance:*** *According to Advanced taxpayers (ATPs): 1. those strategies that are rudimentary-dominated (parts 1 and 3 of proposition 5.1) are also advanced-dominated; and 2. if $s = b$ (signal is medium), then: 2.a. in E treatments, Evasion SDs Compliance: $Y \succ_A Z$; and 2.b. in C treatments, Compliance SDs Evasion: $Z \succ_A Y$.*

At this point, it is worth defining the concepts of Soft, Medium and Tough games, which are simply the games played by the members of a group when the agency is soft, medium and tough, respectively (i.e., they are like the game shown in table 5.3, with $q = A$, B

[10] This does not apply to Lottery treatments for the obvious reason that in those cases, by definition, a taxpayer's payoff does not depend on other people's choices or the taxpayer's beliefs about them.

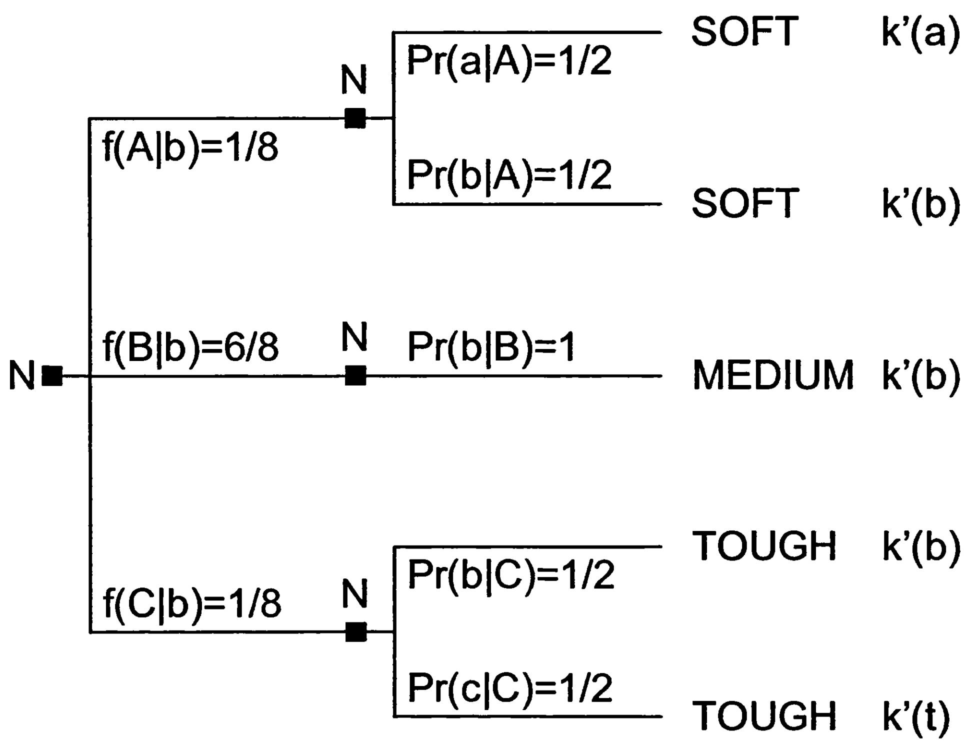

Figure 5.2: Game tree if signal is medium ($s = b$).

and C). Clearly, these games $g \in \mathcal{G} := \{S, M, T\}$ depend on the type of the agency, and so both g and q are subject to the same probabilistic process.

Based on this taxonomy of games and on the conditional probability distribution of q (shown in table 5.4), two different scenarios can be identified: one in which the signals give perfect information about the game being played, and another one in which precision is less than perfect.

In the first iteration, therefore, a taxpayer who receives a soft signal ($s = a$) knows for sure that she is playing the Soft game ($g = S$). Furthermore, because of equation 5.6, she can immediately realise that Evasion SDs Compliance, the very result indicated in part 1 of proposition 5.1. Following a similar argument and using equation 5.7, part 3 is also proved.

When the signal is medium ($s = b$), though, the person does not know the actual game g that is played, but she does know its conditional probability distribution $f(g(q)|b) = f(q|b)$. Thus, the game that she faces is depicted in figure 5.2, and her expected utility from evasion is given by equation 5.9, where s is replaced by b. This expression is an increasing function of the beliefs about the other member's probability of evasion, $k'(s')$, $\forall s' \in \mathcal{S}$, because of the nature of the contingent policy (equation 5.5). The worst-case scenario for the optimising person occurs, therefore, when she expects the other member to choose Compliance irrespective of the signal received ($\mathbf{k'(s')} = \mathbf{0}$, $\mathbf{0} := (0,0,0)$), such that the expected utility from evasion is $Eu(Y, \mathbf{0}|b)$. Analogously, the best case scenario corresponds to that in which the other member always evades ($\mathbf{k'(s')} = \mathbf{1}$, $\mathbf{1} := (1,1,1)$) and expected utility is $Eu(Y, \mathbf{1}|b)$. It is not difficult to see that the no-strict-dominance condition of proposition 5.1 (part 2) requires

$$Eu(Y, \mathbf{0}|b) < u(Z) < Eu(Y, \mathbf{1}|b) \tag{5.10}$$

and if it is satisfied, a Rudimentary taxpayer will act exactly as predicted by the *RDom* proposition.[11]

A Rudimentary taxpayer would stop her analysis here, but an Advanced taxpayer will continue to the next iteration. Furthermore, the ATP will realise that, if the other member of her group is (at least) Rudimentary, then (by symmetry) she would have also worked out that Evasion is the strictly dominant strategy when the signal received is soft (a). Following an analogous argument, the ATP will also realise that the other person will work out that Compliance is the strictly dominant strategy when the signal received is tough (c). Formally, the ATP's beliefs about the other person's choices will have precise numbers attached to them, namely, $k'(a) = 1$ and $k'(c) = 0$. The expected utility will reflect this in general, $Eu(Y, (1, k'(b), 0)|b)$, as well as in the worst- and best-case scenarios, $Eu(Y, \mathbf{c}|b)$ and $Eu(Y, \mathbf{e}|b)$,where $\mathbf{c} := (1,0,0)$ and $\mathbf{e} := (1,1,0)$.

Depending on the position of the safe utility $u(Z)$ with respect to $Eu(Y, \mathbf{c}|b)$ and $Eu(Y, \mathbf{e}|b)$, three cases can arise, of which we are interested only in the following two:[12]

$$Eu(Y, \mathbf{c}|b) > u(Z) \tag{5.12}$$

$$u(Z) > Eu(Y, \mathbf{e}|b) \tag{5.13}$$

[11] An alternative interpretation of this equation that will be used later is the following. Let us construct a new, artificial 2x2 game like the one in table 5.3, but which is a weighted average of the Soft, Medium and Tough games defined above, $A := \sum_{q \in \mathcal{Q}} f(q|b) \cdot g(q)$, so that the corresponding (expected) utility in each of its cells is

$$u(d, d', E(q|b)) := \sum_{q \in \mathcal{Q}} f(q|b) \cdot u(d, d', q) \tag{5.11}$$

It can then be shown that $u(Y, Z, E(q|b)) = Eu(Y, \mathbf{0}|b)$, $u(Y, Y, E(q|b)) = Eu(Y, \mathbf{1}|b)$, and $u(Z, Y, E(q|b)) = u(Z, Z, E(q|b)) = u(Z)$, so that equation 5.10 implies that this "Average game" is a coordination game.

[12] The third one does not lead to a unique solution, which goes against the spirit of the theory of global games. The reason for the non-uniqueness is the discreteness of the model. Having continuous choices may have avoided this problem, but at the cost (considered to be too high) of increasing the complexity of the game and thus the noise in the observations.

Condition 5.12 indicates that even in the worst-case scenario, the expected utility from Evasion is higher than that from Compliance or, equivalently, that the former *SDs* the latter. Equation 5.13, on the other hand, implies that, even in the best-case scenario, the expected utility from Evasion is lower than that from Compliance, and so that the second *SDs* the first.

Which of these two mutually exclusive conditions is satisfied determines the taxpayer's optimal strategy: either $\boldsymbol{\sigma}^* = (Y, Y, Z)$ if (5.12) holds or $\boldsymbol{\sigma}^* = (Y, Z, Z)$ if (5.13) holds. These strategies are of the "threshold" type (Heinemann et al. (2004a), Heinemann et al. (2004b)) but can be indexed by their second component, which is the only one that differentiates one strategy from the other and corresponds to the optimal choice when the signal is medium, $\sigma^*(b)$. The value of this component, therefore, is the one that defines the Evasion, $\sigma^*(b) = Y$, and Compliance, $\sigma^*(b) = Z$, treatments.

The rationale for including these two types of treatments reflects, above all, the lack of theoretical predictions or stylised facts about what strategy we should expect to be played in the medium case. Its presence, however, allows for the testing of some hypotheses regarding the comparative statics of global games: in particular, the change of parameters predicts that the number of people receiving b signals that choose Y should be greater in E-treatments than in C-ones, while no significant difference should exist if signals are soft (a) or tough (c).[13] This is summarised as follows:

Hypothesis 5.1 (OS) ***Optimal strategy:*** *1. If signal is soft ($s = a$) then evade ($d = Y$); 2. if signal is tough ($s = c$) then comply ($d = Z$); and 3. if signal is medium ($s = b$) then: 3.a. in E treatments, evade ($d = Y$); and 3.b. in C treatments, comply ($d = Z$).*

If choices satisfy all three parts of the hypothesis, then one can say they are "consistent with the *ADom* predictions" and label the taxpayer as "Advanced". If they only satisfy the first two parts, they are "consistent with the *RDom* predictions" and the taxpayer can be labelled as "Rudimentary".

[13]For L-treatments, the analysis is greatly simplified since other people's choices do not affect one's decisions. Then, the equivalents of equations 5.12 and 5.13 are, respectively,

$$Eu(Y, b) > u(Z) \tag{5.14}$$
$$u(Z) > Eu(Y, b) \tag{5.15}$$

5.2.2 Selection of payoffs

Turning now to the main prediction of the TEGG model, it is clear that, in order to test which of the auditing rules (contingent or cut-off) is better, one needs a "level playing field". In this setup, it requires the enforcement costs to be the same in G- and L-treatments, which further simplifies to undertaking the same (expected) number of audits (for a given value of q) in each treatment. This way, the mere comparison of the errors made by each type of agency across treatments will indicate which rule is superior to the other (if any). Formally,

Hypothesis 5.2 (SCR) ***Superiority of Contingent Rule:*** *Given a fixed level of enforcement, Global treatments yield less (expected) errors than Lottery ones for all possible types of agency, $q \in \mathcal{Q}$.*

The expected number of audits is

$$Ea\left(d, d', q\right) = \sum_{y \in \mathcal{Y}} \Pr(y) \sum_{s \in \mathcal{S}} \Pr(s|q) \sum_{s' \in \mathcal{S}} \Pr\left(s'|q\right) \sum_{d \in \mathcal{D}} \Pr(d|s, y) \sum_{d' \in \mathcal{D}} \Pr\left(d'|s', y\right) \sum_{a \in \mathcal{A}} \Pr\left(a|d, d', q\right) \sum_{a' \in \mathcal{A}} \Pr\left(a'|d, d', q\right) \cdot \left(a + a'\right) \quad (5.16)$$

where $\mathcal{Y} := \{0, 1\}$ is the set of possible values that income y can take (1 in "good" years and 0 in "bad" ones), $\mathcal{A} := \{0, 1\}$ is the set of possible values that an audit can take (1 if the audit is undertaken and 0 if it is not) and $a, a' \in \mathcal{A}$ are the agency's decisions regarding whether to audit or not the taxpayers of a group.[14]

Thus, the equalisation of enforcement costs demands that, for each and every value of the agency's type $q \in \mathcal{Q}$, the expected number of audits in G-treatments must be equal to the corresponding one in L-treatments. Since payoffs are linear functions of the probabilities, the equalisation requires the following conditions to hold:[15]

[14] In the experiment, however, implicit in the stage game (table 5.3) is the assumption that all taxpayers have high income ($y = 1$). The reasons for this are that introducing the possibility of low income periods will not add to our knowledge (trivially, if $y = 0$ everyone declares truthfully) and that all interesting hypotheses to test are related to the high-income scenario (not to mention the extra cost and time that running this expanded experiment will demand).

[15] See appendix C.2 for the derivataion of these conditions.

LE	GE	LC	GC
$u^{LE}(Y,A) =$	$u^{GE}(Y,Y,A)$	$u^{LC}(Y,A) =$	$\sum_{d'\in\mathcal{D}} \omega(d') \cdot u^{GC}(Y,d',A)$
$u^{LE}(Y,B) =$	$u^{GE}(Y,Y,B)$	$u^{LC}(Y,B) =$	$u^{GC}(Y,Y,B)$
$u^{LE}(Y,C) =$	$\sum_{d'\in\mathcal{D}} \omega(d') \cdot u^{GE}(Y,d',C)$	$u^{LC}(Y,C) =$	$u^{GC}(Y,Y,C)$

Note: The weights, ω(d'), d'$\in$D, are functions of the parameters of the problem. In the experiment, ω(Y)=1/6 and ω(Z)=5/6.

Table 5.5: Payoffs. Conditions for the equalisation of enforcement costs.

Focusing on the Evasion case (columns 1 and 2), it requires that the payoff of evaders when the agency is soft (A) or medium (B) should be the same regardless of the auditing rule. This is logical because in these cases taxpayers can only get soft or medium signals (a or b) and, because of Rudimentary and Advanced dominance (equations 5.6 and 5.12/5.14), they will always evade. Since the agency cannot tell whether the pair of low declarations is the result of a bad year or of evasion, it will audit with the same probability in both G and L treatments (no difference in the information available to the agency) and so the (expected) payoffs of taxpayers are the same as well. When the agency is tough, however, taxpayers could receive medium or tough signals, $s \in \{b, c\}$, with those receiving $s = b$ choosing Y (because of equation 5.12 or 5.14) and the others choosing Z (equation 5.7). In the L-treatments the agency only uses the information derived from an individual's action; in G-treatments, on the other hand, it also uses the information derived from the other person's choice. In particular, if choices are different from each other, the agency knows that the one who chose Y is likely to be lying and can therefore audit her with a higher probability and the other one with lower probability. This is the basic mechanism behind the contingent rule and relies heavily on the extra difficulty taxpayers face when trying to coordinate on the Full Evasion equilibrium (strategic uncertainty). The actual position of $u^{LE}(Y,C)$ between $u^{GE}(Y,Y,C)$ and $u^{GE}(Y,Z,C)$ is determined by the weights $\omega(d') \in [0,1]$, $d' \in \mathcal{D}$, which depend on the parameters of the problem (especially $p(q)$ and $f(q|s)$) and in the experiment take the values $1/6$ (for $d' = Y$) and $5/6$ (for $d' = Z$). The conditions for the C-treatments (columns 3 and 4) are found following a similar argument.

The parameters chosen for the four treatments are therefore the ones shown in table 5.6.

Person 1's choice	Person 2's choice	Type of agency	GC	GE	LC	LE
Y	Y	A	1,000	1,000	715	1,000
Y	Y	B	655	145	655	145
Y	Y	C	579	6	579	1
Y	Z	A	658	156	715	1,000
Y	Z	B	651	135	655	145
Y	Z	B	0	0	579	1
Z	{Y,Z}	{A,B,C}	654	140	654	140

Note: Only payoffs of Person 1 are shown. Those of Person 2 are symmetric.

Table 5.6: Payoffs. All treatments.

These payoffs satisfy all the conditions mentioned so far: "type of agency" (equation 5.3), "global game" (equations 5.6, 5.7 and 5.8), "average game" (equation 5.10), "equal enforcement costs" (table 5.5), and the conditions that define treatments: $L\ v\ G$ (equations 5.4 and 5.5) and $E\ v\ C$ (equations 5.12 and 5.13, or 5.14 and 5.15).

The actual vector of values chosen is just one among many that satisfy the above-mentioned conditions. The feasible set was narrowed down by setting, without loss of generality, the maximum and minimum payoffs in the G-treatments equal to $1,000$ and 0 respectively, and by restricting attention to natural numbers.[16] Noting that payoffs in L-treatments are deterministic functions of those in G-treatments (see table 5.5), only 10 parameters remain to be determined, namely, the intermediate payoffs of the GC and GE treatments (including the safe payoffs). Before getting to it, however, a digression about equilibrium selection is in order here.

The global game (GG) technique selects one of the equilibria of a coordination game, an equilibrium that coincides (for 2×2 games like the ones used here) with the one selected by the "risk dominance" criterion, RD (Harsanyi and Selten (1988)). Intuitively, the latter chooses the equilibrium which, if abandoned, inflicts the highest costs on the players. Since the criterion applies to the global game we need to consider all three possible scenarios it usually entails: the two extreme ones and the intermediate one mentioned on page 66. In this particular case, however, it is enough to concentrate on the "average game" (defined in footnote 11), since it neatly summarises the whole game and thus simplifies the analysis. Because this "average game" is a coordination game, it will have 2 pure-strategy equilibria: one in which both players choose Y and get $Eu(Y, 1|b)$, and another one in which they both choose Z and get $u(Z)$. Which of the two is risk-dominant depends on

[16] To simplify computations and understanding by subjects, as well as to avoid prospect-theoretical interpretations (which, though interesting in themselves, are not the focus of the present analysis).

the relationship between $l_E := Eu(Y, \mathbf{1}|b) - u(Z)$ (the loss from deviating from the Full Evasion equilibrium) and $l_C := u(Z) - Eu(Y, \mathbf{0}|b)$ (the loss from deviating from the Full Compliance one). If deviating from (Y, Y) is more costly than deviating from (Z, Z) (i.e., if $l_E > l_C$), the risk-dominant equilibrium (RDE) is (Y, Y); otherwise, it is (Z, Z). In the experiment, the RDE depends on the treatment: it is (Y, Y) in GE and (Z, Z) in GC. These are, not surprisingly, the choices that equations 5.12 and 5.13 predicted to be optimal in those treatments, thus confirming that both the global game theory and the RD criterion select the same equilibrium.

There is, however, an important competitor for the RD/GG criterion: the payoff-dominance criterion, PD (Harsanyi and Selten (1988)). It simply states that if all equilibria can be Pareto-ranked, players will coordinate on the dominant one. In the experiment, the payoff-dominant equilibrium (PDE) is always (Y, Y) regardless of the treatment, because of the Average game being a coordination game and the contingent policy penalising evaders in case of coordination failure.

Thus, the PD and RD criteria select the same equilibrium (Y, Y) in the GE treatment but different equilibria ((Y, Y) and (Z, Z), respectively) in the GC one. Since the criteria reinforce each other in GE but compete against each other in GC, this suggests an interesting hypothesis to test:

Hypothesis 5.3 (RF) *Relative frequency GE/GC: The frequency of choices consistent with the GG/RD prediction will be (weakly) higher in GE than in GC.*

The main hypothesis of interest, however, is whether data fits the global game predictions (hypothesis 5.1). Thus, the 10 "free" parameters in table 5.6 were chosen to make the satisfaction of the predictions as difficult as possible, i.e., by making the RDE as little risk-dominant as possible. This required minimising l_E and maximising l_C in GE,and the opposite in GC. This way, if the data supports the global game predictions in these most demanding conditions, then the theory could be expected to be an even better predictor in more favourable environments.

Finally, it is important to mention here that risk aversion could dramatically alter the predictions of the model, and this may be especially important since evidence indicates that attempts to induce risk-preferences seem not to work (Selten et al. (1999)). The solution implemented in the experiment was to choose parameters such that all constraints will be satisfied for a large range of risk preferences. In particular, in E-treatments the parameters of table 5.5 are robust to degrees of relative risk aversion as high as 0.4 (about 60% of the population, according to Holt and Laury (2002)). For C-treatments, they are robust for values as low as 0 (about 80% of the population, according to the same study). Also, it is acknowledged in the experimental literature that when playing complex games people often avoid the complications of utility maximisation and instead simply maximise payoffs, which implies that risk preferences should not be an important issue here (probably most of the participants will end up acting as if their degrees of risk aversion were somewhere in the $[0, 0.4]$ range).

5.3 Results

A total of 1,520 observations were collected in the experiment; table 5.7 shows the breakdown by treatment. It also shows summary statistics of the key variables needed for testing the hypotheses of the previous section: "Dominance" and "Errors". The first one measures the coincidence between the data and GG theoretical predictions about the subjects' choices (DOM=1 if data fits predictions and 0 otherwise). Its name reflects the fact that those predictions are based on the concepts of dominance (propositions 5.1 and 5.2). The second one quantifies the number of errors (per observation/datapoint) made by the agency (ERR=1 if an error was made, 0 otherwise). Note that Dominance is never lower than 50% and Errors never above 35%.

Treatment	Observations	Dominance (DOM)		Errors (ERR)	
		Mean	St. Dev.	Mean	St. Dev.
GC	360	0.7722	0.4200	0.1522	0.2252
GE	360	0.8639	0.3434	0.2028	0.3034
LC	400	0.5450	0.4986	0.3473	0.3303
LE	400	0.9300	0.2555	0.3243	0.3726
All	1,520	0.7757	0.4173	0.2608	0.3248

Note: DOM=1 if subject's choice coincides with GG's prediction, 0 otherwise.
Error=1 if agency made an error, 0 otherwise.

Table 5.7: Summary Statistics. Dominance and Errors.

For hypothesis testing, it would be useful to aggregate data in two different ways, depending on the information available to the relevant actor. Thus, for hypotheses related to the decisions of the taxpayers (OS and RF), data are aggregated by signal (columns 3-5 in table 5.8). For those related to actions of the agency (SCR), on the other hand, the aggregation is done according to the type of agency (columns 6-8 in the same table).

Treatment	Observations	Signal (s)			Agency's type (q)		
		a	b	c	A	B	C
GC	360	7	295	58	18	234	108
GE	360	29	292	39	54	234	72
LC	400	29	330	41	60	260	80
LE	400	51	337	12	100	280	20
All	1,520	116	1,254	150	232	1,008	280

Note: The agency can be soft, medium or tough on evasion ($q = A$, B or C resp.). Signals can be soft, medium or tough ($s = a$, b or c resp.).

Table 5.8: Number of observations, aggregated by signal and type of agency.

For the analysis, data from all subjects for all periods were pooled. This is justified by the fact that there is little variability in behaviour after the first few rounds of each treatment,[17] with many people choosing exactly the same option every time they receive a given signal. This lack of variability over time is not a bad thing in itself (since the theory actually predicts such rigidity), but it precludes the possibility of using other econometric techniques (e.g., panel data).

5.3.1 OS and RF hypotheses

The set of variables that is going to be used for testing is described in table 5.9.

[17]Except in the GE one, that requires 10 rounds to become stable. This, however, does not usually have an impact on results, and when it does, it will be mentioned in the text.

Variable	Role	Type	Description
DOM	Dependent	Dummy	1 if choice coincides with prediction, 0 otherwise
DOMs	Dependent	Dummy	Idem DOM, but for a fixed $s \in S$
RDOM	Dependent	Dummy	Idem DOM, but for $s \in \{a, c\}$
ADOM	Dependent	Dummy	Idem DOM, but for $s = b$
g	Explanatory	Dummy	1 if G treatment, 0 otherwise
e	Explanatory	Dummy	1 if E treatment, 0 otherwise
ge	Explanatory	Dummy	Interaction term: 1 if GE treatment, 0 otherwise
a	Explanatory	Dummy	1 if $s = a$, 0 otherwise
b	Explanatory	Dummy	1 if $s = b$, 0 otherwise
c	Explanatory	Dummy	1 if $s = c$, 0 otherwise

Note: "Predictions" as defined in hypothesis 5.1.

Table 5.9: Variables of the model. Dominance.

DOMs measures Dominance when only observations with a given signal s are considered. RDOM means Rudimentary Dominance and considers only observations when signals are soft (a) or tough (c). ADOM measures Advanced Dominance and only takes into account observations with medium signals (hence, it is identical to DOMb). The model used for testing is then

$$DOM = \alpha + \beta_1 g + \beta_2 e + \beta_3 ge + \gamma_1 a + \gamma_2 b + \gamma_3 c + \varepsilon \tag{5.17}$$

(analogous ones are used for the alternative dependent variables) and the estimates are shown in table 5.10.

Table 5.11 shows the results of the test in a schematic way.[18] The first panel tests the OS hypothesis (see note below the table for interpretation of symbols). The null hypothesis is that data are consistent with GG predictions,[19] a hypothesis that is supported in the cases of low and high signals ($s = a$ or c) and that implies that people are, at least, Rudimentary.[20] When the signal is medium, however, the GG's predictions are rejected for all treatments and, therefore, the OS hypothesis is statistically rejected as well (i.e., those aspects related to part 2 of the hypothesis). In terms of the sign of the relationship, however, the results do support the predictions, as can be seen in figures 5.3 and 5.4, where the observed strategies resemble the shape of the predicted ones (except for LC,

[18]The tests are shown in table C.1 in appendix C.3.

[19]The predicted value (following proposition 5.2, AD) is 1 for all cases, which is interpreted as requiring that all observations should match predictions.

[20]The null hypothesis is rejected in the GE case because of an outlier. If eliminated, the hypothesis cannot be rejected.

Dep. Var.:	DOM*a*	DOM*b*	DOM*c*	RDOM	ADOM	DOM
a				1.0205		0.7091
				[0]		[0]
b						0.5030
						[0]
c				0.9855		0.7671
				[0]		[0]
g	0.0000	0.2803	-0.0345	-0.0201	0.2803	0.2227
	[0.082]	[0]	[0.158]	[0.35]	[0]	[0]
e	-0.0196	0.4714	0.0000	-0.0297	0.4714	0.3928
	[0.323]	[0]	[0.706]	[0.072]	[0]	[0]
ge	0.0196	-0.3543	-0.0681	-0.0095	-0.3543	-0.2998
	[0.323]	[0]	[0.217]	[0.804]	[0]	[0]
cons	1.0000	0.4485	1.0000		0.4485	
	[.]	[0]	[0]		[0]	
Obs.	116	1,254	150	266	1,254	1,520
LC	1.0000	0.4485	1.0000	1.0000	0.4485	.5450
LE	0.9804	0.9199	1.0000	0.9841	0.9199	.9300
GC	1.0000	0.7288	0.9655	0.9692	0.7288	.7722
GE	1.0000	0.8459	0.8974	0.9412	0.8459	.8639

Note: Top panel: Probability that estimate =0 is shown in brackets below estimate. Bottom panel displays estimated average values of the dependent variable for each treatment.

Table 5.10: Estimation. Dominance. Overall and by signal.

Dep. Var.:	DOM*a*	DOM*b*	DOM*c*	RDOM	ADOM	DOM
LC		X			X	X
LE		X			X	X
GC		X			X	X
GE		X	X	X	X	X
LC=GC		GC			GC	GC
LE=GE	GE	LE	LE		LE	LE
LC=LE	LC	LE			LE	LE
GC=GE		GE			GE	GE

Note: Top panel: Empty if data fits prediction in hypothesis OS; "X" otherwise. Bottom panel: Empty if no difference, treatment with higher dominance otherwise.

Table 5.11: Dominance tests. Predictions and inter-treatment comparisons.

that will be analysed in detail later).[21] Having in mind the discreteness of the model (which amplifies divergences) and that the parameters were chosen to make the test as difficult to pass as possible for the GG theory, the result is still encouraging.

Result 5.1 (QAD) ***Qualitative Advanced Dominance:*** *People are, at least, Rudimentary: they act as predicted by the Rudimentary Dominance proposition (5.1) when signals are low or high. The hypothesis that they make decisions in a way consistent with the Advanced Dominance predictions (proposition 5.2) is statistically rejected (and so is the OS hypothesis, consequently) but supported qualitatively (i.e., the relationship exhibits the expected sign).*

The bottom panel of table 5.11 compares the levels of Dominance across treatments. The null hypothesis for the first two lines is that Dominance is the same in Global and Lottery treatments, a hypothesis that (following the general result) is supported for *RDom* but not for *ADom*. On the other hand, the result that the G/L comparison depends on whether E or C is played is something that the theory cannot explain (there should be no difference, theoretically). It is important to mention, though, that the difference between GE and LE is drastically reduced if considering only the last 10 periods of the session (see figure 5.6, as well as tables C.3, C.4 and C.5 in appendix C.3), so it could be said that Global treatments foster more Dominance than Lottery ones.

For the last two lines, the null hypothesis is that Dominance is the same in Evasion and Compliance treatments. Once again, *RDom* is satisfied but *ADom* is not, but in the latter case the results are clear now: E treatments are more consistent with predictions than C ones. This can be explained by the coincidence of the risk- and payoff-dominant equilibria in the former ones and the discrepancy between them in the latter ones.[22] This is therefore consistent with the RF hypothesis.

Result 5.2 (R/PDE) ***Risk- v Payoff-Dominant Equilibrium:*** *Choices in E treatments are consistent with predictions more frequently than in C ones.*

These results can also be visualised in figures 5.5 and 5.6. The first one confirms that *RDom* is strongly supported by data and that different treatments do not affect it. The second one focuses on choices when the signal is medium and attests that *ADom* predictions are statistically rejected, though qualitatively supported. It also shows that

[21] In the figures, 1 corresponds to Evasion ($d = Y$) and 0 to Compliance ($d = Z$).

[22] Actually, this only applies to the Global treatments, since clearly there is no coordination game in Lottery ones. There is no similar explanation for the difference between E and C in the Lottery cases.

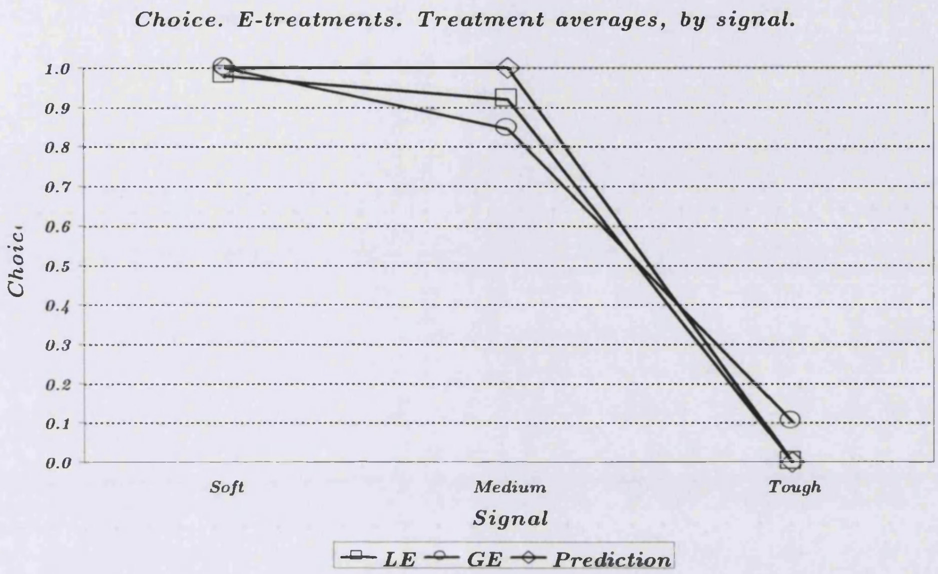

Figure 5.3: Observed and Predicted choices. E-treatments.

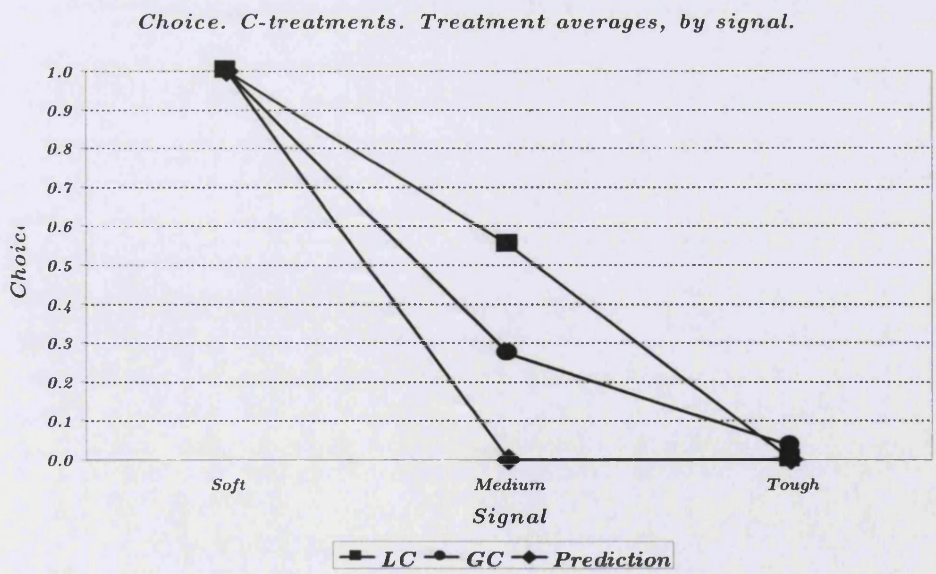

Figure 5.4: Observed and Predicted choices. C-treatments.

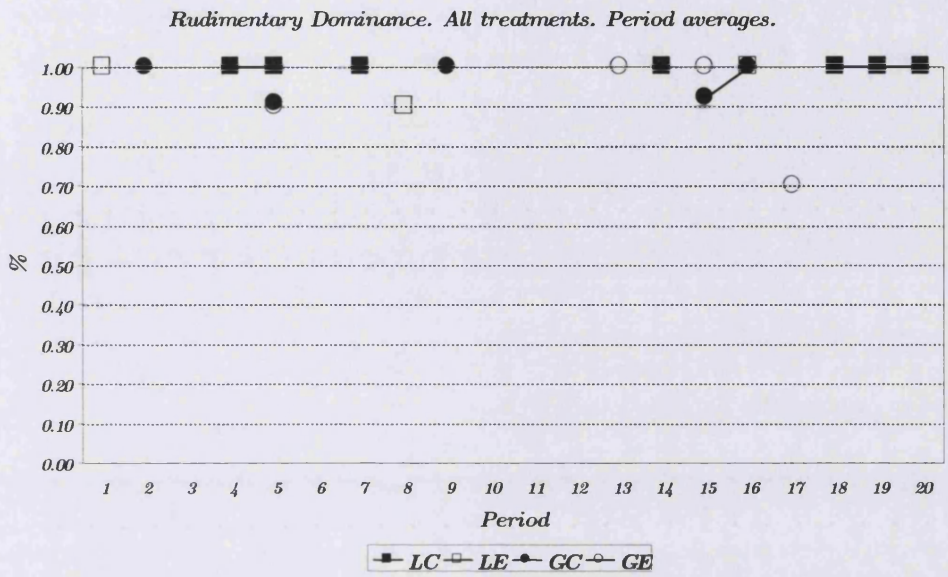

Figure 5.5: Rudimentary Dominance. All treatments. Period averages

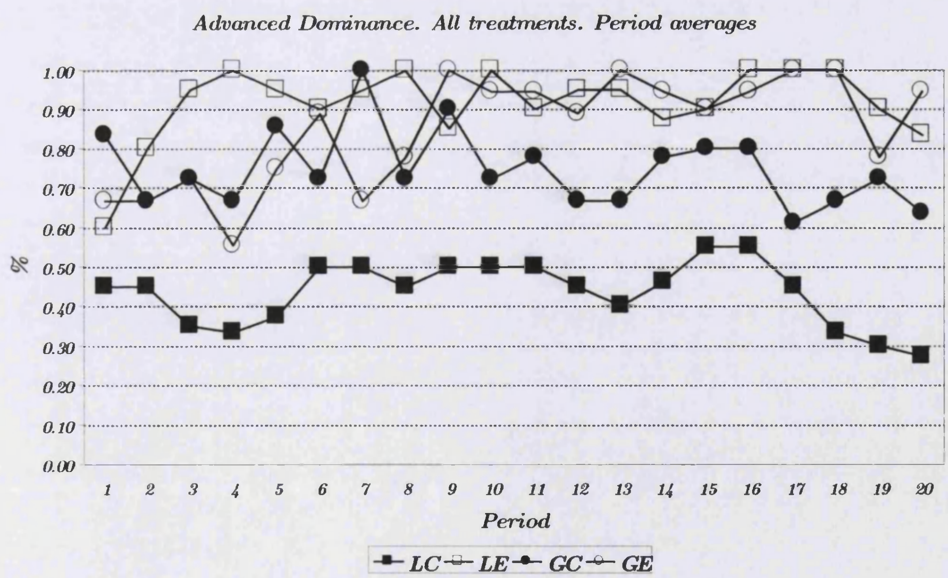

Figure 5.6: Advanced Dominance. All treatments. Period averages.

Variable	Role	Type	Description
ADOM	Dependent	Dummy	1 if data fits prop. AD (part 2), 0 otherwise
g	Explanatory	Dummy	1 if G treatment, 0 otherwise
e	Explanatory	Dummy	1 if E treatment, 0 otherwise
ge	Explanatory	Dummy	Interaction term: 1 if GE treatment, 0 otherwise
gender	Explanatory	Dummy	1 if female, 0 otherwise
age	Explanatory	Natural	
study	Explanatory	Dummy	0 : no study, 1 : non-economics, 0 : economics
# exp	Explanatory	Dummy	0 : none, 1 : 1 to 4, 2 : 5+ experiments
math	Explanatory	Dummy	0 : none, 1 : basic, 2 : advanced knowledge
prob	Explanatory	Dummy	0 : none, 1 : basic, 2 : advanced knowledge
game	Explanatory	Dummy	0 : none, 1 : basic, 2 : advanced knowledge

Note: "Study" refers to "area of study". "Math"/"Prob"/"Game" refer to knowledge of mathematics, probability theory and game theory, respectively.

Table 5.12: Questionnaire variables. Dominance.

treatments can be ranked as determined by the tests, namely, (from higher to lower Dominance), LE, GE, GC and LC.[23]

5.3.2 Characteristics and Decisions

The analysis of choices can be furthered by using the information collected in the questionnaire run after the experimental rounds. The relevant variables are shown in table 5.12.

The analysis will be restricted to that of $ADOM$, that is, to the analysis of Advanced Dominance. The reasons for this are two: first, the previous section proved that $RDom$ is satisfied almost perfectly for the whole sample of participants, regardless of their individual characteristics; and second, $ADOM$ explains most of the variability of the overall dominance (DOM) because in most observations the signal is medium (see table 5.8).

The questionnaire also asked participants about the strategies they followed and the rationale behind them. This information was then used to classify them according to some stylised characteristics, in a fashion similar to the one used by Bosch-Domenech et al. (2002). The distribution of subjects in terms of categories and treatments is shown in table 5.13.

[23]Restricting attention to the last 10 periods so that the learning process in GE converges, the difference between GE and LE vanishes. This is consistent with the RDE/PDE argument (hypothesis 5.3), since it seems that people learn to play the only "reasonable" equilibrium.

Category	GC	GE	LC	LE	All
Expected payoff maximisers (EPM)	10/11	8/11	5	5/13	28/40
Chance maximisers (CM)	1/2	0/3	6/7	0/8	7/20
Learners (L)	0	3	1	1	5
Mixers/Experimenters (M/E)	1	2	0	2	5
Non-independent (NI)	1	0	4	3	8
Randomisers (R)	1	2	1	0	4
Confused (C)	1	0	1/2	1	3/4
Risk-lovers (RL)	2	0	1	0	3
All	18	18	20	20	76

Note: Cells with two numbers separated by "/" reflect uncertainty about the allocation of some subjects to specific categories.

Table 5.13: Questionnaire. Classification of subjects.

The different categories are defined as follows:[24]

- Expected payoff maximisers (EPM): Those who indicated they played either Y in E treatments or Z in C ones, based on expected-payoff maximisation. Note that this category includes everyone who played according to the OS strategy, even though they did not use higher-order beliefs.
- Chance maximisers (CM): Those who only considered the probabilities of outcomes being higher or lower than the safe option, without weighting them using the associated payoffs.[25]
- Learners (L): Those whose decisions varied in the first periods, but chose always the same action afterwards.

[24] Appendix C.4 shows comments from some subjects' questionnaires that are characteristic of each one of these categories.

[25] To see the difference between an EPM and a CM, consider the decisions in the two following cases: In the LE treatment the "safe" payoff is 140 and so, with probability 1/8 the subject gains (receives payoff 1000), with probability 3/4 she gains (payoff 145) and with probability 1/8 she loses (payoff 1). A CM finds that if she evades (chooses Y) the number of "gain" scenarios is greater than the number of "loss" ones, and hence she evades. An EPM computes the expected payoff of evasion 233.875, compares it to the safe payoff 140, and chooses to evade. Thus, in this case both CMs and EPMs would choose the same option and one cannot, based solely on their choices, classify them into one or the other category.

In the LC treatment the safe payoff is 654 and so, with probability 1/8 the subject gains (receives payoff 715), with probability 3/4 she gains (payoff 655) and with probability 1/8 she loses (payoff 579). CMs still prefer evasion over compliance since the number of "gain" cases is greater than that of "loss" cases. EPMs find that the expected payoff of evasion is 653, which is lower than the safe payoff 654, and so they comply. In this case, therefore, choices can distinguish those who belong to one category from those who belong to the other one.

- Mixers/Experimenters (M/E): Those that deviated just once or twice from the predictions of the OS hypothesis but, unlike the Learners, did so at times other than the first periods (Experimenters). An alternative rationale could be that they followed a strategy such that they evaded and complied with probabilities that usually replicated the relevant odds ($(1/8, 7/8)$ in C treatments and $(7/8, 1/8)$ in E ones), and so could be labelled "Mixers".

- Non-independent (NI): Those who (despite the instructions clearly stating that rounds were independent from each other) followed some kind of history-dependent strategy.

- Randomisers (R): Those who chose randomly between Y and Z. Also called "Guessers" (G).

- Confused (C): Those who seemed to be (or acknowledge they were) confused.

On top of these strategies, the degree of risk aversion is expected to play a role as well. In particular, risk aversion fosters compliance (*ceteris paribus*) and hence makes Global Game's predictions easier to be satisfied in *Compliance* treatments, but works against them in *Evasion* ones. Combining the strategies defined above and the degree of risk aversion, one can usually categorise all subjects and find some interesting stylised facts.

The first stylised fact is that categories seem to order themselves in three "Dominance bands" according to their degree of coincidence with the GG predictions (see figures 5.7, 5.8, 5.9 and 5.10). Near the top we can find the $EPMs$ (high dominance). In the middle-ground there is a mixed bag of types (M/E, L, NI and C) who chose different actions in different periods, even though they always got the same signal b. Risk lovers (RL) are close to the top in E treatments and to the bottom in C treatments, and the opposite is true for risk averse (RA) people.

All these results, however, are not surprising. The category that is really exciting to analyse in detail, on the other hand, is that of the Chance Maximisers, since it is behind the case with the largest deviations from predictions (the GC treatment). Now, the first thing to notice is that in some cases CMs cannot be distinguished from $EPMs$, because the observed data are consistent with the predictions of both criteria (expected-payoff and probability maximisation) and the questionnaire information is vague (this is the rationale for the ambiguity in table 5.13). For this very reason, the most interesting scenarios are those where the two criteria prescribe different actions, as is the case in C treatments (global game theory predicts Compliance, chance maximisation predicts Evasion). Focusing on these treatments, it can be seen that significant deviations from

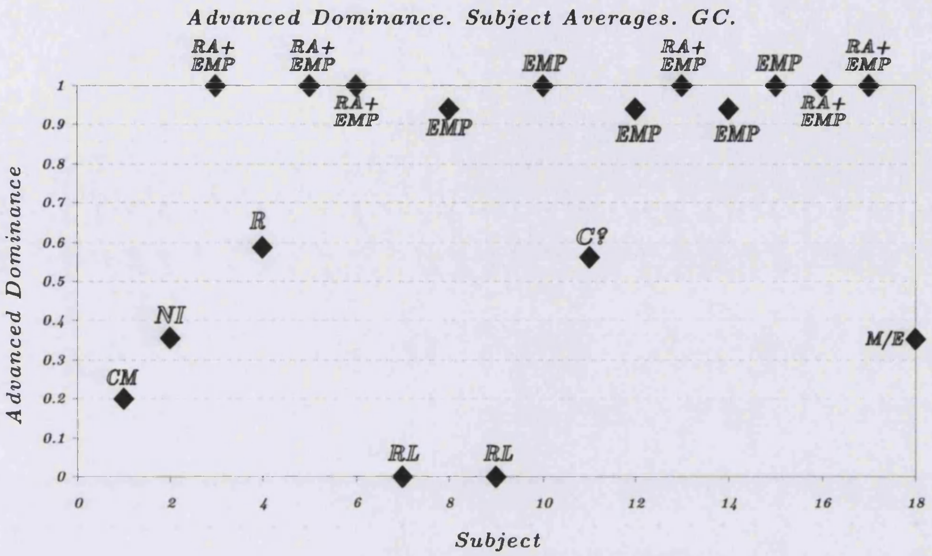

Figure 5.7: Advanced Dominance. Subject averages. GC treatment.

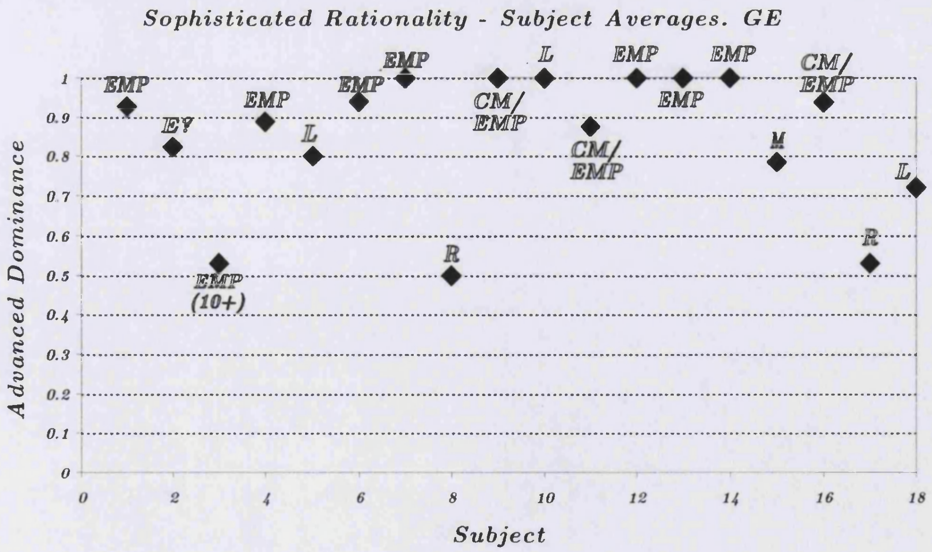

Figure 5.8: Advanced Dominance. Subject averages. GE treatment.

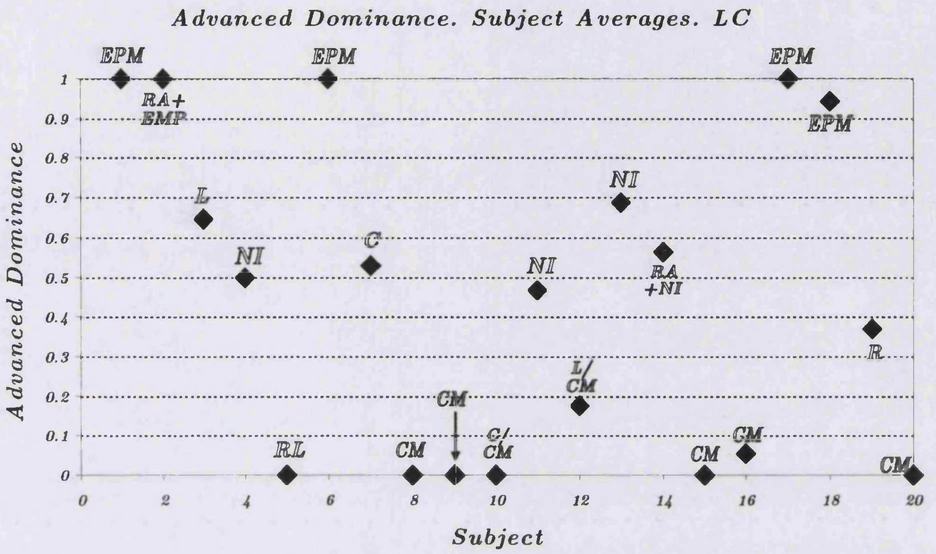

Figure 5.9: Advanced Dominance. Subject averages. LC treatment.

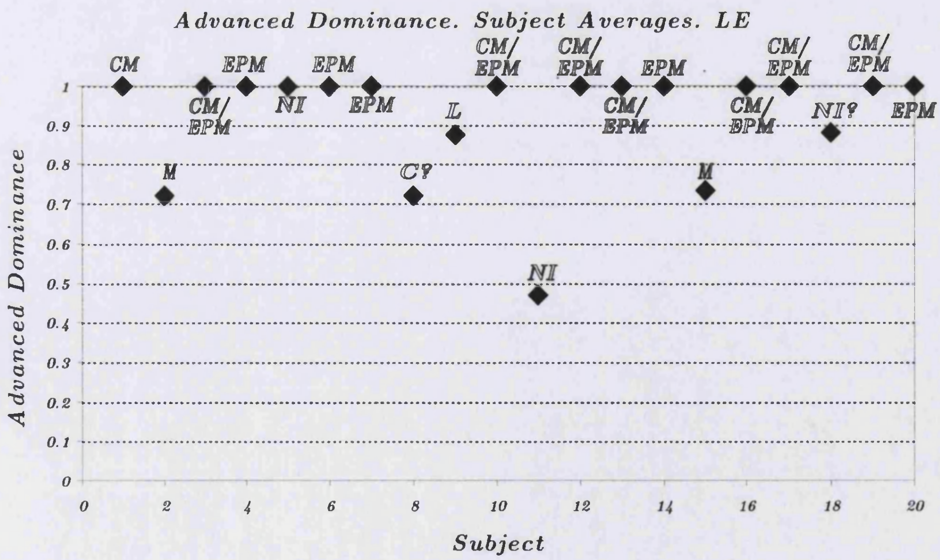

Figure 5.10: Advanced Dominance. Subject averages. LE treatment.

the GG predictions take place, thus confirming the results of the tests that compare the levels of dominance in C and E treatments (table 5.11). Also, since CM's prescription to evade depends on what the other person does in GC but not in LC, it is not surprising that the degree of dominance in the former is greater than in the latter: the uncertainty about the other person's action in GC works against the incentives to evade and (as seen in figure 5.7) only risk loving people end up evading in all periods. Since this interdependence does not play a role in LC, the number of subjects that evade in all periods is far greater (see figure 5.9), and explains the huge divergence between predictions and data (and confirms the ranking of treatments according to Dominance found on page 82).

The stylised facts shown so far give us a snapshot of the data, but the question that remains to be answered is: what is behind these choices? What (if any) are the personal characteristics that drive them? To answer these questions, the variables defined in table 5.12 were used to estimate the following model:

$$\begin{aligned} ADOM = {} & \alpha + \beta_1 g + \beta_2 e + \beta_3 ge + \\ & + \gamma_1 gender + \gamma_2 age + \gamma_3 \# \exp + \gamma_4 math + \gamma_5 prob + \gamma_6 game + \varepsilon \quad (5.18) \end{aligned}$$

The results (shown in table 5.14) indicate that estimates are robust to the specification of the model (last three columns)[26] and usually there is not much difference between treatments or between individual treatments and the whole sample. The analysis finds that being male, young, bad-at-maths and good-at-game-theory makes a subject more likely to make decisions that coincide with the global game predictions. There is no rationale for the gender effect (which, apart from the whole sample, is only significant in one treatment, anyway), though it is important to note that a similar result is found by Heinemann et al. (2004a). The age effect may seem to reflect that most subjects are university students, but actually it is driven by a few older outliers: if the analysis restricts its attention to "up-to-25-year-olds", age becomes non-significant (see table C.6 in appendix C.3). The negative sign of the mathematics coefficient is surprising and may be the result of people's mis-estimation of their mathematical knowledge. Area of study is not significant and, surprisingly, knowledge of probability theory or participation in other experiments are not significant either (though Heinemann et al. (2004b) find the same result regarding experience).

The only robustly significant variable seems to be knowledge of game theory, which increases Dominance. Furthermore, it is significant in both treatments in which strategic (i.e., game theoretic) interactions took place. This may indicate that some degree of in-

[26] For this very reason, only OLS estimates are shown throughout the whole paper.

doctrination may have played a role and so that training can breed "sophistication". This suggests that a typical population (where average knowledge of game theory is expected to be negligible) would make choices quite different from the ones suggested by the GG theory. However, if one considers that firms are sophisticated, then the theory should be a good predictor of their behaviour. Moreover, a similar result could be achieved if individual taxpayers had access to sophisticated professional advice, something that is indeed likely to occur.

	OLS					Probit	Logit
	GC	GE	LC	LE	All	All	All
g					0.2914	0.8573	1.3932
					[0]	[0]	[0]
e					0.4895	1.7349	3.0951
					[0]	[0]	[0]
ge					-0.3894	-1.4149	2.4897
					[0]	[0]	[0]
gender	-0.0306	0.0078	0.0937	-0.0738	-0.0616	-0.2761	0.4752
	[0.745]	[0.853]	[0.31]	[0.005]	[0.011]	[0.003]	[0.004]
age	-0.0385	-0.0282	-0.0251	0.0039	-0.0078	-0.0304	0.0540
	[0]	[0]	[0]	[0]	[0]	[0]	[0]
study	0.1916	0.0349	-0.5397	-0.0149	-0.0306	-0.1446	0.2348
	[0.004]	[0.514]	[0]	[0.787]	[0.369]	[0.312]	[0.382]
#exp	0.0006	0.1402	-0.0677	-0.0160	-0.0059	-0.0416	0.0729
	[0.988]	[0]	[0.161]	[0.575]	[0.738]	[0.513]	[0.507]
maths	-0.5119	0.0418	0.1660	-0.0788	-0.0993	-0.3749	0.6982
	[0]	[0.432]	[0.196]	[0.099]	[0.002]	[0.001]	[0.002]
prob	-0.0454	-0.0060	0.0033	0.1204	0.0047	0.0219	0.0314
	[0.635]	[0.893]	[0.963]	[0.007]	[0.866]	[0.834]	[0.867]
game	0.3409	0.1840	0.0362	0.0242	0.0961	0.4249	0.6918
	[0]	[0]	[0.464]	[0.193]	[0]	[0]	[0]
cons	2.0118	1.2868	1.3768	0.8555	0.7772	1.1536	2.1058
	[0]	[0]	[0]	[0]	[0]	[0]	[0]
Obs.	295	292	330	337	1,254	1,254	1,254

Note: Probability that estimate =0 is shown in brackets below estimate.

Table 5.14: Estimation. Effect of personal characteristics on choices.

5.3.3 SCR hypothesis

The key prediction of the TEGG model is that an agency would be advised to use the contingent auditing rule and to discard the cut-off one. This means that the agency would make fewer targeting errors if implementing the former than if using the latter, given that enforcement costs are the same in both cases. These are the Zeal and Negligence errors defined on page 39, though –for the reasons explained in footnote 14– the analysis will focus on the Negligence errors only.

The expected loss of a 2-person group can then be expressed as

$$EL\left(d,d',q\right) = \sum_{y\in\mathcal{Y}} \Pr\left(y\right) \sum_{s\in\mathcal{S}} \Pr\left(s|q\right) \sum_{s'\in\mathcal{S}} \Pr\left(s'|q\right) \sum_{d\in\mathcal{D}} \Pr\left(d|s,y\right) \sum_{d'\in\mathcal{D}} \Pr\left(d'|s',y\right)$$
$$\sum_{a\in\mathcal{A}} \Pr\left(a|d,d',q\right) \sum_{a'\in\mathcal{A}} \Pr\left(a'|d,d',q\right) \left[\left(1-a\right)\left(1-d\right)y + \left(1-a'\right)\left(1-d'\right)y\right] \quad (5.19)$$

where the expression in square brackets is the sum of negligence errors for the 2-person group. Armed with this information, the model to be estimated is therefore

$$ERR = \beta_1 g + \beta_2 e + \beta_3 ge + \gamma_1 A + \gamma_2 B + \gamma_3 C + \varepsilon \quad (5.20)$$

where the variables are defined as in table 5.15.

Variable	Role	Type	Description
ERR	Dependent	Dummy	1 if an error was made, 0 otherwise
ERRq	Dependent	Dummy	Idem ERR, but for a fixed $q \in Q$
g	Explanatory	Dummy	1 if G treatment, 0 otherwise
e	Explanatory	Dummy	1 if E treatment, 0 otherwise
ge	Explanatory	Dummy	Interaction term: 1 if GE treatment, 0 otherwise
A	Explanatory	Dummy	1 if $q = A$, 0 otherwise
B	Explanatory	Dummy	1 if $q = B$, 0 otherwise
C	Explanatory	Dummy	1 if $q = C$, 0 otherwise

Note: ERR measures errors per person in a 2-person group.

Table 5.15: Variables of the model. Errors.

The estimates can be seen in table 5.16.

Dep. Var.:	ERRA	ERRB	ERRC	ERR
A				0.8610
				[0]
B				0.2886
				[0]
C				0.1526
				[0]
g	0.0059	-0.1610	-0.1847	-0.1242
	[0.956]	[0]	[0]	[0]
e	0.3718	-0.2130	-0.1950	-0.1007
	[0]	[0]	[0]	[0]
ge	-0.0967	0.1466	0.1856	0.0804
	[0.423]	[0]	[0]	[0]
cons	0.5482	0.3477	0.1954	
	[0]	[0]	[0]	
Obs.	232	1,008	280	1,520
LC	0.5482	0.3477	0.1954	0.3473
LE	0.9200	0.1346	0.0004	0.3243
GC	0.5541	0.1866	0.0107	0.1522
GE	0.8293	0.1203	0.0013	0.2028

Note: Top panel: Probability that estimate =0 is shown in brackets below estimate. Bottom panel displays estimated average values of the dependent variable for each treatment.

Table 5.16: Estimation. Errors. Overall and by type of agency.

In a fashion similar to the one used in section 5.3.1, several tests are shown in a schematic form in table 5.17 (the values of the tests can be found in table C.2 in appendix C.3).

The top panel tests the accuracy of predictions and shows that the data do not fit them. In particular, errors are usually higher than predicted in C treatments but lower than predicted in E ones. This is consistent with the Dominance results, which indicate that "too many" people evade when they should comply (C treatments) and comply when they should evade (E treatments), as shown in figures 5.3 and 5.4. The main conclusion, thus, is basically the same as the one found for Dominance in Result 5.1, and subject to the same qualifications.

The first two lines of the bottom panel are the important ones: they show the tests for the SCR hypothesis. Given the minimum variability in the extreme cases (when the agency

Dep. Var.:	ERRA	ERRB	ERRC	ERR
LC	+	+	+	+
LE	-	-		
GC	+	+		+
GE	-	-		-
LC=GC		GC	GC	GC
LE=GE		GE	LE	GE
LC=LE	LC	LE	LE	LE
GC=GE	GC	GE		

Note: Top panel: Empty if data fits predictions; "+" if observed errors are higher than predicted; "-" otherwise. Bottom panel: Empty if no difference, treatment with less errors otherwise.

Table 5.17: Errors tests. Predictions and inter-treatment comparisons.

is too soft, $q = A$, or too tough, $q = C$), the relevant tests are those for the medium one, and this one shows clearly that the *Global* treatments lead to less errors per capita than the *Lottery* ones. In other words, the SCR hypothesis is strongly supported.

Result 5.3 (SCR) ***Superiority of Contingent Rule:*** *From the agency's perspective, the contingent rule is better than the cut-off rule.*

The last two lines test whether there are significant differences between E and C treatments and show (again focusing on the medium case) that the first lead to less errors than the second. Again, this can be linked to the Dominance analysis, where E treatments show a higher degree of coincidence with predictions than C ones. This means, in other words, than in the latter many people evaded when they should have complied, and the higher number of associated errors thus explains the present result.

Finally, it is important to notice that all these findings are also supported graphically, as shown in figures 5.11 to 5.16. It can be clearly seen there that G treatments (i.e., those in which the contingent policy is implemented) lead to (weakly) less errors than L ones (those in which the cut-off one is used). It also shows the (expected) result that errors are a decreasing function of the agency's "toughness", which is consistent with the global game comparative statics.

Result 5.4 (EAT) ***Effect of agency's type:*** *Errors decrease with the agency's "toughness".*

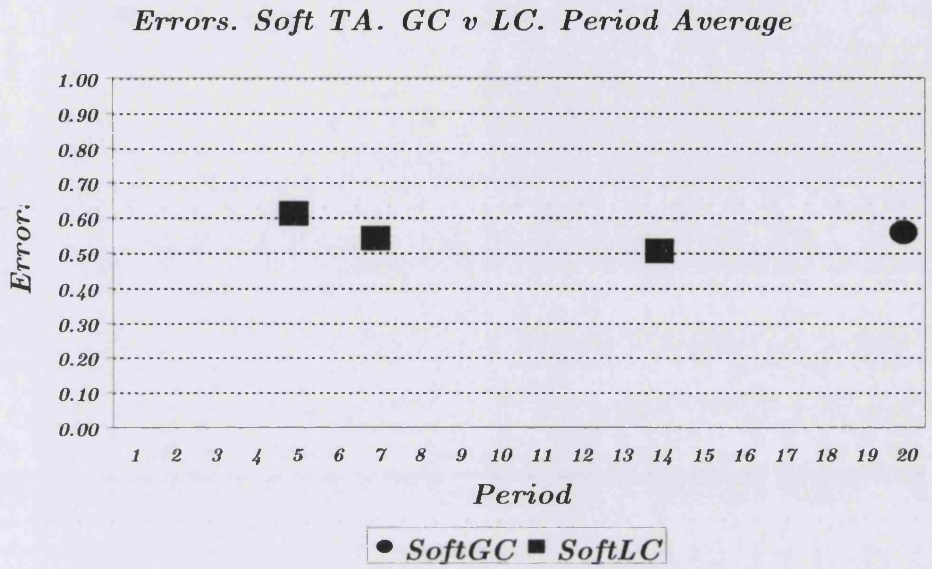

Figure 5.11: Errors. Soft agency. GC v LC.

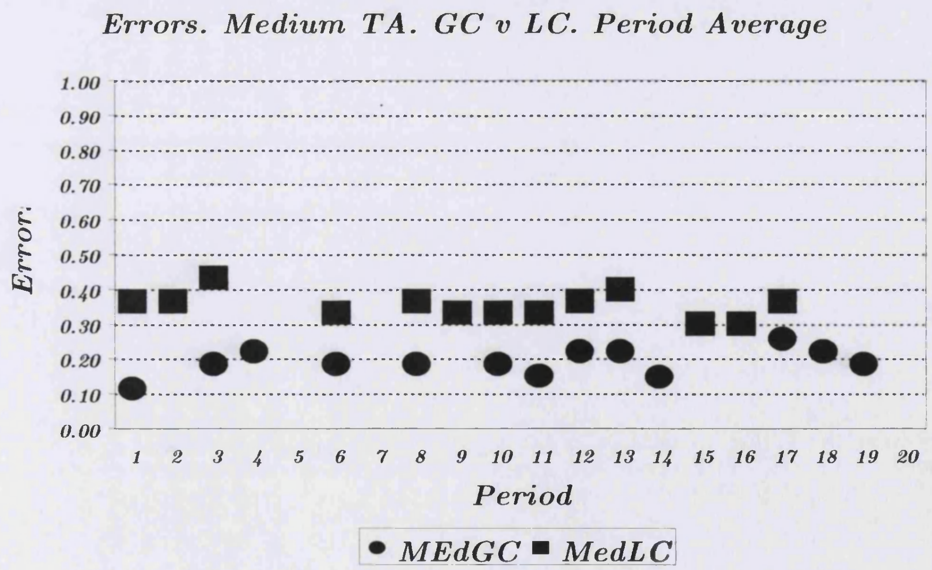

Figure 5.12: Errors. Medium agency. GC v LC.

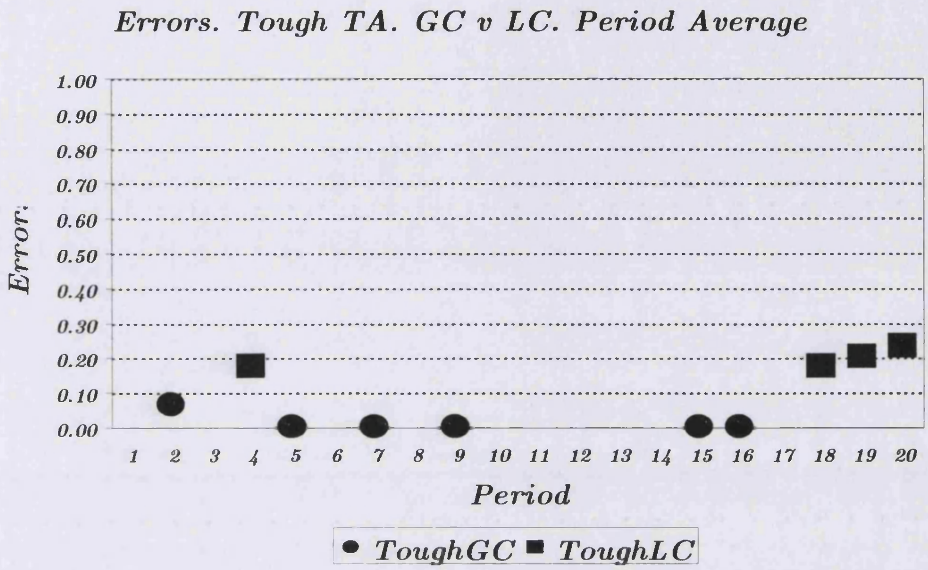

Figure 5.13: Errors. Tough agency. GC v LC.

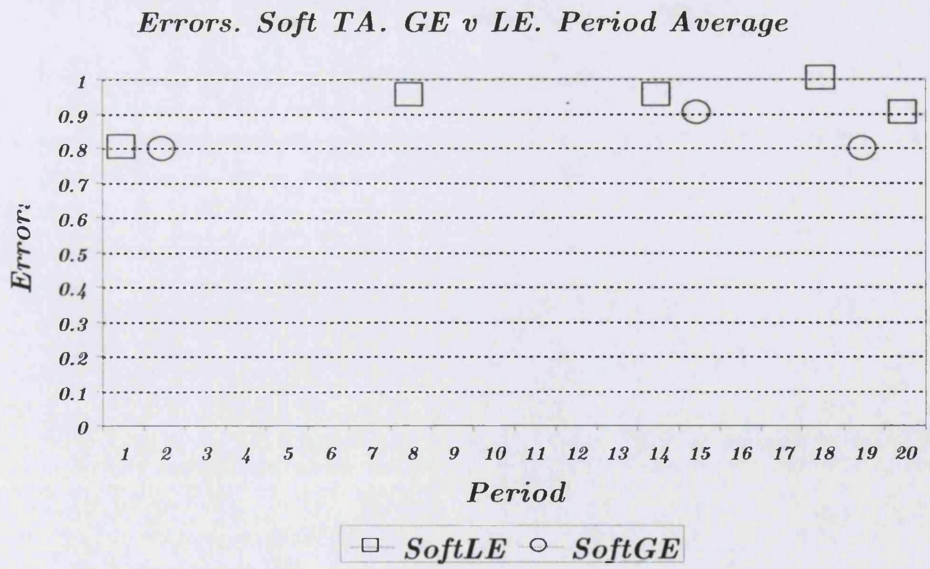

Figure 5.14: Errors. Soft agency. GE v LE.

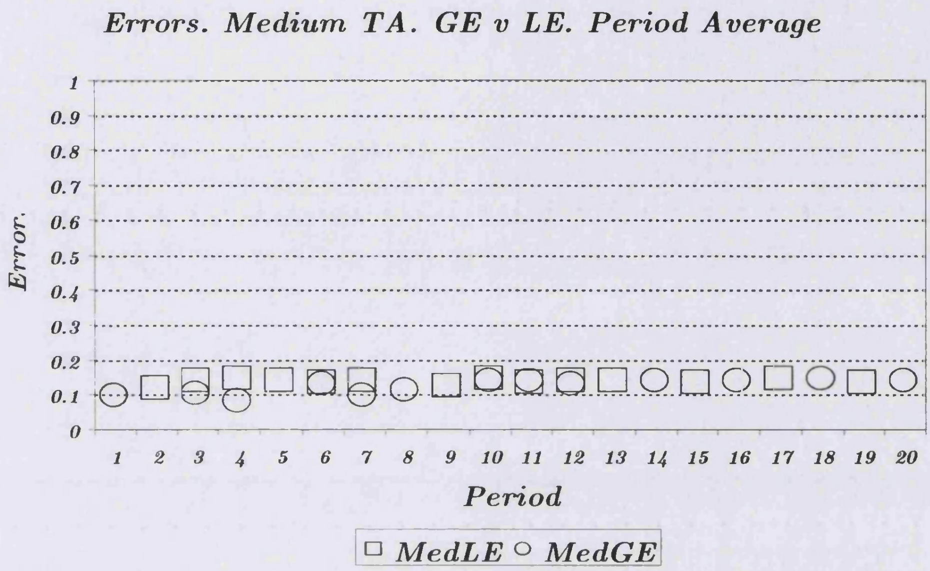

Figure 5.15: Errors. Medium agency. GE v LE.

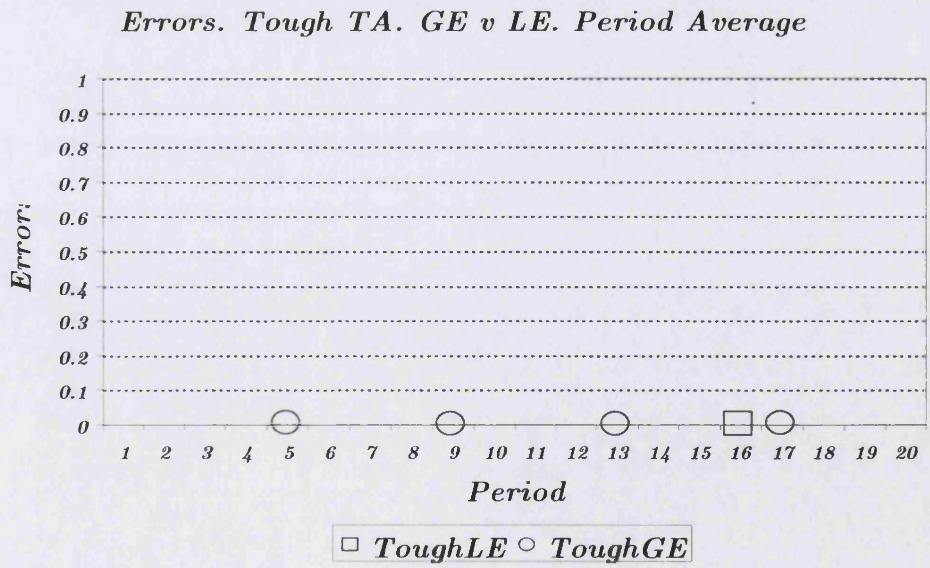

Figure 5.16: Errors. Tough agency. GE v LE.

5.4 Conclusions

The empirical analysis of tax evasion is problematic because of the reluctance of both taxpayers and tax agencies to provide the relevant information. This study, therefore, uses experimental data as a second-best alternative and focuses on the testing of some of the theoretical predictions of the model developed in chapter 4, though the richness of the dataset also allows for the investigation of other interesting hypotheses related to decision-making processes and the global game theory.

Results are strongly supportive of the main prediction of the TEGG model, namely, that a tax agency using a contingent auditing policy would do better than if it used the standard cut-off one. The negative *externality* between taxpayers generated by the contingent policy and the associated strategic uncertainty it creates seem to be the powerful forces behind this result.

Also supported by the data are the predictions derived from the comparative statics of global games: evasion is higher in Evasion treatments than in Compliance ones, errors decrease with the agency's "toughness", and "tougher" signals lead to lower evasion.

The picture, so encouraging in terms of the sign of the relationships, is however radically different when considering it in terms of statistical significance: in general, the numerical predictions of the theory are rejected by the data. This is true for the medium cases (when the signal is medium), but not for the extreme ones though: in the latter, results are as expected and support the idea that people are, at least, "Rudimentary" and (intuitively) understand the concept of dominance in simple scenarios. Medium cases, on the other hand, show that most people do not use higher-order beliefs when making their decisions (not even in this simple experiment, in which only two iterations are needed). In spite of this, many times they do choose the actions predicted by the theory of global games, usually after playing the game a few times. This "learning" result is not so surprising, as it was already hinted by Carlsson and van Damme (1993) and found experimentally by Cabrales et al. (2002). Other factors also seem to affect decisions, like the tension between the risk-dominant and payoff-dominant equilibria, with their predicted effects closely mimicked by the data. More worrying, however, is the apparently pervasive presence of a significant group of people ("chance maximisers") who choose their strategies without taking into account all the available information (in this particular experiment, the payoffs in different scenarios) and that lead to the largest differences between observed and predicted actions (treatment LC). This concern is connected to the main result derived from the analysis of questionnaire data, which suggests that those with knowledge of game theory ("sophisticated" agents) are more likely to play according to predictions than those without that knowledge ("simple" agents). This indicates that game theoretical "in-doctrination" helps subjects to analyse strategic interaction in a quick and standardised way. In terms of policy, this suggests that firms will react to the implementation of the contingent rule almost exactly as predicted by the theory, while the responses of individual taxpayers (who are expected to be less "sophisticated" than firms) will be more erratic (though the difference in the behaviour of the two groups can be greatly reduced if the latter have access to sophisticated professional advice).

The bottom line is, therefore, that though people may not use higher-order beliefs, many times their decisions are indistinguishable from those of people who do use higher-order beliefs. Consequently, predictions are usually supported in terms of the sign of the coefficients (comparative statics and inter-treatment comparisons) but rejected in terms of statistical significance. The latter problem is, however, mitigated by two factors: First, the discreteness of the model can work against it because it amplifies small differences and thus makes the data-predictions matches more difficult (something already highlighted by Heinemann et al. (2002)). Second, the parameters of the model were explicitly chosen to discourage said matches. This may indicate that, since the estimated coefficients have the signs predicted by the theory in these most demanding conditions, the model would be a better predictor in more favourable environments. On the other hand, the present analysis only compared two possible auditing strategies (the "cut-off" and the "contingent" rules), and a proper test of the optimality (or not) of the latter demands further comparisons against other rules. This testing is something that I plan to undertake as part of my future research agenda regarding this topic. In particular, I intend to design experiments such that, for a given situation, the contingent rule and the alternative ones predict different behaviour by the subjects, so that I will be able to determine which of them reflects more accurately the empirical evidence.

Chapter 6

Conclusion

Improved communication and lower transportation costs have practically eliminated the concept of distance in the present, globalised world. Networks have multiplied and got bigger (e.g., Facebook), thus increasing the size of not only a person's direct, first-order circle of acquaintances, but also of second- and higher-order ones: for a sample, just check the number of times a chain email was forwarded before it got to you.

This means that interconnectivity has soared, but with it have also soared the associated external effects: simply think about that chain email once again...

Going back to the pre-globalised world of independent, self-contained "islands" is, however, not an option. But continuing to use pre-globalised incentive schemes in a globalised world does not seem a smart alternative either: they will probably do more harm than good.

Analysing how those mechanisms are affected by the increased connectivity and how to devise new ones that are better adapted to it are therefore the goals of this thesis. Such enterprise, however, is a vast one, and so the present dissertation focused instead on a few but thoroughly studied cases, using the British railway system and the design of anti-evasion auditing policies as leading examples. The first one helped to illustrate how delegation of contracting rights can improve the efficiency of an organisation in which jointly produced output is the only contractible variable. The second one showed that sometimes the optimal incentive scheme requires the creation of an externality by the designer.

The results suggest that the presence or absence of externalities is an important factor to consider when designing incentive schemes. Thus, externalities can radically modify the schemes (like in the British railway example, in which delegation is preferred over centralisation) or become an integral part of them, transforming the nature of the game in the process (like in the anti-evasion case in which a coordination game is created).

Bibliography

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

Alm, J., B. R. Jackson, and M. McKee (1992). Institutional uncertainty and taxpayer compliance. *American Economic Review 82*, 1018–1026.

Alm, J. and M. McKee (2004). Tax compliance as a coordination game. *Journal of Economic Behavior and Organization 54*, 297–312.

Andreoni, J., B. Erard, and J. Feinstein (1998). Tax compliance. *Journal of Economic Literature 36*, 818–860.

Atkeson, A. (2000). Discussion of morris and shin's "rethinking multiple equilibria in macroeconomic modelling".

Baldry, J. C. (1986). Tax evasion is not a gamble. *Economics Letters 22*, 333–335.

Baliga, S. and T. Sjolstrom (1998). Decentralization and collusion. *Journal of Economic Theory 83*, 196–232.

Basseto, M. and C. Phelan (2004). Tax riots. *Federal Reserve Bank of Minneapolis. Mimeo.*

Benjamini, Y. and S. Maital (1985). Optimal tax evasion and optimal tax evasion policy: behavioral aspects. In W. Gaertner and A. Wenig (Eds.), *The Economics of the Shadow Economy.* Berlin, Germany: Springer Verlag.

Bosch-Domenech, A., J. G. Montalvo, R. Nagel, and A. Satorra (2002, December). One, two, (three), infinity,...: Newspaper and lab beauty-contest experiments. *American Economic Review 92*(5), 1687–1701.

Cabrales, A., R. Nagel, and R. Armenter (2002, February). Equilibrium selection through incomplete information in coordination games: An experimental study.

Carlsson, H. and E. van Damme (1993). Global games and equilibrium selection. *Econometrica 61*(5), 989–1018.

Cornes, R. and T. Sandler (1996). *The Theory of Externalities, Public Goods and Club Goods* (second ed.). Cambridge University Press.

Cowell, F. A. (1990). *Cheating the Government.* Cambridge, Massachusetts: The MIT Press.

Cronshaw, M. B. and J. Alm (1995, April). Tax compliance with two sided uncertainty. *Public Finance Quarterly 23*(2), 139–166.

Devereux, M. and M. Pearson (1990). Harmonising corporate taxes in europe. *Fiscal Studies 11*(1), 21–35.

Faure-Grimaud, A., J. Laffont, and D. Martimort (2003). Collusion, delegation and supervision with soft information. *Review of Economic Studies 70*, 253–279.

Feld, L. P. and J.-R. Tyran (2002). Tax evasion and voting: An experimental analysis. *Kyklos 55*(2), 197–222.

Felli, L. and R. Hortala-Vallve (2007). Preventing collusion through discretion.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Fortin, B., G. Lacroix, and M.-C. Villeval (2004). Tax evasion and social interactions. *CIRPEE Working Paper 04-32.*

Hardin, G. (1968, December). The tragedy of the commons. *Science 162*, 1243–1248.

Harsanyi, J. C. and R. Selten (1988). A general theory of equilibrium selection in games. *Cambridge, MA. MIT Press.*.

Heinemann, F., R. Nagel, and P. Ockenfels (2002, June). Speculative attacks and financial architecture: Experimental analysis of coordination games with public and private information.

Heinemann, F., R. Nagel, and P. Ockenfels (2004a, December). Measuring strategic uncertainty in coordination games.

Heinemann, F., R. Nagel, and P. Ockenfels (2004b, September). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica 72*(5), 1583–1599.

Holmstrom, B. (1982, Autumn). Moral hazard in teams. *The Bell Journal of Economics 13*(2), 324–340.

Holmstrom, B. and P. Milgrom (1987, March). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55*(2), 303–328.

Holt, C. and S. K. Laury (2002, December). Risk aversion and incentive effects. *American Economic Review 92*(5), 1644–1655.

Itoh, H. (1991, May). Incentives to help in multi-agent situations. *Econometrica 59*(3), 611–636.

Itoh, H. (1994). Job design, delegation and cooperation: a principal-agent analysis. *European Economic Review 38*, 691–700.

Katz, M. and C. Shapiro (1985, June). Network externalities, competition, and compatibility. *American Economic Review 75*(3), 424–440.

Kim, Y. (2005). Audit misperception, tax compliance, and optimal uncertainty. *Journal of Public Economic Theory 7*(3), 521–541.

Krugman, P. (2005). *Geography and Trade.* Cambridge,MA: MIT Press.

Macho-Stadler, I. and J. D. Jelovac (2002). Comparing organizational structures in health services. *Journal of Economic Behavior and Organization 49*, 501–522.

Macho-Stadler, I. and J. D. Pérez-Castrillo (1998, December). Centralized and decentralized contracts in a moral hazard environment. *Journal of Industrial Economics 46*(4), 489–510.

Melumad, N., D. Mookherjee, and S. Reichelstein (1995, Winter). Hierarchical decentralization of incentive contracts. *RAND Journal of Economics 26*, 654–672.

Morgan, J. (2000). Public goods and lotteries. *Review of Economic Studies 67*, 761–784.

Morris, S. and H. S. Shin (1997, August). Unique equilibrium in a model of self-fulfilling currency attacks. *Mimeo.*

Morris, S. and H. S. Shin (2002a). Global games: Theory and applications.

Morris, S. and H. S. Shin (2002b). Social value of public information. *Mimeo.*

Myles, G. D. and R. A. Naylor (1996). A model of tax evasion with group conformity and social customs. *European Journal of Political Economy 12*, 49–66.

Prendergast, C. (1995, July). A theory of responsibility in organizations. *Journal of Labor Economics 13*(3), 387–400.

Ratto, M., R. Thomas, and D. Ulph (2005). Tax compliance as a social norm and the deterrent effect of investigations.

Reinganum, J. and L. Wilde (1988). A note on enforcement uncertainty and taxpayer compliance. *The Quarterly Journal of Economics 103*(4), 793–798.

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

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

Scotchmer, S. and J. Slemrod (1989). Randomness in tax enforcement. *Journal of Public Economics 38*, 17–32.

Selten, R., A. Sadrieh, and K. Abbink (1999). Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory and Decision 46*, 211–249.

Stahl, D. O. and P. W. Wilson (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization 25*, 309–327.

Stella, P. (1991). An economic analysis of tax amnesties. *Journal of Public Economics (46)3*, 383–400.

Thaler, R. H. and H. M. Shefrin (1981, April). An economic theory of self-control. *Journal of Political Economy 89*(2), 392–406.

Tirole, J. (1986, Autumn). Hierarchies and bureaucracies: On the role of collusion in organizations. *Journal of Law Economics and Organization 2*(2), 181–214.

Appendix A

Exogenous externality

A.1 Limited liability

Since adding limited liability constraints ($w_i(x) \geq 0, \forall x$) greatly complicates the problem, we explore here an alternative that intends to capture the intuition of what could happen if we introduced such constraints. Thus, we compute the maximum value of x for which the contracts offered to the agents are non-positive (i.e. the point where the limited liability constraint bites) for both the Centralised Second Best and Delegated Second Best contracting structures. In figure A.1 we have depicted these thresholds in terms of the parameter γ.

Note that the curve that corresponds to the hierarchic net contract of the first agent ($X(w_1^h(x))$) is always lower than the two remaining curves. That is, the limited liability constraint of Agent 1 under the delegated structure bites at a lower value of realized output x than the one of Agent 2 ($X(w_2^h(x))$). This indicates that Agent 1 is able to ensure that she gets a positive payment more often than Agent 2, a result that is consistent with proposition 3.2 –Agent 1 bears less risk than Agent 2– and with the fact that the Principal is able to induce Agent 1 to exert a high level of effort while transferring her a relatively low level of risk. When comparing the Centralised Second Best ($X(w^{**}(x))$) and Delegated Second Best ($X(w_1^h(x))$, $X(w_1^h(x))$) structures, we can see that for low levels of γ the latter is more affected by the limited liability constraint (in particular, Agent 2's realization of wages is negative for a larger range of output realizations). However, for large values of γ, the limited liability constraint affects the Centralised Second Best structure more than the Delegated Second Best one.

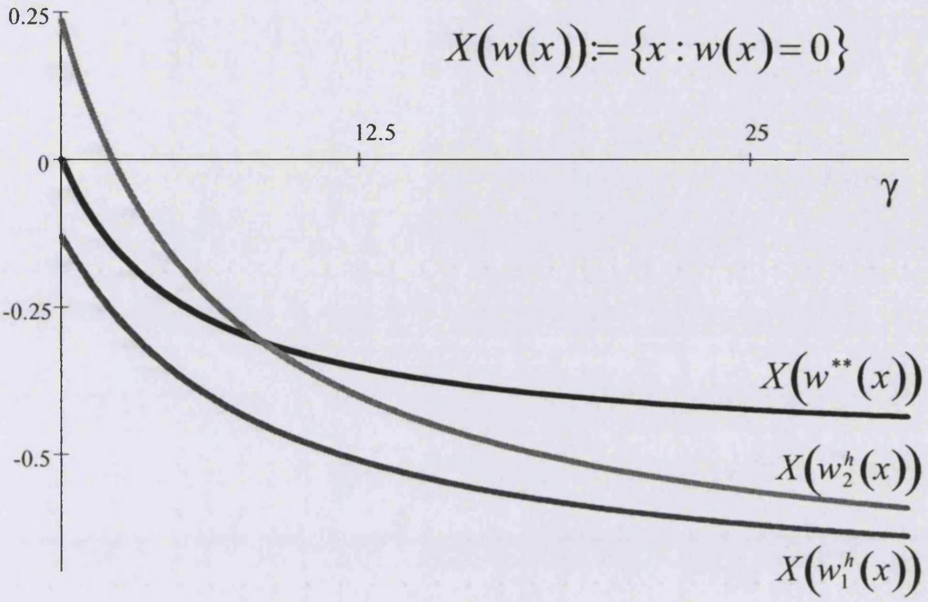

Figure A.1: Limited liability

This result, together with the fact that the relative gains of the delegated structure are strictly increasing in γ (proposition 3.1), suggest that delegated structures should be frequently found when agents are very risk averse and/or the project is very risky.

A.2 Profit-sharing schemes

Empirical evidence shows that many agents are rewarded not in relation to their output, but to profits generated. This is the standard practice, for example, in the case of CEOs. From a theoretical point of view, allowing for profit-based incentive schemes does not represent much of a difficulty, as profits are verifiable by court if disagreement between the parties arises. How would these compensation schemes affect the results?

Let us keep the assumption of linear contracts, but now instead of linear in output, they are assumed to be linear in profits. Hence, the wage schedule faced by agent i is given by

$$w_i(x) = a_i + c_i \pi(x) \qquad \forall i \in \{1, 2\} \tag{A.1}$$

Profits are defined as output minus the total wage bill

$$\pi(x) := x - \sum_{i=1}^{2} w_i(x) \tag{A.2}$$

Plugging equation $A.1$ into equation A.2 and rearranging the terms, we get profits as a linear function of output:

$$\pi(x) = \frac{1}{1 + \sum_{i=1}^{2} c_i} x - \frac{\sum_{i=1}^{2} a_i}{1 + \sum_{i=1}^{2} c_i} \tag{A.3}$$

Plugging now equation $A.3$ into equation $A.1$, we can express the wage schedule of agent i as a linear function of output only:

$$w_i(x) = \left[a_i - c_i \frac{\sum_{i=1}^{2} a_i}{1+\sum_{i=1}^{2} c_i}\right] + \frac{c_i}{1+\sum_{i=1}^{2} c_i} x \qquad \forall i \in \{1, 2\} \tag{A.4}$$

Defining

$$\alpha_i = a_i - c_i \frac{\sum_{i=1}^{2} a_i}{1 + \sum_{i=1}^{2} c_i} \tag{A.5}$$

$$\beta_i = \frac{c_i}{1 + \sum_{i=1}^{2} c_i} \tag{A.6}$$

the wage schedule can be further simplified to

$$w_i(x) = \alpha_i + \beta_i x \qquad \forall i \in \{1, 2\} \tag{A.7}$$

Hence, it is clear that a linear contract contingent only on profits can be rewritten as another linear contract, this one contingent only on output. This means that our analysis can be applied to situations where the incentive schemes are based on profits rather than on output if a simple change of variables is undertaken.[1]

[1] In fact, any linear contract contingent on both profits and output can be rewritten as another linear contract contingent only on output.

That is, every contract of the form

$$w_i(x) = a_i + b_i x + c_i \pi(x) \qquad \forall i \in \{1, 2\} \tag{A.8}$$

can be analysed as a contract contingent only on output (by undertaking the appropriate change of variables).

A.3 Cartel

When the agents make their decisions jointly, the Principal's programme becomes

$$\max_{\{w_1(x),e_1,w_2(x),e_2\}} E_x\{\pi(x)\}$$
$$\text{s.t.}\begin{cases} PC^c: EU^c := \sum_{i=1}^{2} E_x\{U(w_i(x),e_i)\} \geq \underline{U} \\ IC^c: (e_1,e_1) \in \arg\max_{\dot{e}_1,\dot{e}_2} EU^c = \sum_{i=1}^{2} E_x\{U(w_i(x),\dot{e}_i)\} \end{cases} \tag{A.9}$$

where the superindex c labels the Cartel variables, such as the expected utility of the cartel, EU^c, which is defined as the sum of the agents' expected utilities.[2] The objective function of the Principal remains unchanged:

$$E\pi = E_x\{\pi(x)\} = E_x\{x - W(x)\} \tag{A.10}$$

The incentive compatibility constraint yields the result that both agents choose the same level of effort, B:

$$e_1^c = e_2^c = B \tag{A.11}$$

which means that both agents fully internalise the externality they generate on the other.

Given the linearity of the incentive scheme, the participation constraint will be binding, and so the expected total wage bill $E_x\{W(x)\}$ will cover exactly the Cartel's disutility of effort and risk. Using the certainty equivalent defined on page 22, the expected total bill is given by the following expression

$$E_x\{W(x)\} = \frac{\gamma}{2}(B-b)^2 + \frac{1}{2}e_1^2 + \frac{\gamma}{2}b^2 + \frac{1}{2}e_2^2 \tag{A.12}$$

which, using equation A.11, simplifies to

$$E_x\{W(x)\} = \frac{\gamma}{2}(B-b)^2 + \frac{\gamma}{2}b^2 + B^2 \tag{A.13}$$

The Principal's expected profits are therefore given by

$$E_x\{\pi(x)\} = 2B - \left[\frac{\gamma}{2}(B-b)^2 + \frac{\gamma}{2}b^2 + B^2\right] \tag{A.14}$$

[2] Introducing inter-agent transfers to ensure their participation constraints are satisfied does not change the results of this section. Therefore, for simplicity, such transfers will be ignored here.

Solving for the optimal values of b and B gives

$$b^c = \frac{2}{4+\gamma} \tag{A.15}$$

$$B^c = \frac{4}{4+\gamma} \tag{A.16}$$

and so the Principal's expected profits are

$$E\pi^c = \frac{4}{4+\gamma} \tag{A.17}$$

The first conclusion one can draw is that both agents are treated equally: they exert the same level of effort (equation A.11) and bear the same level of risk (from equations A.15 and A.16, the risk borne by Agent 1, $B^c - b^c$, is equal to the risk borne by Agent 2, b^c). This highlights the fact that the delegated structure presented in section 3.2.3 of chapter 3 increases the efficiency of the system (compared to the centralised structure) *in spite of* allocating efforts and risks asymmetrically between the agents.

In terms of expected profits, the Principal prefers the Cartel case, then the Delegated Second Best case, and finally the Centralised Second Best (see figure A.2).

The difference $E\pi^c - E\pi^{**}$ represents the profits lost due to the externality when contracting in a centralised manner. Defining the "Internalisation ratio" as

$$\xi := \frac{E\pi^d - E\pi^{**}}{E\pi^c - E\pi^{**}} \tag{A.18}$$

ξ can be interpreted as the proportion of the loss that is "recovered" (or the proportion of the externality that is internalised) when using delegation instead of centralised contracting. The ratio is depicted in figure A.3 and ranges from (approx.) 0.4 to (approx.) 0.68,[3] suggesting that delegation allows for the recovery of a significant part of the profits overlooked by centralised contracting. Furthermore, the ratio is greater than $\frac{1}{2}\ \forall\gamma < \frac{7+\sqrt{57}}{2} \approx 7.275$.

An alternative way to analyse the effectiveness of the DSB is by using the ratio of (expected) profits: the DSB's in the numerator and the Cartel's in the numerator. Formally,

$$v := \frac{E\pi^d}{E\pi^c} \tag{A.19}$$

[3]Actually, the ratio is equal to 1 when $\gamma = 0$, but in that the case the analysis is purely academic, since then First Best feasible.

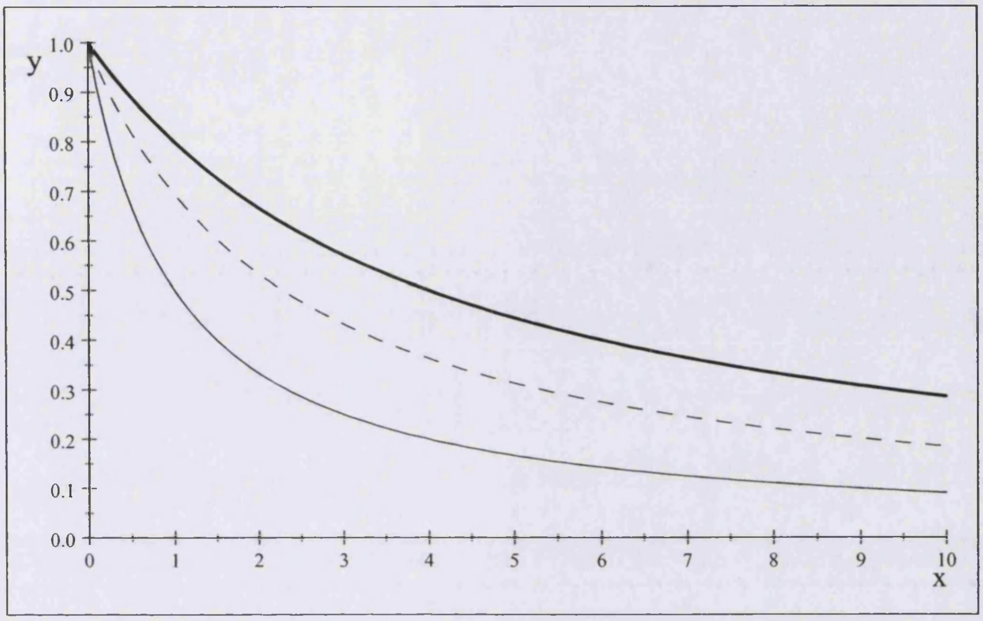

Figure A.2: Expected Profits.
Horizontal axis: γ. Vertical axis: Thick line: $E\pi^c$; Dotted line: $E\pi^d$; thin line: $E\pi^{**}$.

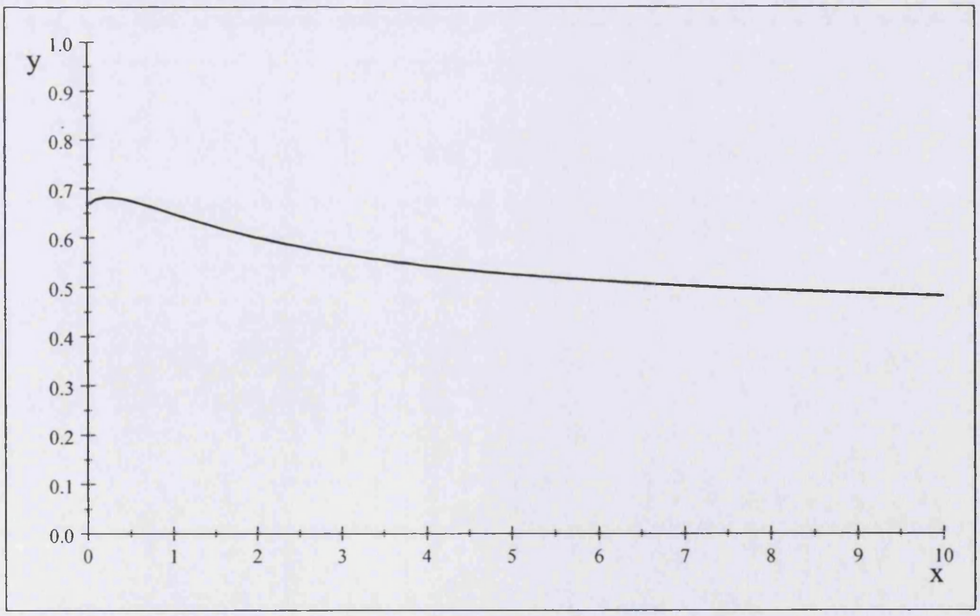

Figure A.3: Internalisation ratio.
Horizontal axis: γ. Vertical axis: ξ.

This ratio measures how close the DSB gets to the optimal scenario (Cartel) and is shown in figure A.4. It is equal to 1 when $\gamma = 0$ and tends to $\frac{9}{16} = 0.5625$ as $\gamma \to \infty$, thus confirming the effectiveness of the DSB suggested above.

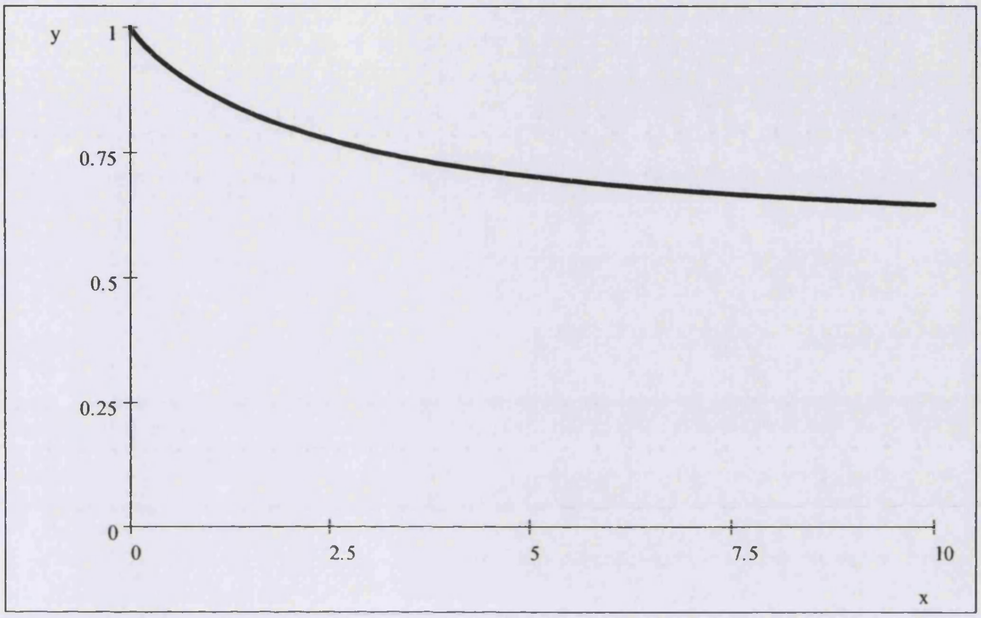

Figure A.4: DSB to Cartel ratio.
Horizontal axis: γ. Vertical axis: v.

In conclusion, when agents make their decisions jointly, the Principal offers contracts that treat them equally. The joint decision ensures that both agents fully internalise the externality, thus increasing the expected profits of the Principal (compared to those she gets under Centralised contracting). The Delegated Second Best can therefore be seen as an intermediate scenario in which only one agent fully internalises the externality (Agent 1), so that the expected profits under the DSB structure are higher than under the CSB but lower than under the Cartel. The effect of partial internalisation is, however, powerful, as suggested by the fact that the rate of "recovery" of "profits lost due to Centralised contracting" is higher than 50% for a large range of values of γ.

A.4 Composition of the "disutility portfolio"

It may be important to consider how the composition of an agent's disutility (or cost) varies as different contracting structures are used.

An agent's cost consists of two components: the disutility due to the effort exerted

$$\phi_i := \tfrac{1}{2}e_i^2 \qquad \forall i \in \{1,2\} \tag{A.20}$$

and the disutility generated by the risk borne

$$R_i := \tfrac{\gamma}{2}b_i^2 \qquad \forall i \in \{1,2\} \tag{A.21}$$

Thus, in order to analyse the behaviour of these components, a "Risk ratio" can be constructed as follows

$$\rho_i := \tfrac{R_i}{R_i+\phi_i} \qquad \forall i \in \{1,2\} \tag{A.22}$$

Equation A.23 and figure A.5 show the values of the ratios for each agent under the two relevant contracting structures, namely, Centralised Second Best and Delegated Second Best:

$$\begin{array}{cc} \rho_1^d = \frac{\gamma^3}{(\gamma+1)(3\gamma+\gamma^2+1)} & \rho_2^d = \frac{\gamma}{\gamma+1} \\ \rho_1^{**} = \frac{\gamma}{\gamma+1} & \rho_2^{**} = \frac{\gamma}{\gamma+1} \end{array} \tag{A.23}$$

It can be noted that, irrespective of the contracting structure, the ratios are equal to 0 when $\gamma = 0$ and they tend to 1 when $\gamma \to \infty$. This is consistent with intuition: when risk is not an issue ($\gamma = 0$), the ratio is 0 and only the disutility of effort matters; on the other hand, when risk is infinitely more important than effort ($\gamma \to \infty$), then the ratio tends to 1.

As expected, the equal treatment under CSB leads to the equality of agents' ratios ($\rho_1^{**} = \rho_2^{**}$). Under DSB, on the other hand, risk accounts for a greater fraction of total disutility for Agent 2 than for Agent 1 ($\rho_1^d < \rho_2^d$).

The comparison CSB v DSB yields the results that Agent 2's ratios are the same under both structures ($\rho_2^{**} = \rho_2^d$) and that Agent 1's ratio is lower under DSB than under CSB ($\rho_1^{**} > \rho_1^d$). The intuition behind these results relates to the trade-off between risk and incentives that is a standard feature of moral hazard settings: Agent 2's perception of the

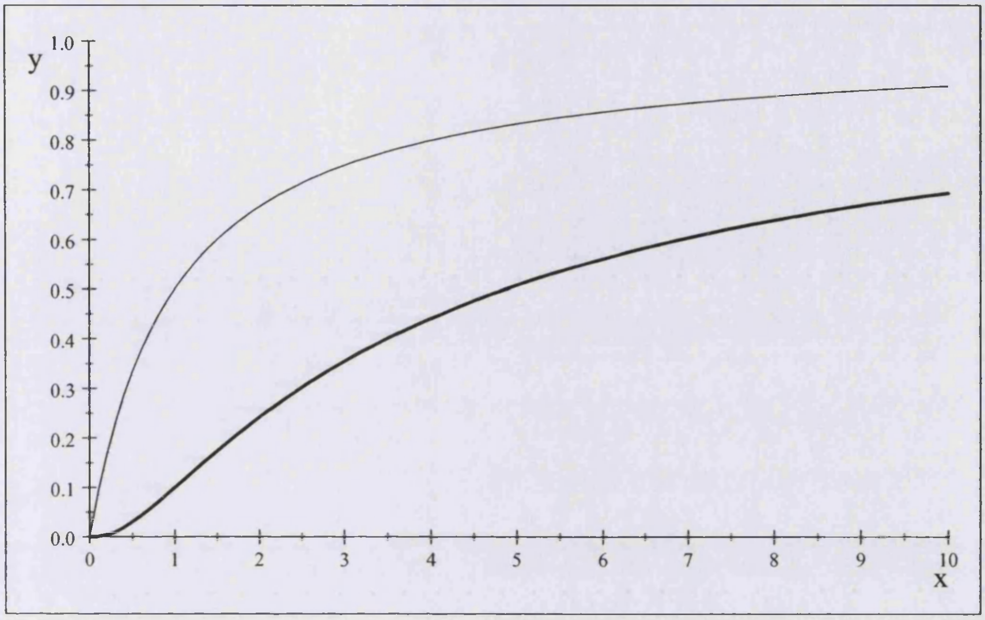

Figure A.5: Risk ratios.
Horizontal axis: γ. Vertical axis: Thick line: ρ_1^d; thin line: $\rho_2^d = \rho_1^{**} = \rho_2^{**}$

trade-off is unaffected by the change in structure, and so her decisions (and "portfolio") are unaffected. On the other hand, Agent 1 faces a "looser" trade-off under DSB than under CSB because she can transfer part of her risk to Agent 2 when choosing Agent 2's incentive scheme and, thus, she can decrease the importance of risk as a proportion of her total disutility.

All these results are in line with those presented in proposition 3.2 and, furthermore, they will also hold if the wage schedule of Agent 2, $w(x)$, were included as part of Agent 1's cost in the DSB case. This is so because it only increases the numerator of the latter's risk ratio.

An important caveat of the present analysis, however, is that different functional forms could lead to slightly different results, though the most important one (that Agent 1's ratio is lower under DSB than under CSB, $\rho_1^{**} > \rho_1^d$) is expected to be robust to the abovementioned modifications.

A.5 No replicability result

The basic (intuitive) proof derives from the Centralised Second Best in section 3.2.2: the Principal can choose any pair of contracts $(\tilde{a}_1, \tilde{b}_1)$, $(\tilde{a}_2, \tilde{b}_2)$, including those corresponding to the delegated structure, (a_1^d, b_1^d), (a_2^d, b_2^d), but she chooses the CSB contracts (a_1^{**}, b_1^{**}), (a_2^{**}, b_2^{**}) instead.

Formally, consider the case when the Principal offers the agents the Delegated Second Best contracts:

$$
\begin{aligned}
(\tilde{a}_1, \tilde{b}_1) &= (A^d - a^d, B^d - b^d) && \text{(A.24)}\\
(\tilde{a}_2, \tilde{b}_2) &= (a^d, b^d) && \text{(A.25)}
\end{aligned}
$$

Agents then choose their levels of effort optimally, which requires their efforts to be equal to the power of their respective incentive schemes:

$$
\begin{aligned}
\tilde{e}_1 &= B^d - b^d && \text{(A.26)}\\
\tilde{e}_2 &= b^d && \text{(A.27)}
\end{aligned}
$$

Notice, however, that while Agent 2 will choose the same level of effort as in the DSB setting, Agent 1 will exert less effort than in the DSB case:

$$
\begin{aligned}
\tilde{e}_2 &= b^d = e_2^d && \text{(A.28)}\\
\tilde{e}_1 &= B^d - b^d < B^d = e_1^d && \text{(A.29)}
\end{aligned}
$$

As a consequence, expected output will be strictly lower than in the DSB case:

$$\tilde{x} = \tilde{e}_1 + \tilde{e}_2 = b^d < B^d + b^d = e_1^d + e_2^d = x^d \tag{A.30}$$

Now, in the DSB case, the participation constraints of both agents are binding. In the case of Agent 2, this requires

$$EU_2^d = a_2^d + b_2^d x^d - \frac{\gamma}{2}\left(b_2^d\right)^2 - \frac{1}{2}\left(e_2^d\right)^2 = 0 \tag{A.31}$$

so that

$$a_2^d = -b_2^d x^d + \frac{\gamma}{2}\left(b_2^d\right)^2 + \frac{1}{2}\left(e_2^d\right)^2 \tag{A.32}$$

The expected utility of Agent 2 when offered contract $(\tilde{a}_2, \tilde{b}_2)$ is therefore

$$E\tilde{U}_2 = \tilde{a}_2 + \tilde{b}_2\tilde{x} - \frac{\gamma}{2}\left(\tilde{b}_2\right)^2 - \frac{1}{2}\left(\tilde{e}_2\right)^2 \tag{A.33}$$

Using the information about the contract, optimal choice of effort and value of a_2^d (equations A.25, A.29 and A.32) equation A.33 becomes

$$E\tilde{U}_2 = -b_2^d\left(x^d - \tilde{x}\right)$$

and so (due to equation A.30) the expected utility of Agent 2 is negative.

That is, if the Principal offered the DSB contracts to the agents, Agent 2 would reject the offer. This means that the Principal cannot replicate the Delegated Second Best result because the contract offered to Agent 2 is not individually rational. The rationale for this is that Agent 1 has no incentives to exert extra effort ($\tilde{e}_1 < e_1^d$) because she is not residual claimant and, therefore, she cannot take advantage of choosing simultaneously both her effort and Agent 2's incentive scheme. As a consequence, expected output is lower and Agent 2's expected income does not compensate for the disutility of effort and risk she suffers, so that she would rather reject the offer.

Appendix B

Endogenous Externality: Theory

B.1 Expected net revenue

The existing literature (Reinganum and Wilde (1986), Cronshaw and Alm (1995), etc.) usually considers the expected net revenue (taxes plus fines minus auditing costs) as the government's objective function, justifying it by arguments that range from the willingness to avoid using normative social welfare functions to the assumption that, by nature, a tax agency's goal is to collect as much revenue as possible.

In the model used in chapter 4, expected net revenue (ENR) is given by the expression

$$ENR = -(1-\gamma)\,pc + \gamma\,(1-\kappa)\,t + \gamma\kappa p\,((1+\varsigma)\,t - c) \tag{B.1}$$

The first term corresponds to the revenue lost (c) when income is low (which happens with probability $1-\gamma$) and the agency audits with probability p. The second term corresponds to the expected net revenue when income is high (which occurs with probability γ) and the taxpayer complies (which occurs with probability $1-\kappa$): the taxpayer pays tax t and is not audited. The final term corresponds to the case when income is high (which occurs with probability γ), the taxpayer evades (which occurs with probability κ) and the agency audits (with probability p): the taxpayer is caught and pays a fine $(1+\varsigma)\,t$. In this setting, the government's "type" (its private information) can be interpreted as the cost of an audit, c. This assumption is a plausible one, since usually taxpayers do not know how costly an audit it, while they are more likely to know other parameters of the problem like the tax rate t or the surcharge rate ς.

In order to find the relationship between the two competing objective functions, it is necessary to construct their respective "gap" functions similar to those used by HM Revenue and Customs in the United Kingdom (see Ratto et al. (2005)). These functions measure the gap between the first-best (perfect information) outcome and the actual value of ENR (or EL). Formally,

$$xGap := x - x^* \qquad x \in \{ENR, EL\} \tag{B.2}$$

In the case of the expected loss, an agency that knew with certainty taxpayers' incomes would make no targeting error, so the expected loss would be zero ($EL^* = 0$) and the $ELGap$ would be equal to the EL function:

$$\begin{aligned} ELGap \ &: = EL - EL^* && \text{(B.3)} \\ &= EL && \text{(B.4)} \\ &= (1-\gamma)\,p\,(1-\lambda) + \gamma\kappa\,(1-p)\,\lambda && \text{(B.5)} \end{aligned}$$

where the first term corresponds to the loss due to zeal errors and the second one to the loss due to negligence errors.

Analogously, the first-best ENR is given by

$$\begin{aligned} ENR^* &= (1-\gamma)\cdot 0 + \gamma\,\{(1-\kappa)\,t + \kappa\,[(1-0)\cdot 0 + 1\cdot((1+\varsigma)\,t - c)]\} && \text{(B.6)} \\ &= \gamma\,\{(1-\kappa)\,t + \kappa\,[(1+\varsigma)\,t - c]\} && \text{(B.7)} \end{aligned}$$

since the government will not audit anyone who declares her true income (be it low or high) and will audit everyone who evades. The $ENRGap$ function is then

$$\begin{aligned} ENRGap \ &: = ENR^* - ENR && \text{(B.8)} \\ &= (1-\gamma)\,pc + \gamma\kappa\,(1-p)\,[(1+\varsigma)\,t - c] && \text{(B.9)} \end{aligned}$$

A straightforward comparison between equations B.5 and B.9 suggests a high degree of similarity between the two objective functions. Formally, the ENRGap can be expressed as a linear function of the ELGap (and vice versa):

$$ENRGap = \alpha + \beta \cdot ELGap \tag{B.10}$$

or, using the equations above,

$$(1-\gamma)\,c\cdot p+\gamma\,((1+\varsigma)\,t-c)\cdot\kappa\,(1-p)=$$
$$\alpha+\beta\,(1-\gamma)\,(1-\lambda)\cdot p+\beta\gamma\lambda\cdot\kappa\,(1-p) \quad \text{(B.11)}$$

This relationship holds for every $(\kappa, p) \in [0,1] \times [0,1]$ if and only if

$$\lambda = 1 - \frac{c}{(1+\varsigma)\,t} \quad \text{(B.12)}$$

and

$$\alpha = 0 \quad \text{(B.13)}$$
$$\beta = (1+\varsigma)\,t \quad \text{(B.14)}$$

It can be seen that the greater is the cost of an audit, the lower is λ (i.e., the importance attached by the agency to negligence errors). This is understandable, as both situations (low audit cost and high importance of negligence errors) lead the government to undertake the same action, namely, to follow a high-intensity auditing policy. Thus, if $c = 0$ (auditing is costless), then $\lambda = 1$ (government is only concerned with negligence errors), and so the agency will audit as much as possible. On the other hand, if $c = (1+\varsigma)\,t$ (i.e., when the fine paid by a discovered evader just covers the cost of the audit), $\lambda = 1$ (government is only concerned with zeal errors), and so the agency will audit nobody.

The interpretation of equation B.12 is easier if it is re-organised as follows:

$$\frac{1-\lambda}{\lambda} = \frac{c}{(1+\varsigma)\,t - c} \quad \text{(B.15)}$$

where the left hand side measures the relative importance of zeal errors $(1-\lambda)$ with respect to the importance of negligence errors (λ). The condition in equation B.15, hence, is a logical one: the relative importance of zeal *vis-à-vis* negligence, $\frac{1-\lambda}{\lambda}$, must be equal to the ratio between the cost of an audit, c (the loss in case of a zeal error) and the net fine extracted from a caught evader, $(1+\varsigma)\,t-c$ (the revenue not collected in case of a negligence error).

B.2 Imperfect correlation

When imperfect correlation is allowed, the results of chapter 4 will hold as long as the common shocks are the main source of income variability.

In the simplest case, assume that a taxpayer's income y_i consists of two elements: a common component $y \in \{0, 1\}$ and an idiosyncratic one $v_i \in \{-v, v\}$. The idiosyncratic shock is assumed small compared to the common shock. In this particular case, $v < \frac{1}{2}$. This means, in particular, that the government can still detect the common shock if at least one person declares $1 - v$ or $1 + v$. The probability of a good common shock (good year) is $\gamma \in (0, 1)$ as before; the probability of a good idiosyncratic shock is $\frac{1}{2}$ (i.e., the idiosyncratic shock v_i is a white noise variable).

Audits are assumed to be profitable if someone with a positive common shock evades and is audited; otherwise, they are unprofitable. That is, the important case is when the common shock is missed. The idiosyncratic shocks are so small that even if the agency knew for sure that someone underdeclared their idiosyncratic positive shock (though declaring their common shock truthfully), it does not pay off for the government to undertake the audit. This is a radical assumption, but highlights the importance of the common shock, which is the focus of the present analysis. Formally, this requires the following condition regarding the cost of an audit c:

$$(1 + \varsigma)\, t \cdot 2v < c < (1 + \varsigma)\, t \tag{B.16}$$

It implies that the fine collected from someone who underdeclares the common component but tells the truth about her idiosyncratic component (the expression on the right hand side) is greater than the cost of the audit, which in turn is greater than the fine collected from someone who tells the truth about the common component but underdeclares her idiosyncratic one (the expression on the left hand side).

In such scenario, as in the perfect correlation case, nobody who declares the maximum possible income $(1 + v)$ is ever audited: $a(1 + v, \mathbf{d}_{-i}) = 0$. The relative position of someone's declaration also matters, though the relationship is a bit more complex than in the perfect correlation scenario: anyone who declares $-v$ or v when at least one other person declares $1 - v$ or $1 + v$ will be audited with certainty, because the government knows that they underdeclared the common shock: $a\,(\pm v, 1 \pm v) = 1$

From the previous assumptions, it is also clear that nobody who declares $1 - v$ will ever be audited: the government knows the person has high income, and auditing the person in order to recover a small fine from the (potentially) underdeclared idiosyncratic shock is unprofitable: $a(1 - v, \mathbf{d}_{-i}) = 0$.

The only cases that are left for analysis, therefore, are those when:

1. both taxpayers declare v,
2. both taxpayers declare $-v$, and
3. one taxpayer declares v, the other $-v$.

The three cases, however, lead to the same result. This is the case because only the common shock matters when computing the negligence and zeal errors, which are therefore the same in all three cases. The agency's objective function is

$$EL(\pm v, \pm v) = \gamma \cdot 2\lambda \cdot [1 - a(\pm v, \pm v)] + (1-\gamma) \cdot 2(1-\lambda) \cdot a(\pm v, \pm v) \tag{B.17}$$

and so the optimal strategy is

$$a(\pm v, \pm v) = \begin{cases} 0 & \text{if } \lambda < 1-\gamma \\ \in \{0,1\} & \text{if } \lambda = 1-\gamma \\ 1 & \text{if } \lambda > 1-\gamma \end{cases} \tag{B.18}$$

Thus, summarizing all the results, the agency's optimal policy is

$$a(d, d', \lambda) = \begin{cases} 0 & \text{if } d = 1 \pm v \\ 1 & \text{if } d = \pm v,\ d' = 1 \pm v \\ 1 & \text{if } d = \pm v,\ d' = \pm v,\ \lambda > 1-\gamma \\ 0 & \text{if } d = \pm v,\ d' = \pm v,\ \lambda < 1-\gamma \end{cases} \tag{B.19}$$

This auditing policy is (weakly) decreasing in a taxpayer's declaration and increasing in the other taxpayer's declaration, a result that replicates the one in proposition 4.4.

It is straightforward to show that taxpayers' declarations are (weakly) increasing functions of their expected probabilities of detection, and so the equivalent of proposition 4.5 is obtained.

The combination of these two results yield the strategic complementarity analysed in chapter 4 and which generates the associated coordination game.

B.3 More levels of income

The analysis with more levels of income is more cumbersome than when only two levels of income are considered, so for the time being the only situation that will be investigated is the one in which income can take three possible levels: $y \in \{0, \frac{1}{2}, 1\}$, which occur with probabilities l, m and h, respectively $(l + m + h := 1)$. In such scenario the intermediate case $(y = \frac{1}{2})$ can be interpreted as the "status quo" while the lower and higher ones ($y = 0$ and $y = 1$, respectively) can be interpreted as the bad and good years, respectively. Let us consider the case with just two taxpayers, i and j.

A taxpayer can therefore declare only 0, $\frac{1}{2}$ or 1. Expected utility in each case is given by

$$EU_i\left(d_i \mid y_i\right) = y_i - td_i - (1 + \varsigma)\, t\left(y_i - d_i\right) \cdot E_i\left[a\left(d_i, \mathbf{d}_{-i}\right)\right] \tag{B.20}$$

and it is straightforward to show that declared income is an increasing function of the expected probability of detection, just as in the case with just two levels of income (equation 4.38.1).

Considering now the tax agency's problem, two results are self-evident:

- Anyone who declares 1 will never be audited:

$$a\left(1, d_j\right) = 0 \tag{B.21}$$

- If different taxpayers declare different incomes, the one who declares less income will always be audited (assuming that audits are profitable):

$$a\left(d_i, d_j\right) = 1 \qquad \forall d_i < d_j \tag{B.22}$$

Hence, we need only to consider what happens in the following cases:

1. Both taxpayers declare 0,
2. Both taxpayers declare $\frac{1}{2}$, and
3. One taxpayer declares 0 and the other $\frac{1}{2}$.

In the first case, the government's expected loss is

$$\begin{aligned} EL_{TA}(\mathbf{d}=(0,0)) &= l\left[2(1-\lambda)\cdot a(0,0)+0\cdot(1-a(0,0))\right]+ \\ &\quad +m\left[0\cdot a(0,0)+2\lambda\cdot(1-a(0,0))\right]+ \qquad \text{(B.23)} \\ &\quad +h\left[0\cdot a(0,0)+2\lambda\cdot(1-a(0,0))\right] \\ &= 2\left[(l-\lambda)\cdot a(0,0)+(1-l)\lambda\right] \qquad \text{(B.24)} \end{aligned}$$

where the first term is the loss if $y=0$ (if both are audited: two zeal errors, if none is audited: no errors) and the second and third the respective losses when $y=\frac{1}{2}$ and $y=1$ (if both are audited: no losses, if none is audited: two negligence mistakes). The agency therefore chooses the auditing strategy such that

$$a(0,0)=\begin{cases} 0 & \text{if } \lambda<l \\ \in\{0,1\} & \text{if } \lambda=l \\ 1 & \text{if } \lambda>l \end{cases} \qquad \text{(B.25)}$$

In the second case, the government's expected loss is

$$\begin{aligned} EL_{TA}\left(\mathbf{d}=\left(\frac{1}{2},\frac{1}{2}\right)\right) &= \phi\left[2(1-\lambda)\cdot a\left(\frac{1}{2},\frac{1}{2}\right)+0\cdot\left(1-a\left(\frac{1}{2},\frac{1}{2}\right)\right)\right]+ \\ &\quad +(1-\phi)\left[0\cdot a\left(\frac{1}{2},\frac{1}{2}\right)+2\lambda\cdot\left(1-a\left(\frac{1}{2},\frac{1}{2}\right)\right)\right] \qquad \text{(B.26)} \\ &= 2\phi(1-\lambda)\cdot a\left(\frac{1}{2},\frac{1}{2}\right)+ \\ &\quad +2(1-\phi)\lambda\cdot\left(1-a\left(\frac{1}{2},\frac{1}{2}\right)\right) \qquad \text{(B.27)} \end{aligned}$$

where

$$\phi:=\frac{m}{m+h} \qquad \text{(B.28)}$$

is the probability of $y=\frac{1}{2}$ conditional on $\mathbf{d}=\left(\frac{1}{2},\frac{1}{2}\right)$. That is, the government can discard $y=0$ given that at least one person declared $\frac{1}{2}$. The optimal strategy is now

$$a\left(\frac{1}{2},\frac{1}{2}\right)=\begin{cases} 0 & \text{if } \lambda<\phi \\ \in\{0,1\} & \text{if } \lambda=\phi \\ 1 & \text{if } \lambda>\phi \end{cases} \qquad \text{(B.29)}$$

In the final case, the expected loss is

$$
\begin{aligned}
EL_{TA}\left(\mathbf{d}=\left(\frac{1}{2},0\right)\right) &= \phi\left[(1-\lambda)\cdot a\left(\frac{1}{2},0\right)+0\cdot\left(1-a\left(\frac{1}{2},0\right)\right)\right]+ \\
&\quad +(1-\phi)\left[0\cdot a\left(\frac{1}{2},0\right)+\lambda\cdot\left(1-a\left(\frac{1}{2},0\right)\right)\right] \quad \text{(B.30)} \\
&= \phi(1-\lambda)\cdot a\left(\frac{1}{2},0\right)+(1-\phi)\lambda\cdot\left(1-a\left(\frac{1}{2},0\right)\right) \quad \text{(B.31)}
\end{aligned}
$$

and the optimal strategy is

$$
a\left(\frac{1}{2},0\right)=\begin{cases} 0 & \text{if } \lambda<\phi \\ \in\{0,1\} & \text{if } \lambda=\phi \\ 1 & \text{if } \lambda>\phi \end{cases} \quad \text{(B.32)}
$$

Getting all cases together, the government's optimal strategy is

$$
a\left(d,d',\lambda\right)=\begin{cases} 0 & \text{if } d=1 \\ 1 & \text{if } d=\frac{1}{2},\ d'=1 \\ 1 & \text{if } d=\frac{1}{2},\ d'<1,\ \lambda>l \\ 0 & \text{if } d=\frac{1}{2},\ d'<1,\ \lambda<l \\ 1 & \text{if } d=0,\ d'>0 \\ 1 & \text{if } d=0,\ d'=0,\ \lambda>\phi \\ 0 & \text{if } d=0,\ d'=0,\ \lambda<\phi \end{cases} \quad \text{(B.33)}
$$

It is straightforward to show that, keeping d and λ fixed, as d' increases, $a(d,d')$ also (weakly) increases. Also, under mild assumptions[1], it can be shown that (for fixed d' and λ), higher d leads to a (weakly) lower probability of detection $a(d,d')$. These two results are the counter-parts of proposition 4.4 in the two-income case.

Similarly mild conditions (i.e., based on hazard rates) are expected to be necessary for the results to hold when more than three levels of income are considered. Thus, the optimal

[1] The condition needed is

$$
l<\phi \quad \text{(B.34)}
$$

In a more general environment, it requires the "hazard rate" to be increasing in income:

$$
\frac{\partial\left(\frac{f(x)}{1-F(x)}\right)}{\partial x}>0 \quad \text{(B.35)}
$$

a condition that is satisfied as long as the probability distribution of shocks is not too skewed to the right. It is satisfied, for example, by the uniform distribution and for a symmetric distribution where $l=h<m$ (which could replicate a situation with a "most likely scenario" (m) and both "bad" and "good" ones (l and h)).

policy being non-decreasing in the average declaration and the associated coordination game it creates seem to be robust features of the analysis.

B.4 Proofs

Proof. Proposition 4.2

Derive the expected loss function (equation 4.27) with respect to the agency's two policy variables, namely, $a_0 := a(0, D, \lambda)$ and $a_1 := a(1, D, \lambda)$.

For the first part of the proposition, compute the derivative of the expected loss with respect to a_1:

$$\frac{\partial E_{TA}(L)}{\partial a_1} = (1 - \lambda) D \tag{B.36}$$

which is positive[2], so that the optimal strategy in order to minimise expected losses is to set

$$a_1^* = 0 \tag{B.37}$$

That is, the agency must not audit anyone who declares high income.

For the last three parts of the proposition, that determine the value of a_0, it is necessary to distinguish two cases: one when the average declaration is zero ($D = 0$, parts 2 and 3 of the proposition) and another when the average declaration is positive ($D > 0$, part 4 of the proposition).

Consider first the scenario in which the average declaration is positive ($D > 0$). Since it is common knowledge that declaring low income is the dominant strategy for taxpayers when income is low (proposition 4.1), the agency is able to infer that whoever declares high income says the truth, i.e., the posterior probability of the taxpayer having high income conditional on the taxpayer having declared high income is 1:

$$\Pr(y_i = 1 \mid d_i = 1) = 1 \tag{B.38}$$

Furthermore, given the perfect correlation between incomes, if the agency observes at least one high declaration (i.e., if $D > 0$), it is able to infer that the year is a good one with probability one: the posterior probability of a good year ($y = 1$) conditional on the

[2]Variable a_1 is the audit decision regarding a taxpayer who declares high income. This means that at least one person declared high income, and so the average declaration D is strictly positive.

average declaration being strictly positive is 1:

$$\Pr\left(y = 1 \mid D > 0\right) = 1 \tag{B.39}$$

(thus the first part of equation 4.16). From the definition of a negligence error (equation 4.8), therefore, the agency's best strategy is to audit everyone who declares low income ($d_i = 0$) when average income is strictly positive ($D > 0$). This proves the last part of the proposition.

Consider now the case when every taxpayer declares low income, so that the average declaration is zero ($D = 0$). In this scenario, the agency does not know if income is truly low ($y = 0$) or if it is high and every taxpayer evaded ($y = 1$ and $d_i = 0$ for every $i \in [0, 1]$). The agency's posterior belief over the probability of a good year conditional on the average declaration being equal to zero is therefore:

$$\Pr\left(y = 1 \mid D = 0\right) = \gamma \tag{B.40}$$

where γ is the prior probability of a good year (thus the second part of equation 4.16). Hence, deriving the expected loss function (equation 4.27) with respect to the audit decision a_0 yields

$$\frac{\partial E_{TA}\left(L\right)}{\partial a_0} = 1 - \lambda - \gamma \tag{B.41}$$

This expression is positive if $\lambda < 1 - \gamma$ and negative otherwise, so the agency will audit in the latter case ($\lambda > \tilde{\lambda} := 1 - \gamma$) and will not audit in the the first case ($\lambda < \tilde{\lambda} := 1 - \gamma$). This proves parts 2 and 3 of the proposition ■

Proof. Proposition 4.7

The proof is similar to that in Morris and Shin (1997), so I will concentrate on those elements specific to my model.

Propose a taxpayer strategy (for good years) of the following type:

$$d_i\left(s_i\right) = \begin{cases} 0 & \text{if} \quad s_i < \hat{s} \\ \in [0,1] & \text{if} \quad s_i = \hat{s} \\ 1 & \text{if} \quad \hat{s} < s_i \end{cases} \tag{B.42}$$

that is, a strategy according to which the taxpayer declares high income (comply) when the signal is sufficiently high and low income (evades) when it is sufficiently low.

To be optimal, such strategy must satisfy the following condition

$$\hat{s} \gtrless s_i \quad \text{if and only if} \quad E_i\left(u\left(evasion\right)\right) \gtrless u\left(compliance\right) \tag{B.43}$$

where $E_i\left(u\left(evasion\right)\right)$ and $u\left(compliance\right)$ are given by equations 4.34 and 4.35.

The average declaration in the economy is therefore given by the proportion of taxpayers that receive signals greater than $\hat{s}$. Since signals are uniformly distributed around the true type of the agency λ (with support $[\lambda - \varepsilon, \lambda + \varepsilon]$), there are three cases to consider:

1. If $\hat{s} \leq \lambda - \varepsilon$, then $D = 1$;
2. If $\lambda - \varepsilon < \hat{s} < \lambda + \varepsilon$, then $D = \int_{\hat{s}}^{\lambda+\varepsilon} \frac{1}{2\varepsilon} ds = \frac{\lambda+\varepsilon-\hat{s}}{2\varepsilon}$; and
3. If $\lambda + \varepsilon \leq \hat{s}$, then $D = 0$;

That is, the average declaration in the economy D is a weakly increasing function of the type of the agency λ. In particular, if the agency is so tough (case 1 above) that even the person with the lowest signal $(\lambda - \varepsilon)$ complies (see equation B.42), then everyone declares high income and $D = 1$. On the other hand, when the agency is so soft (case 3 above) that even the person with the highest signal $(\lambda + \varepsilon)$ evades, then everyone declares low income and $D = 0$. In intermediate cases, some people declare high income and others declare low income, so $D \in (0, 1)$. Formally,

$$D(\lambda) = \begin{cases} 0 & \text{if} \quad \lambda \leq \hat{s} - \varepsilon \\ \in [0,1] & \text{if} \quad \hat{s} - \varepsilon < \lambda < \hat{s} + \varepsilon \\ 1 & \text{if} \quad \hat{s} + \varepsilon \leq \lambda \end{cases} \tag{B.44}$$

Equation 4.28 gives the agency's optimal strategy. Such strategy depends on the individual taxpayer's declaration d_i, the declarations of every other taxpayer in the category D, and the type of agency λ. In particular, it states that the agency will audit a person who declares low income $(d_i = 0)$ only if one of the two following scenarios occur:

1. The average declaration in the economy is strictly positive: $D > 0$.
2. The average declaration in the economy is zero and the agency is "tough": $D = 0$ and $\lambda > \tilde{\lambda}$ (where $\tilde{\lambda}$ is defined as in proposition 4.2).

Consider first case 1: From equation B.44, this scenario occurs when $\lambda > \hat{s} - \varepsilon$.

Case 2, on the other hand, requires both $\lambda \leq \hat{s} - \varepsilon$ (from equation B.44) and $\lambda > \tilde{\lambda}$ (and therefore, implicitly, that $\tilde{\lambda} < \hat{s} - \varepsilon$). Combining the two, case 2 occurs when $\tilde{\lambda} < \lambda \leq \hat{s} - \varepsilon$.

Therefore, the agency audits a taxpayer who declares low income only if $\lambda > \hat{\lambda}$, where $\hat{\lambda} := \min\left\{\hat{s} - \varepsilon, \tilde{\lambda}\right\}$.

The agency's optimal strategy can then be reduced to the following expression

$$a_i(d_i, \lambda) = \begin{cases} 0 & \text{if } d_i = 0 \\ \in [0,1] & \text{if } d_i = 1 \text{ and } \lambda \leq \hat{\lambda} \\ 1 & \text{if } d_i = 1 \text{ and } \lambda > \hat{\lambda} \end{cases} \tag{B.45}$$

The expected utility of a taxpayer i who evades is given by expression

$$E_i(u(0, a_i, 1)) = E\left[u(0, a_i(0, \lambda), 1) \mid s\right] \tag{B.46}$$

Taxpayer i's posterior distribution of λ conditional on her private signal s is uniformly distributed around s, so there are three cases to consider:[3]

1. If $\hat{\lambda} \leq s - \varepsilon$, then $a_i = 1$ and $E_i(u(0, a_i, 1)) = \int_{s-\varepsilon}^{s+\varepsilon} \frac{u(0,1,1)}{2\varepsilon} d\lambda = \int_{s-\varepsilon}^{s+\varepsilon} \frac{1-f}{2\varepsilon} d\lambda = 1 - f$;
2. If $s - \varepsilon < \hat{\lambda} < s + \varepsilon$, then $E_i(u(0, a_i, 1)) = \int_{s-\varepsilon}^{\hat{\lambda}} \frac{u(0,0,1)}{2\varepsilon} d\lambda + \int_{\hat{\lambda}}^{s+\varepsilon} \frac{u(0,1,1)}{2\varepsilon} d\lambda = \int_{s-\varepsilon}^{\hat{\lambda}} \frac{1}{2\varepsilon} d\lambda + \int_{\hat{\lambda}}^{s+\varepsilon} \frac{1-f}{2\varepsilon} d\lambda = 1 - \frac{s+\varepsilon-\hat{\lambda}}{2\varepsilon} f$; and
3. If $s + \varepsilon \leq \hat{\lambda}$, then $a_i = 0$ and $E_i(u(0, a_i, 1)) = \int_{s-\varepsilon}^{s+\varepsilon} \frac{u(0,0,1)}{2\varepsilon} d\lambda = \int_{s-\varepsilon}^{s+\varepsilon} \frac{1}{2\varepsilon} d\lambda = 1$.

Intuitively, if the signal is so high (case 1) that even her lowest estimate of λ, $s - \varepsilon$, is high enough as to trigger an audit (from equation B.45), then her payoff from evasion is $1 - f$ with certainty. On the other hand, if the signal is so low (case 3) that even her highest estimate of λ, $s + \varepsilon$, is low enough as to avoid triggering an audit, then her payoff from evasion is 1 with certainty. In intermediate cases, the taxpayer's payoff is not certain, and the expected utility takes values in the $(1 - f, 1)$ range.

Thus, the taxpayer's expected utility of evasion is given by the following expression:

$$E_i(u(0, a_i, 1)) = E\left[u(0, a_i(0, \lambda), 1) \mid s\right] = \begin{cases} 1 & \text{if } s \leq \hat{\lambda} - \varepsilon \\ 1 - \frac{s+\varepsilon-\hat{\lambda}}{2\varepsilon} f & \text{if } \hat{\lambda} - \varepsilon < s < \hat{\lambda} + \varepsilon \\ 1 - f & \text{if } \hat{\lambda} + \varepsilon \leq s \end{cases} \tag{B.47}$$

[3] f is the fine paid if caught evading, as defined in equation 4.5 .

Note that this function is a continuous and (weakly) decreasing function of the taxpayer's signal and that it can take values in the range $[1-f, 1]$. From the definition of f (equation 4.5), it is straightforward to show that

$$1 - f < 1 - t < 1 \tag{B.48}$$

This means that the utility of compliance (the term in the centre, from equation 4.35) is higher than the utility if caught evading (the left-hand side term) but lower than the utility if evasion goes undetected (the right-hand side term). This implies that exists a signal $\hat{s}$ such that the expected utility of evasion (equation B.47) equals the utility of compliance $1 - t$. Formally

$$E\left[u\left(0, a_i\left(0, \lambda\right), 1\right) \mid \hat{s}\right] := 1 - t \tag{B.49}$$

or, equivalently, using equation B.47,

$$\hat{s} := \hat{\lambda} - \varepsilon + 2\varepsilon P \tag{B.50}$$

where P is defined as in equation 4.36.

When $\hat{\lambda} = \hat{s} - \varepsilon$, then equation B.50 becomes

$$\hat{s} = (\hat{s} - \varepsilon) - \varepsilon + 2\varepsilon P \tag{B.51}$$

which simplifies to

$$2\varepsilon\left(1 - P\right) = 0 \tag{B.52}$$

and is only satisfied in extreme and rather uninteresting cases: when $\varepsilon = 0$ (no fundamental uncertainty regarding the type of agency λ) and/or $P := \frac{1}{1+\varsigma} = 1$ (when no fine is paid if caught evading). Thus, this case will be ignored.

When $\hat{\lambda} = \tilde{\lambda}$, on the other hand, the switching point becomes

$$\hat{s} := \tilde{\lambda} - \varepsilon + 2\varepsilon P \tag{B.53}$$

and, since the expected utility of evasion (equation B.47) is a weakly decreasing function of the private signal received by the taxpayer and $1 - f < 1 - t < 1$, this means that $\hat{s}$ is unique. This proves the first part of the proposition: every taxpayer follows the same threshold strategy.

Furthermore, it is now straightforward to show that the expected utility of evasion is higher (respectively, lower) than the utility of compliance when the private signal s is lower (respectively, higher) than the threshold $\hat{s}$, thus proving that the threshold strategy of equation 4.40.1 is indeed optimal (i.e., satisfies the condition in equation B.43) and hence, trivially, that the optimal declaration strategy is a weakly increasing function of the private signal received.This proves the second part of the proposition. ■

Proof. Proposition 4.8

The average declaration is defined as

$$D := \int_s d_i(1,s)\, dG(s \mid \lambda) \tag{B.54}$$

where $G(s \mid \lambda)$ is the probability distribution of signals, conditional on the type of agency being λ. From equation 4.13, $s \mid \lambda$ is uniformly distributed on the $[\lambda - \varepsilon, \lambda + \varepsilon]$ segment. Note that, because of the taxpayer's optimal strategy in good years (proposition 4.7), the average declaration can be interpreted as the fraction of the population that gets a signal above the threshold $\hat{s}$.

Depending on the value of λ, three cases can occur:

1. Full evasion ($\lambda < \hat{s} - \varepsilon$): Even the person with highest signal (i.e., $s_i = \lambda + \varepsilon$) would evade. Formally, $D := \int_{\lambda-\varepsilon}^{\lambda+\varepsilon} (0) \frac{1}{2\varepsilon} ds = 0$.
2. Partial evasion ($\hat{s} - \varepsilon < \lambda < \hat{s} + \varepsilon$): Those with signals between $\lambda - \varepsilon$ and $\hat{s}$ evade, those with signals between $\hat{s}$ and $\lambda + \varepsilon$ comply. Formally, $D := \int_{\lambda-\varepsilon}^{\hat{s}} (0) \cdot \frac{1}{2\varepsilon} ds + \int_{\hat{s}}^{\lambda+\varepsilon} (1) \frac{1}{2\varepsilon} ds = \frac{\lambda+\varepsilon-\hat{s}}{2\varepsilon}$.
3. Full compliance ($\hat{s} + \varepsilon < \lambda$): Even the person with the lowest signal (i.e., $s_i = \lambda - \varepsilon$) would comply. Formally, $D := \int_{\lambda-\varepsilon}^{\lambda+\varepsilon} (1) \frac{1}{2\varepsilon} ds = 1$.

The level of evasion in good years is simply the fraction of the population that gets a signal below the threshold $\hat{s}$. That is, $\kappa = 1 - D$. ■

Proof. Proposition 4.11

Consider first the full evasion case ($\lambda < \hat{s} - \varepsilon$). Since the threshold $\hat{s}$ is defined as in equation B.53, the condition $\lambda < \hat{s} - \varepsilon$ becomes

$$\varepsilon < \frac{\tilde{\lambda} - \lambda}{2(1-P)} \tag{B.55}$$

where P is the auditing intensity that eliminates evasion.

Since $P \in \left(\frac{1}{2}, 1\right)$, $1 - P$ can only take values in the interval $\left(0, \frac{1}{2}\right)$. Also, since the noise of the signals cannot be negative, it must be the case that

$$0 < \varepsilon \tag{B.56}$$

Combining equations B.55 and B.56, the full evasion case requires $0 < \varepsilon < \frac{\tilde{\lambda}-\lambda}{2(1-P)}$, which is only feasible if $\lambda < \tilde{\lambda}$ (i.e., full evasion is only feasible if the government is soft).

In the full compliance case ($\hat{s} + \varepsilon < \lambda$), the condition $\hat{s} + \varepsilon < \lambda$ becomes

$$\varepsilon < \frac{\lambda - \tilde{\lambda}}{2P} \tag{B.57}$$

Combining equations B.57 and B.56, the full compliance case requires $0 < \varepsilon < \frac{\lambda-\tilde{\lambda}}{2P}$, which is feasible only if $\tilde{\lambda} < \lambda$ (i.e., full compliance is only feasible if the government is tough).

Finally, the condition needed for the existence of the partial evasion case ($\hat{s}-\varepsilon < \lambda < \hat{s}+\varepsilon$) becomes $\varepsilon > \max\left\{\frac{\tilde{\lambda}-\lambda}{2(1-P)}, \frac{\lambda-\tilde{\lambda}}{2P}\right\}$. If $\lambda < \tilde{\lambda}$ it becomes $\varepsilon > \frac{\tilde{\lambda}-\lambda}{2(1-P)}$. If $\tilde{\lambda} < \lambda$, it is $\varepsilon > \frac{\lambda-\tilde{\lambda}}{2P}$.

Summarising the results so far, there are two cases to consider: (1) if the government is soft ($\lambda < \tilde{\lambda}$) the full evasion case arises when the noise is low ($\varepsilon < \frac{\tilde{\lambda}-\lambda}{2(1-P)}$) and the partial evasion one when it is high; and (2) if the government is tough ($\tilde{\lambda} < \lambda$) the full compliance case occurs when the noise is low ($\varepsilon < \frac{\lambda-\tilde{\lambda}}{2P}$) and the partial evasion when it is high.

Hence, using proposition 4.8, the average declaration in each of the two cases is given by

$$\begin{array}{ll} (1) & D^* = \begin{cases} 0 & \text{if} \quad 0 < \varepsilon < \frac{\tilde{\lambda}-\lambda}{2(1-P)} \\ 1 - P + \frac{\tilde{\lambda}-\lambda}{2\varepsilon} & \text{if} \quad \frac{\tilde{\lambda}-\lambda}{2(1-P)} < \varepsilon \end{cases} \\ (2) & D^* = \begin{cases} 1 & \text{if} \quad 0 < \varepsilon < \frac{\lambda-\tilde{\lambda}}{2P} \\ 1 - P + \frac{\tilde{\lambda}-\lambda}{2\varepsilon} & \text{if} \quad \frac{\lambda-\tilde{\lambda}}{2P} < \varepsilon \end{cases} \end{array} \tag{B.58}$$

It is straightforward from here to prove the first part of the proposition by simply computing the derivative of D with respect to ε.

For the second part, using the two cases considered above and proposition 4.10, the expected loss of the agency is as follows

$$\begin{array}{ll} (1) & EL^* = \begin{cases} \gamma\lambda & \text{if} \quad 0 < \varepsilon < \frac{\tilde{\lambda}-\lambda}{2(1-P)} \\ 0 & \text{if} \quad \frac{\tilde{\lambda}-\lambda}{2(1-P)} < \varepsilon \end{cases} \\ (2) & EL^* = (1-\gamma)(1-\lambda) \end{array} \tag{B.59}$$

Deriving with respect to ε yields the result stated in the proposition. ■

Appendix C

Endogenous Externality: Experiment

C.1 Instructions for treatment *GC*

Introduction[1]

First of all, thank you very much for taking part in this experiment. It is important to start by saying that, though part of a serious research programme, this experiment is NOT a test. There are no "right" or "wrong" answers.

How it works

Before we do anything, we have to run through a few ground rules and instructions. After that we will move to the experiment proper, where you will be asked to make decisions in a number of economic situations presented to you. Finally you will get paid: on top of a show-up fee of £5, you will get a sum of money that will depend on your performance in the situations mentioned before.

The experiment consists of 5 stages:

- Instructions
- Trial rounds

[1] Instructions for the other treatments were similar to these ones, with the logical changes in rules and parameters needed in each case.

- Experiment rounds
- Questionnaire
- Payment

We will go through these in detail below.

Ground rules

For the experiment to work we need to run it according to fairly strict rules, but there are not too many:

- From now until the end of the experiment, please do not talk (it will not take long!)
- If there is something you need to ask about the way the experiment works just raise your hand -the experimenter will come to your desk.
- Please do not use the computer until you are told to.

The Six Stages

1 Instructions

The experimenter will read out the instructions. If you have questions, this is the time to deal with them. Just raise your hand and the experimenter will answer them privately.

2 Short quiz

This is to ensure that you understand the instructions.

3 Trial rounds

The experiment is organised in a series of rounds. Each round is a period in which you interact –via the computer only– with the other participants and make decisions that determine the amount of money you will get at the end of the session.

As a warm-up you will first take part in 2 trial rounds. These trial rounds are identical to the experiment rounds in every respect with one exception: the effect on payment. Trial rounds do NOT affect your reward at the end of the experiment. They allow you to check out the interface and familiarise yourself with the screen tables, buttons and commands. They also allow you to make mistakes without losing money.

4 Experiment rounds

This is the real thing. What you do during these rounds will determine the total amount of money you will get.

The following "Frequently Asked Questions" will lead you through the basic mechanics of the rounds.

4.1. What is this all about?

Let us start by saying that the experiment will consist of 20 experiment rounds. In each one of them the computer will pair you up with one other participant. Each of the other participants in the room is equally likely to be paired up with you.

4.2. What do I have to do?

You have to choose one of two possible actions, namely Y or Z. You choose one or the other by clicking on your preferred option in the bottom left panel of the choice screen (see figure C.1) and then pressing the "OK" button in the same panel.

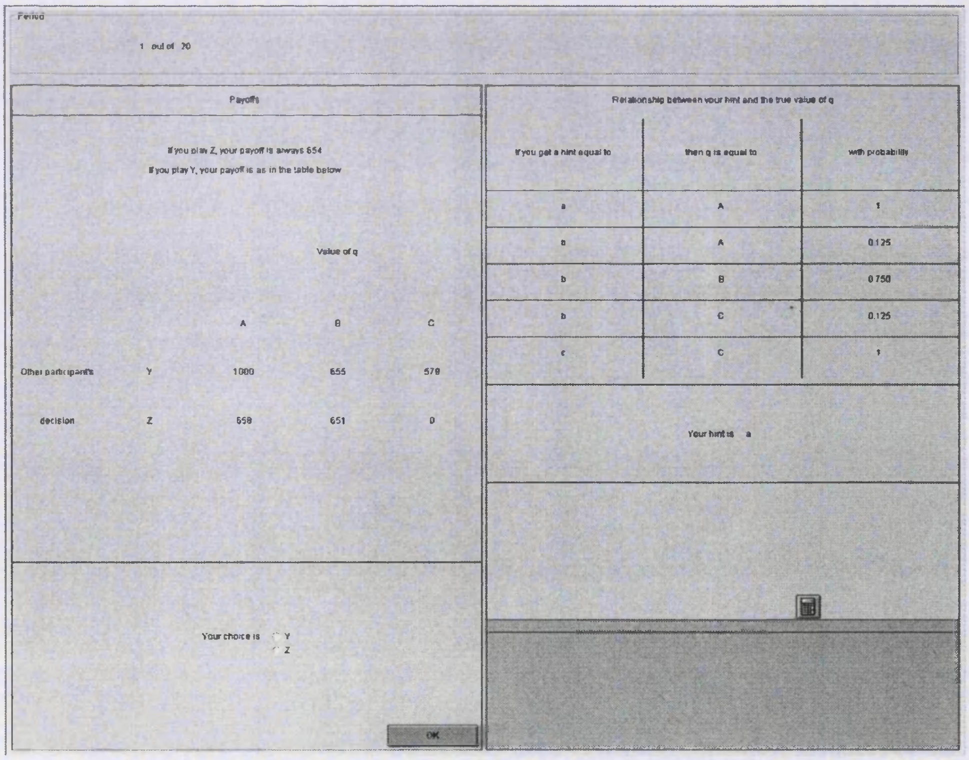

Figure C.1: Choice screen

4.3. How is my payoff for the round determined?

Your payoff for the round depends on your own action, the action of the other participant, and an unknown parameter called q.

4.4. But exactly how *is my payoff for the round determined?*

There are two cases to consider:

a. If you choose action Z, your payoff is 654 "experimental points" with certainty.

b. If you choose action Y, your payoff depends on both the value of q and the action of the other participant, as shown in the table below (and also in the top-left panel of the choice screen (see figure C.1)):

		Value of q		
		A	B	C
Other participant's choice	Y	1000	655	579
	Z	658	651	0

That is, if you choose Z, you always get 654 "experimental points", regardless of what the other participant does and what the value of q is. But if you choose Y, then there are several cases to consider. Let us see some of them (remembering that in all of them you choose Y and your payoff is measured in "experimental points"):

If the other participant chooses Y and q equals A, then your payoff is 1000.

If the other participant chooses Y and q equals B, then your payoff is 655. And so on.

4.5. So how much money do I get then?

Your payoffs are transformed into money at a rate of: 1000 "experimental points" = 50 pence

That is, if your payoff for the round is, for example, 655 "experimental points", your corresponding money earnings are $655 \times 50/1000 = 32.75$pence.

Your session earnings are computed by adding up the money you got during the 20 experiment rounds.

4.6. But, what is q?

q is a parameter that can only take one of 3 values: A, B or C. In any given round, your computer will choose one of these 3 values, with probabilities 0.20, 0.60 and 0.20, respectively.

Intuitively, you can think of these probabilities in the following way: Consider an urn with 100 balls. 20 of them are labelled "A", 60 "B" and 20 "C". The value of q will be determined by the label of one of the 100 balls in the urn, chosen randomly (by the computer).

4.7. Is there anything I could use to make a more informed decision?

Yes, there is. Before you make a decision you will get a "hint". This hint will be known only to you and can only take one of 3 values: a, b or c. It provides some information about the value of the unknown parameter q, as shown in the following table (and in the top-right panel of the choice screen (figure C.1)):

If hint is...	...then q is...	...with probability...
a	A	1.000
	A	0.125
b	B	0.750
	C	0.125
c	C	1.000

For any given round, your hint can be found immediately below this table in the choice screen (figure C.1).

The table may seem a bit complicated but do not worry, it is not. It simply says that if your hint is equal to a, then you can be sure that q is equal to A. Analogously, if your hint is equal to c, then q is equal to C. When your hint is equal to b, however, you do not know for sure what the value of q is, but you can tell how likely each value is: q is equal to B with probability 0.750, while it is equal to A or C with probabilities 0.125 and 0.125, respectively.

Important note: Although q is the same for you and the other participant, your hints may differ from each other.

4.8. Anything else I should know before making my choice?

If you want to make some computations before choosing your action, you can press the calculator button on the choice screen (the small square button just above the darker area (see figure C.1)). Pens and paper are available for those who prefer them: raise your hand and an experimenter will take them to your desk.

Also, it is worth mentioning that there is no "Back" button, so please make your decisions carefully and only press the "OK" or "Continue" buttons when you are sure you want to move to the next screen.

4.9. So I made my decision, what now?

After you submit your decision, you will be shown the action you chose and the payoff you got for the round, as well as the value that q took (see figure C.2). By clicking on the "Continue" button you will move to a new round (if there is any still to be played).

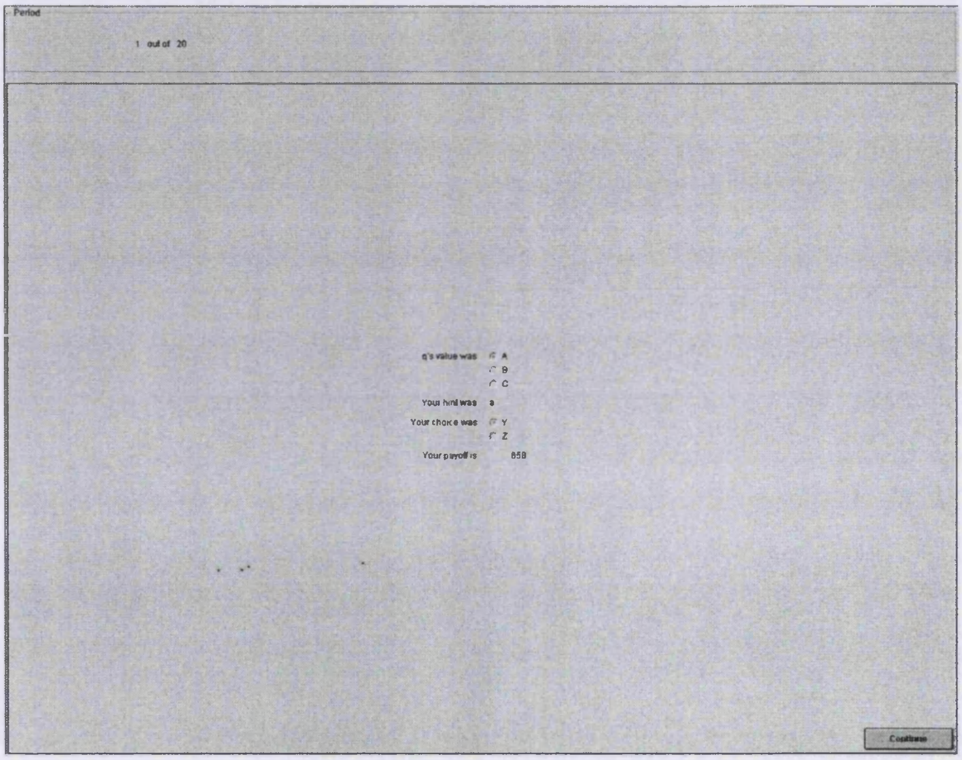

Figure C.2: Results screen

4.10. And then? Is it the same over and over again?

Basically, yes. In every round, the structure is identical to the one described above: first a new q will be selected by the computer and you will be paired up with another participant, then you will be assigned a hint and will have to make a decision, and finally your payoff will be shown on the results screen.

You can check what happened in previous periods by taking a look at the darker area in the bottom-right panel of the choice screen (see figure C.1). It includes information about the values adopted by q, the hints you got and the actions you chose in earlier rounds.

Important note: Every period is like a clean slate: the value of q, the participant you are paired up with and the hint you get may vary from round to round, but the RULES that determine them (explained in questions 4.6., 4.1. and 4.7.) do not. In short, rounds are independent: for example, you can think that in every round a new urn with 100 balls -20 "As", 60 "Bs" and 20 "Cs" - is used to determine the value of q, as explained in question 4.6. Similarly, the pairings and hints of a given round are independent of the pairings and hints of previous rounds.

5 Questionnaire

We will ask you a few questions that will help us to further understand the data collected in the session.

6 Payment

Finally! You will be paid a show-up fee of £5 plus the sum earned during the session, as explained in question 4.5.

And that is it. Once again, thank you very much for participating!

SHORT QUIZ

1. What is your payoff (in "experimental points") if you choose Y, the person paired-up with you chooses Z and q is equal to A? ....

2. What is your payoff (in "experimental points") if you choose Z, the person paired-up with you chooses Y and q is equal to C? ....

3. If your hint is equal to b, what is the probability that q is equal to A? ....

C.2 Equal enforcement cost

In order to test hypothesis 5.2 (on page 71), it is necessary to ensure that the (expected) enforcement cost in both treatments (*GE* and *LE*, or *GC* and *LC*) are the same, so that the comparison is a valid one. This, in turn, requires analysing the problem of the government in each of the 12 possible scenarios: those generated by the combination of the four treatments (GE, LE, GC and LC) and the type of the agency (Soft, Medium and Tough).

The task is greatly simplified, though, thanks to the result mentioned on page 65, namely, that L treatments can be seen as special cases of the G ones.

Before getting into the computations, however, it is important to notice that some actions that in the experiment lead to a fixed, certain payoff, in the real world lead to uncertain ones.

In the experiment, the payoff generated by a given combination of parameters and choices (d, d', q) is certain: for example, if the agency is soft ($q = A$), the taxpayer evades ($d = Y$) and the other taxpayer complies ($d' = Z$), her payoff is $x(Y, Z, A)$ with certainty, and the associated utility is $u(Y, Z, A)$. In real life, however, the payoff $x(Y, Z, A)$ (actually, any payoff other than the *safe* one) is not certain, but rather the *expected* payoff the taxpayer gets when $(d, d', q) = (Y, Z, A)$. Indeed, in such scenario, what actually happens is the following: with probability $a(Y, Z, A)$ the taxpayer is audited and gets utility $u^c(Y)$, and with probability $(1 - a(Y, Z, A))$ she is not audited and gets utility $u^n(Y)$. Thus, utility from payoff $x(Y, Z, A)$ can be defined as

$$u(Y, Z, A) = a(Y, Z, A) \cdot u^c(Y) + (1 - a(Y, Z, A)) \cdot u^n(Y) \tag{C.1}$$

where $u^c(Y)$ is the utility of the taxpayer when she evades and is caught and $u^n(Y)$ is the utility she gets when she evades and is not caught. This means that there is a linear relationship between utility and probability, which can therefore be used to analyse the conditions for equal enforcement cost. In particular, one can define the following relevant probabilities:[2]

- p_q: the probability of detection in L treatments when the agency is of type $q \in Q$;
- ρ_q: the probability of detection in G treatments when the other person complies and the agency is of type $q \in Q$; and
- π_q: the probability of detection in G treatments when the other person evades and the agency is of type $q \in Q$.

Since in G treatments payoffs satisfy equation 5.5, it is necessary that

$$\pi_q < \rho_q \qquad \forall q \in Q \tag{C.2}$$

[2]Since the government will never audit anyone who declares high income (anyone who chooses Z in the experiment), only the probabilities corresponding to low declarations (i.e., Y-choices) are important for the analysis.

That is, the probability of detection for a taxpayer that evades is higher when the other taxpayer complies (ρ_q) than when the other taxpayer evades (π_q). This is exactly the type of relationship expected to exist between these two probabilities when the agency chooses its auditing strategy optimally, as found in chapter 4.

Turning back to the problem of ensuring equal enforcement costs, what is needed is to equalise the expected number of audits, Ea, in G and L treatments, for each possible value of q. Computing the Ea for each of the 12 cases above mentioned yields the following results:

1. *GE*, *Soft*: With probability γ (equal to 1/2 in the experiment) true income is high and both taxpayers get low or medium signals and evade, so $Ea = 2\pi_S$. With probability $1 - \gamma$ (= 1/2) true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_S$ again. The overall (expected) number of audits is therefore

$$Ea_{GE,S} = 2\pi_S \tag{C.3}$$

2. *GE*, *Medium*: With probability $\gamma = 1/2$ true income is high and both taxpayers receive medium signals and evade, so $Ea = 2\pi_M$. With probability $1-\gamma = 1/2$ true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_M$. The overall (expected) number of audits is therefore

$$Ea_{GE,M} = 2\pi_M \tag{C.4}$$

3. *GE*, *Tough*: With probability 1/8 true income is high and both taxpayers receive medium signals and evade, so $Ea = 2\pi_T$. With probability 1/4 true income is high and one taxpayer receives a medium signal and evades while the other receives a tough signal and complies, so that the first one is audited and the second is not: $Ea = \rho_T$. With probability 1/8 true income is high and both taxpayers receive tough signals and comply, so nobody is audited and $Ea = 0$. With probability 1/2 true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_T$. The overall (expected) number of audits is therefore

$$Ea_{GE,T} = \frac{1}{2}\left(\frac{1}{2}\pi_T + \frac{1}{2}\rho_T\right) + \pi_T \tag{C.5}$$

4. *LE*, *Soft*: Since *Lottery* treatments can be interpreted as special cases of the *Global* ones where the probabilities are independent of the other taxpayer's choice, then,

from equation C.3, the overall (expected) number of audits is

$$Ea = 2p_S \tag{C.6}$$

5. *LE*, *Medium*: From equation C.4,

$$Ea = 2p_M \tag{C.7}$$

6. *LE*, *Tough*: From equation C.5,

$$Ea = \frac{1}{2}\left(\frac{1}{2}p_T + \frac{1}{2}p_T\right) + p_T = \frac{3}{2}p_T \tag{C.8}$$

7. *GC*, *Soft*: With probability 1/8 true income is high and both taxpayers receive soft signals and evade, so $Ea = 2\pi_S$. With probability 1/4 true income is high and one taxpayer receives a soft signal and evades while the other receives a medium signal and complies, so that the first one is audited and the second one is not: $Ea = \rho_S$. With probability 1/8 true income is high and both taxpayers receive medium signals and comply, so nobody is audited and $Ea = 0$. With probability 1/2 true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_S$. The overall (expected) number of audits is therefore

$$Ea = \frac{1}{4}(\pi_S + \rho_S) + \pi_S \tag{C.9}$$

8. *GC*, *Medium*: With probability 1/2 true income is high and both taxpayers receive medium signals and comply, so nobody is audited and $Ea = 0$. With probability 1/2 true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_M$. The overall (expected) number of audits is therefore

$$Ea = \pi_M \tag{C.10}$$

9. *GC*, *Tough*: With probability 1/2 true income is high and both agents receive medium or tough signals and comply, so $Ea = 0$. With probability 1/2 true income is low and both taxpayers truthfully declare low income, so $Ea = 2\pi_T$. The overall (expected) number of audits is therefore

$$Ea = \pi_T \tag{C.11}$$

10. *LC*, *Soft*: From equation C.9,

$$Ea = \frac{1}{4}(p_S + p_S) + p_S = \frac{3}{2}p_S \tag{C.12}$$

11. *LC*, *Medium*: From equation C.10,

$$Ea = p_M \tag{C.13}$$

12. *LC*, *Tough*: From equation C.11,

$$Ea = p_T \tag{C.14}$$

The constraints in table 5.5 are therefore obtained by equalising the relevant equations (C.3 and C.6, C.4 and C.7, etc.) and using the relationship between utility and probability mentioned on page 137.

C.3 Extra Tables

	DOMa	DOMb	DOMc	RDOM	ADOM	DOM
LC	.	0.0000	.	1.0000	0.0000	0.0000
LE	0.3231	0.0000	.	0.3180	0.0000	0.0000
GC	.	0.0000	0.1578	0.1552	0.0000	0.0000
GE	.	0.0000	0.0390	0.0418	0.0000	0.0000
LC=GC	.	0.0000	0.1578	0.1552	0.0000	0.0000
LE=GE	0.3231	0.0042	0.0390	0.1920	0.0042	0.0029
LC=LE	0.3231	0.0000	.	0.3180	0.0000	0.0000
GC=GE	.	0.0005	0.2170	0.4359	0.0005	0.0014

Note: Values of F-tests. Values below 5% imply the null hypothesis is rejected.
Dots mean there is no variability in data as to compute the statistics.

Table C.1: Dominance tests. Predictions and inter-treatment comparisons.

	ERR*A*	ERR*B*	ERR*C*	ERR
LC	0.3456	0.0000	0.0000	0.0518
LE	1.0000	0.1450	0.0004	0.3515
GC	0.2939	0.0000	0.0000	0.0147
GE	1.0000	0.145	0.0010	0.2445
LC	0.0000	0.0000	0.0000	0.0000
LE	0.0038	0.0000	1.0000	0.1441
GC	0.0092	0.0000	0.1575	0.0000
GE	0.0006	0.0000	0.2612	0.0092
LC=GC	0.9556	0.0000	0.0000	0.0000
LE=GE	0.1082	0.0000	0.0034	0.0000
LC=LE	0.0000	0.0000	0.0000	0.3555
GC=GE	0.0135	0.0000	0.2160	0.0112

Note: Top panel: Predicted values of dependent variable. Middle and bottom panels: Values of F-tests. Values below 5% imply the null hypothesis is rejected.

Table C.2: Errors tests. Predictions and inter-treatment comparisons.

Dep. Var.:	DOM*a*	DOM*b*	DOM*c*	RDOM	ADOM	DOM
LC		X			X	X
LE		X			X	X
GC		X			X	X
GE		X			X	X
LC=GC		GC			GC	GC
LE=GE						
LC=LE		LE			LE	LE
GC=GE		GE			GE	GE

Note: Top panel: Empty if data fits prediction in hypothesis OS; "X" otherwise. Bottom panel: Empty if no difference, treatment with higher dominance otherwise.

Table C.4: Dominance tests. Predictions and inter-treatment comparisons. Last 10 periods.

Dep. Var.:	DOM*a*	DOM*b*	DOM*c*	RDOM	ADOM	DOM
a				1.0473		0.6725
				[0]		[0]
b						0.5053
						[0]
c				0.9890		0.7641
				[0]		[0]
g	dropped	0.2561	-0.0476	-0.0393	0.2561	0.2024
	[-]	[0]	[0.321]	[0.281]	[0]	[0]
e	dropped	0.4821	0.0000	-0.0310	0.4821	0.3982
	[-]	[0]	[0]	[0.073]	[0]	[0]
ge	dropped	-0.2485	-0.1103	-0.0295	-0.2485	-0.2158
	[-]	[0]	[0.265]	[0.586]	[0]	[0]
cons	1.0000	0.4479	1.0000		0.4479	
	[.]	[0]	[0]		[0]	
Obs.	62	616	82	144	616	760
LC	1.0000	0.4479	1.0000	1.0000	0.4479	0.5500
LE	1.0000	0.9300	1.0000	1.0000	0.9300	0.9300
GC	1.0000	0.7040	0.9524	0.9643	0.7040	0.7722
GE	1.0000	0.9376	0.8421	0.9167	0.9376	0.9333

Note: Top panel: Probability that estimate =0 is shown in brackets below estimate. Bottom panel displays estimated average values of the dependent variable for each treatment.

Table C.3: Estimation. Dominance. Overall and by signal. Last 10 periods.

	OLS				
	GC	GE	LC	LE	All
g					0.2653
					[0]
e					0.3871
					[0]
ge					-0.2935
					[0]
gender	0.0052	-0.1064	0.1704	-0.1118	-0.0785
	[0.958]	[0.039]	[0.059]	[0.002]	[0.005]
age	-0.0068	0.0264	-0.0339	-0.0089	0.0020
	[0.742]	[0.236]	[0.142]	[0.416]	[0.781]
study	0.2118	0.0060	-0.3627	-0.1777	-0.1052
	[0.011]	[0.906]	[0.02]	[0.052]	[0.002]
#exp	-0.0615	-0.0275	0.0619	0.0718	0.0401
	[0.175]	[0.658]	[0.118]	[0.24]	[0.055]
maths	-0.4264	-0.1135	0.4849	0.0345	-0.0750
	[0]	[0.326]	[0.001]	[0.404]	[0.036]
prob	-0.1368	0.0990	0.8721	-0.0770	0.0323
	[0.181]	[0.219]	[0.303]	[0.298]	[0.302]
game	0.3613	-0.0432	-0.0825	0.0729	0.0917
	[0]	[0.722]	[0.218]	[0.005]	[0]
cons	1.3190	0.4526	0.9149	1.2914	0.6329
	[0]	[0.224]	[0.178]	[0]	[0]
Obs.	231	177	251	239	898

Note: Probability that estimate =0 is shown in brackets below estimate.

Table C.6: Estimation. Effect of personal characteristics on choices. Age<26.

	DOM*a*	DOM*b*	DOM*c*	RDOM	ADOM	DOM
LC	.	0.0000	.	.	0.0000	0.0000
LE	.	0.0006	.	.	0.0006	0.0007
GC	.	0.0000	0.3207	0.3171	0.0000	0.0000
GE	.	0.0021	0.0694	0.0766	0.0021	0.0004
LC=GC	.	0.0000	0.3207	0.3171	0.0000	0.0001
LE=GE	.	0.7927	0.0694	0.0766	0.7927	0.6365
LC=LE	.	0.0000	.	.	0.0000	0.0000
GC=GE	.	0.0000	0.2645	0.4187	0.0000	0.0000

Note: Values of F-tests. Values below 5% imply the null hypothesis is rejected.
Dots mean there is no variability in data as to compute the statistics.

Table C.5: Dominance tests. Predictions and inter-treatment comparisons. Last 10 periods.

C.4 Examples of categories

Expected-Payoff Maximisers (EPM): "If the hint was a, I selected action Y; otherwise, I selected action Z. There are only three outcomes that generate more than 654 points, and two of them only generate a negligible increase (relative to their risk). The only way to "gamble and win" is to play Y when the hint is b, and in that case, I am gambling that either my "opponent" has a hint of a (very unlikely), or my opponent has a hint of b, is risking that q is really A, and is right (also very unlikely). My risk is that my opponent plays Z, which is safer, and that q is B or C, which is likely. The risk/reward is far too high. When my hint is a or c the correct play is obvious - in the former case, playing Y always nets me more than 654, and in the latter, playing Y always nets me less than 654, no matter what my opponent does." (Subject #10, GC).

Chance Maximisers (CM): "If the hint is c, the best decision is always Z with a higher payoff. If the hint is b, it worths choosing Y, because thereis a probablility of 0.875 getting A or B, which are both higher than Z(654). If the hint is a, my decision is definitely Y." (Subject #20, LC).

Learners (L): "At first i played it safe and went with the guarantee button z and then i took more of a risk by chosing the y button every time i got the hint "a" or "b". because there was a higher probability of gaining more points." (Subject #18, GE).

Mixers/Experimenters (M/E): "If the hint was A, choice was Y. If the hint was C, choice was Z. If the hint was B, 80% of the time choice was Y and 20%, B." (Subject #15, LE).

Non-independent (NI): "If the hint came up as A i always selected choice Y as I would be better off (ie gaining more money) through doing so regardless of what the other participant chose. Conversely, if the value of q was C i always chose Z since I would be worse off if i choice Y despite what the other person selected. If the value of q came up as b i would go systematically throught the choices Y,Y,Z. This was my order since if q=b and q=a i would be better off selecting Y and if q=c i would be better off selecting Z. Since the probability of q=b was the highest i put Y at the beginning of the order. I used my knowledge of maths and probablities to calculate the order in which to place my choices." (Subject #2, GC).

Randomisers (R): "If the hint was a then i chose Y if the hint would have been c then i would have chosen Z. apart from this i just guessed randomly. the last 3 i thought i may as well take the risk as it was the end of the experiment." (Subject #19, LC).

Confused (C): "If the probability was lower than the other option, i chose the other option. I did not take risks in the cases where the probability could also go for the lowest amount. Becasue i dont know much about the probability theory so i decided to go for the safest method." (Subject #7, LC).

Risk-lovers (RL): "I chose Y every time unless I knew it was C. I was not given the hint a at any time. The difference between playing it safe and gambling with the Y option was small enough to make the experiment slightly more fun. I knew that I could lose 579, but only gain 421, but preferred the gamble." (Subject #7, GC).