Tax Audit Rules and Firm Behaviour

Tax Audit Rules and Firm Behaviour

David John Hoey

THE UNIVERSITY of ADELAIDE

A thesis submitted to the University of Adelaide in fulfillment of the requirements for the degree of Doctor of Philosophy

October 2015

Declaration

I, David John Hoey certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree. I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968. I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Acknowledgements

I would like to thank my supervisor Professor Ralph Bayer who has been immensely generous with his encouragement, advice and support. I cannot thank him enough. I acknowledge those who helped with the running of the experiments used in this thesis: Sookie Zhang, Mickey Chan and Alex Jarvis. I thank Pru Kerr, especially for her comments on my writing. I thank all those who have inspired my learning over the years, including Professor Ralph Bayer (again), Dr Virginie Masson, Dr Nadya Baryshnikova, Dr Eran Binenbaum, Professor Mark Weder, Dr Jacob Wong, Ian Graham and Dr Seungmoon Choi. Also, I thank everyone in my coursework cohort including Dr Anne Arnold, Dr Mark Dodd, Dr Brita Pekarsky, Sumit Dutta, Dr Kofi Otumawu-Apreku, Niti Rathod and Dr Hang Wu.

Finally, I would like to thank my family and friends—Mum, Dad and Dr Anne Hoey for everything they have done for me. Also, Nick Birch and Brian Monger for their explanations of comparative advantage and sunk cost! Last but definitely not least, my beautiful partner Carlyn Leaver. She has been unfailing in her love and support. I genuinely cannot thank her enough.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis we study the effects of tax audit policy on firm behaviour. We study firm behaviour in two dimensions: (1) the firm's tax reporting behaviour, and (2) the firm's market behaviour. We are interested in the extent to which firms try to evade tax by concealing income when they face different audit policies; and how firms make, for example, pricing or production decisions when they face different audit policies. We examine how tax audit policy affects: the way firm's compete, how they report income for tax purpos and how efficient the market is.

In Chapters 2, 3 and 4, we investigate the effects of audit policy on firm behaviour different ways. Chapters 2 and 3 are theoretical and Chapter 4 is empirical. Chapter 2 looks at how different audit policies affect the way firms behave when they are operating under different market structures. We consider the effects of tax audit policy, not just in a standard Cournot or Bertrand setting, but in markets where firms interact in other ways as well. To obtain general results, we categorize market structures based on the way firms interact. We then we use examples to illustrate how they work in more specific cases. Some of our examples include markets where firms produce complement goods, markets where firm's compete vertically in a supply chain, and markets where firm's compete by choosing an advertising strategy. In Chapter 3, we look at the effect of an audit policy that is at the opposite extreme to the one most commonly used in the literature. Audit policy is commonly modeled by assuming that the firms each face a fixed probability of being selected for audit[1][2]. This means the probability of being selected is completely independent of the tax reporting behaviour of the different firms competing in the market. We consider the opposite extreme where, not only does audit selection depend on the reporting behaviour of each firm, but it does so deterministically. We model the tax reporting game as a rank-order tournament—or more specifically, a tournament where the firm that reports the lowest income is audited with certainty. The reason for studying this type of tax audit policy is to understand the way firms might behave when the tax authority's best strategy is simply to target the firm that signals the weakest information (the lowest declaration of income), and focus all of their audit resources on this one firm[3]. In Chapter 4, we take the model to the data. We use data from a laboratory experiment to understand how human subjects behave when faced with different audit policies. We compare this behaviour with the behaviour of rational firms as predicted by the model.

In the rest of this introduction we do the following. Firstly in Section 1, we provide an overview of the topic of this thesis. In particular, we explain what is meant by "tax auditing", "tax evasion" and "firm behaviour" in this thesis. Then, in Section 2 we provide some motivation for why we study this topic. In Section 3 we describe the general model used throughout this thesis. This allows us then in Section 4 to explain our program of work for the remaining

[1] We discuss the literature and develop an argument to support this in Section 2.1

[2] An alternative interpretation of the fixed rule is that firms face an audit where the tax authority allocates a fixed amount of resources to the audit.

[3] We discuss this further in Section 2.2

chapters.

1.1 Tax auditing, tax evasion and firm behaviour

Since the whole thesis deals with the same subject matter, just in different ways, we now explain what this subject matter is in more detail. Specifically, we explain what is meant here by: "tax audit policy", "corporate tax evasion" and "market behaviour". The aim is to create some context and make these ideas a little more concrete. We take each in turn.

1.1.1 Tax auditing

According to the IRS, a tax audit is "a review/examination of an organization's or individual's accounts and financial information to ensure information is being reported correctly, according to the tax laws, to verify the amount of tax reported is accurate." (IRS website). Since we will be studying *corporate* tax evasion, the "organization's or individual's" mentioned in the quote, will be firms. A **tax audit** for our purposes will mean (from the quote): *a review or examination of a firm's accounts to verify that their information has been reported correctly.*

There are different ways to select firms for audit. One of the simpler way is to select firms uniformly at random according to a fixed probability. In this case, each firm faces the same audit probability regardless of what they, or other firms in the industry, report in their tax returns. An alternative selection method is to make the probability of audit conditional on information that is available to the tax authority. For example, information about items from past or present tax returns and information about past compliance behaviour can be used by the tax authority to determine which firms should be chosen to receive an audit. Conditioning on this information allows the tax authority to make a more sophisticated assessment of compliance risk for each firm, and allows the tax authority to recover (ideally) a greater portion of unpaid taxes.

The terms **audit policy** and **audit rule** will be used synonymously throughout this thesis and will refer to the *functional form of the detection probability—the probability that the firm is audited and that any unreported income is found out*. There are two interpretations of these audit rules. An audit rule can be thought of as the probability that a particular firm is audited and their unreported income is found out, or alternatively, it can be thought of as the *faction* of resources allocated for auditing a particular firm and therefore the *fraction* of underreported income that is found out. Either interpretation is fine.

Audit policy in practice, requires decisions to be made about other things besides how audit resources are allocated between firms, or which firms should be audited. It also requires for example: decisions to be made about how resources should be allocated between different jobs within an audit, or about what information the tax authority should make available to firms. Since we do not deal with this however, when we refer to "audit policy" we mean specifically, which audit rule is used. An "audit policy" or an "audit rule" will be the functional form of the probability that the tax fraud of a particular firm will be detected. Again, in much of the economic literature, the audit rule or audit probability is modeled simply as a constant[4]. We pit against this fixed audit probability some more sophisticated audit rules, where the probability of audit is endogenous and depends on the reporting behaviour of firms.

1.1.2 Tax evasion

There is a distinction between tax evasion and tax avoidance. Tax avoidance is the use of *legal* means to reduce the amount of tax payable. We deal only with tax evasion in this thesis—which is *illegal*. Tax evasion is defined by the IRS as

[4]We discuss the literature in 2.1

follows. In section 7201 of the Internal Revenue Code there are two offenses (as outlined by the IRS tax crimes handbook): (a) the willful attempt to evade or defeat the assessment of a tax, and (b) the willful attempt to evade or defeat the payment of a tax. We will deal with the first kind only—part (a). For the purpose of this thesis we assume that once the tax authority understands how much tax is owed by a firm (by conducting an audit), they have no problem in recovering this amount (i.e. we do not deal with part (b)). We model tax evasion by allowing firms to declare less income to the tax authority than what they have truly earned. **Tax evasion** is then: *where firms declare less income than what they earn and they escape audit and their under-reporting is not found out.*

Understanding tax evasion is important because it means there is a redistribution of resources within the economy. Those that evade tax benefit from doing so. Those that do not evade tax however carry a higher proportion of the tax burden. This redistribution may well be considered unfair because those that behave honestly must bear the burden. Also, since we are dealing with *corporate* tax evasion, and since firms that evade tax are the ones to gain, it means that shareholders, managers and owners of the firms are the ones that benefit. The redistribution then is likely to be regressive[5]. To get an idea about the size of the redistribution, some tax authorities publish estimates of the "tax gap". The tax gap is the amount of tax liability faced by taxpayers that is not paid on time. In the US the most recent estimate of the tax gap was $450bn in 2006; $67bn of which was due to under reporting by corporations and another $4bn due to underpayment by corporations[6]. The opportunity cost of this forgone tax revenue is therefore very large. It might include, for example, the value

[5]Since people in these groups tend to be quite well-off on average. Interestingly, there is some evidence that the evasion of personal income tax however, rises with income but actually at a less than proportionate rate. See Christian (1994).

[6]http://www.irs.gov/uac/The-Tax-Gap

of a quite considerable increase in the provision of public goods or a reduction in the tax burden for more honest tax payers.

1.1.3 Firm behaviour

As we have already mentioned, we study the effect of audit policy on firms rather than individuals. The difference between the tax compliance problem for firms as opposed to individual taxpayers is two-fold. Firstly, the internal organisation is much different. Firms have many stakeholders and many decision makers, whereas the individual taxpayer, is clearly, just an individual. Although we look at two different types of organisations in Chapter 4 in the laboratory (owner-controlled firms and CEO/CFO-controlled firms), we do not explicitly model these differences theoretically. For the majority of the thesis, each firm is assumed to behave as a single rational decision maker. Secondly however, the nature of the the firm's market activity is different to that of the individual. An individual faces a labour supply decision, whereas a firm must make complicated decisions about, pricing, marketing, production, and about which projects to invest in. The different types of decisions that firms may face and how these decisions affect other firms (as we will see in Chapter 2) have important implications for the effectiveness of different audit rules.

As already mentioned, we examine two aspects of firms' behaviour: their tax behaviour; and their market behaviour. Tax behaviour, since corporate income tax is levied on profits, will be modeled by the firm's decision about how much profit to declare to the tax authority. Market behaviour on the other hand, will be modeled by a choice variable describing the firm's decision related to the way they compete in the market. For example, firms might need to choose how much output to produce (as they do in a Cournot oligopoly), or what price to set (as in a Bertrand oligopoly), or how much to invest (in advertising or R&D

for example). The action that describes the firms' market behaviour, for the majority of this thesis, will be the choice of what quantity of goods to produce since the Cournot framework will be convenient for Chapter 3 and 4. In Chapter 2 however, it will depend on the type of market we are discussing.

1.2 Motivation

Why do we examine different audit policies? Why examine audit policies that differ to the simple constant probability case? We discuss two reasons. Firstly, the theoretical literature now presents a good rationale for why audit policy might affect firm behaviour. If audit policy affects firm behaviour then different groups of people will be affected, including: consumers, managers, policy makers and other corporate stakeholders. We would like to examine audit policy to better understand how it affects firms—and especially for how it affects these different groups of people.

Secondly, audit policies that differ from the simple constant probability case are important because they capture the fact that easily available information can be used to condition audit selection to reduce the chances of unreported income going undetected. It is not realistic to assume that the tax authority would not take this information into account. We argue that although sometimes tax reports may be chosen unconditionally, audit selection is most often conditional on at least some type of information. We address each reason in turn.

1.2.1 Audit policy affects firm behaviour

Audit policy, we argue, may affect the way firms behave. It may affect firm behaviour, not just in terms of the reporting decisions that they make, but also affect the market-based decisions that firms make. If audit policy can affect firms' market behaviour, it follows that audit policy may in turn have important implications for consumers, corporate decision makers and policy makers. Understanding the effects of firm behaviour and the implications of audit policy for these different groups is our first motivation for studying this topic.

For some time the literature on corporate tax evasion did not provide a good rationale for how or why audit policy might affect the way firms behave in the market. The formal literature on tax evasion extends back to Allingham and Sandmo's (1972) seminal paper. In models that were developed from this framework, taxpayers are depicted as gamblers and they face a fixed probability of audit. They choose how much income to declare by weighing up the prospect of successfully evading tax with the prospect of receiving an audit and having to pay a penalty if they are caught. Tax evasion then, boils down to a simple decision theory problem under uncertainty. Subsequent work extended the model to corporate settings where firms decide simultaneously how much to evade as well as how much output to produce (see for example Virmani (1989) where firms face perfect competition; Marrelli & Martina (1988) where firms compete in a duopoly; and Wang & Conant (1988),Marrelli (1984), Yaniv (1996) and Lee (1998) for the case of monopoly). A result that is very prominent in this literature is the separability of evasion decisions and output decisions when the audit probability is fixed (constant). This means that when firms face a fixed a probability of audit, the presence of taxation in the model does not affect their output decisions. Firms still produce the Cournot quantity in a Cournot environment, monopolists still produce the monopoly quantity and competitive firms still produce at the minimum average cost level. The result is robust to different types of taxes as well (for example the sales tax in Marrelli (1984); profit tax in Wang & Conant (1988); and withholding tax in Yaniv (1988, 1995)).

The separation result is not robust to other types of audit rules however.

Various attempts were made to endogenize the audit rule, mostly by specifying it as a function of the firm's own profit, revenue, cost, output level or on the amount that the firm over- or understates on their tax report. Since these rules still depend only on the firm's own actions though, they are not strategic. One paper by Marrelli & Martina (1988), examines an audit rule that makes the problem strategic. Their audit rule is a function of the choice variables of both firms in the market and it boils down to a function of the differences in the amount understated on their returns. The separation result fails to hold when this more sophisticated specification of the audit probability is assumed[7]. Assuming an audit rule that is a function of the amount that is over- or understated however, as Lee (1998) points out, does not appear to be a very reasonable assumption. If the tax authority can condition the audit probability on the understated or overstated amount, then it must have this information ex-ante and the purpose of the audit then is not really clear.

(Bayer and Cowell 2009, 2006, 2010) consider a similar model but with the more reasonable assumption that the audit probability is a function of the declared profit of each firm in the industry. Under this assumption, using a standard Cournot or Bertrand framework, Bayer and Cowell show that when the audit probability is decreasing in the firm's declared profit and increasing in the profits of other firms, firms behave more competitively. More importantly, they provide a clear rationale for *why* firms behave more competitively: they would like to earn more profit relative to their competitors so that they can declare relatively more profit and reduce their audit probability[8]. Since in the literature,

[7]There are other specifications of the audit rule where taxes are no longer neutral, and where it depends on the tax base and how the penalty is formulated. If the expected penalty is not based on the tax base, the separation result fails to hold. SeeLee (1998) for a good summary for the monopolistic case.

[8]Another difference is that the game is played in two stages so that production decisions are made first, profits are observable to the firms (or at least more observable to other firms within the industry than to the tax authority), and then tax decisions are made. The reasoning for modeling the problem as a two-stage game is to capture the information asymmetry between the firms and the tax authority—firms within the industry, it is argued, are often better

it is now clear that audit policy can affect firm decision making in the market, and that there is a good rationale for *why* it affects firm decision-making in the market, it makes sense to study, and tease out the implications of this in more detail. This is exactly what we aim to do.

1.2.2 Utilizing reported information can reduce tax evasion

Our second motivation for studying the effect of different audit rules is that we would like to capture a more realistic approximation of the tax authority's real-life behaviour. The aim is to bring the theoretical model closer to the compliance problem in practice, so that, hopefully, we can obtain more relevant insights into the way firms might be expected to behave (assuming rationality). Using audit rules that take into account reported information, we argue, is a more sophisticated approach to modeling audit selection since they capture the fact that easily available information can be used to condition audit selection to reduce the chances of unreported income going undetected. Although sometimes a tax authority might select firms using a fixed, unconditional probability, audit selection is most often dependent on certain types of reported information. The tax authority can use this information in a variety of ways. We discuss these now.

Firstly, the tax authority can use document matching to match reported information with payor records. If information in a taxpayer's report cannot be verified by third party documentation, then an audit may be required. Information that does not match the information already available to the tax authority will raise a "red-flag", as it suggests the taxpayer will need to be investigated. Also, related examinations may occur when returns selected for audit involve

informed about the market conditions and profits than the authority.

issues or transactions with other taxpayers. Reported information that links the taxpayer with reports by other taxpayers that are in dispute may also raise red-flags and need examination.

Another way audit selection can depend on reported information is that it may be conditional on the history of reporting behaviour of the taxpayer. The tax authority can use information about the taxpayer's past reporting behaviour to determine whether or not they should conduct an audit. A taxpayer that has been consistently non-compliant in the past would then face a much higher audit probability than a taxpayer with a better track record. Similarly, if the taxpayer's current reports have been shown to be non-compliant, then past reports might be audited retrospectively. In this case, taxpayers that have been non-compliant in the current period would be more likely to receive an audit of their previous reports.

Still another method of selecting firms for audit is by comparing data from the taxpayer's report with data from the reports of other firms in the same industry. The tax authority can use discriminant function analysis to determine which reports deviate furthest from the average reports of the group[9]. A Discriminant Index Function (DIF) score can be calculated to assess the non-compliance risk of each taxpayer. A higher DIF score means that there is a higher likelihood that their return will need further examination. Exactly how these DIF scores are calculated is not typically made apparent—such information would benefit those seeking to evade tax. But the idea is very simple: those that deviate furthest from the average reporting behaviour of the group will be much more likely to receive an audit. The reporting decision thus becomes a game between tax payers. We can model this situation in two ways: as a coordination problem or as a prisoner's dilemma.

If we model the problem as a coordination game, let us assume there are

[9]See U.S. Government Accountability Office 1999

two firms and each can opt to be honest (by reporting their true income) or dishonest (by reporting low income—in an attempt to evade some tax). Let us assume that if both firms choose to report their true income, they break even and receive a payoff of zero. We normalize their payoff to zero. However, if they both choose to report low income, they both succeed in evading tax because neither is deviating from the average—they receive a positive amount, y. If one firm reports low income and the other high, the firm that reports high receives the break even amount, 0, since they are being honest, whereas the firm that reports low income and deviated below the average, receives an audit. This firm must pay all of the tax owing and a fine for non compliance. They receive the break even amount, 0, less a fine, x, for trying to evade. In matrix form, we can write the game as follows:

	Report high income	Report low income
Report high income	$0, 0$	$0, -x$
Report low income	$-x, 0$	y, y

Table 1.1: Tax evasion as a coordination game. Firms receive: 0 if honest; $-x$ if found out; and y if they can successfully evade.

Notice here, there are two pure-strategy equilibria ($highincome, highincome$) and ($lowincome, lowincome$). If firms' have the ability to coordinate on the ($lowincome, lowincome$) equilibrium, then the audit rule has no power to reduce evasion. It may not always be possible to coordinate however. Alm & McKee (2004) examine tax evasion as a coordination game in the laboratory. We discuss their results in Chapter 4, especially as they relate to our experiment. They find that for some treatments, collusion on the low declarations is not observed. This means that audit rules that create a coordination problem for the firms may be beneficial.

Suppose however, that if both firms report the same income, rather than having no audit, each firm receives a 50-50 split of the audit resources. This could mean that each firm receives an audit with a 50% probability; or maybe each firm has half of their accounts audited and half of their under reported income is discovered. In either case, the (expected) payoffs for two firms that report low income become $\frac{y-x}{2}$ and $\frac{y-x}{2}$. The tax evasion game is now a prisoner's dilemma if the penalty for evading exceeds the reward for successfully evading, $x > y$. The 2x2 game is:

	Report high income	Report low income
Report high income	$0, 0$	$0, -x$
Report low income	$-x, 0$	$\frac{y-x}{2}, \frac{y-x}{2}$

Table 1.2: Tax evasion as a prisoner's dilemma game.

This time, there is only one equilibrium, $(highincome, highincome)$. Under this audit rule, we have no more tax evasion! We examine audit rules in this thesis that are similar to this in that they also create a prisoner's dilemma-type situation between firms. It means that each firm has an incentive to report more income than their competitor. Firms compete away potential gains from successfully evading tax and they are more honest than if audit selection were simply random and unconditional. Although the audit rules we examine are not nearly as complex as the DIF analysis used by tax authorities, they aim to capture this strategic aspect of the problem which is missing when the audit rule is modeled as a fixed, unconditional probability.

1.3 The compliance problem

We now explain the baseline model since we use the same framework throughout the whole thesis. We then use it, in Section 4, to describe the work that follows in the remaining chapters.

1.3.1 The Bayer-Cowell model

The framework that we use throughout this thesis builds on the familiar story of the compliance problem as seen in the literature, except that firms play in two-stages (as in Bayer and Cowell (2006, 2009 & 2010)). The first stage is the *market stage*, where firms compete strategically in some form of (usually oligopolistic) competition. Bayer and Cowell (2009), for example, consider the case where firms compete in Cournot competition. Bayer and Cowell (2006) look also at Bertrand competition with differentiated goods. In Chapter 2 we examine the effect of tax audit policy in different types of markets, beyond just these two cases. In Chapter 3 and 4, we revert back to a simple Cournot framework.

The second stage of the game is the tax stage. After competing in the first stage, firms must submit a declaration of the profit they have earned. Based on this declaration, the amount of tax they must pay is determined. There is also a chance that they may be audited and forced to pay a fine if they are caught declaring less than what they earned in an effort to evade tax. The firm's payoff is as follows:

$$EU_i = \pi_i(\mathbf{x}) - td_i - \beta_i(\mathbf{d})[f + t](\pi_i(\mathbf{x}) - d_i) - C(\pi_i(\mathbf{x}) - d)$$

The firm's payoff in words, is the sum of: the firm's gross profit earned from the first stage, $\pi_i(\mathbf{x})$; less the firm's tax bill, td_i; less the expected penalty if caught trying to evade tax, $\beta_i(\mathbf{d})[f+t](\pi_i(\mathbf{x})-d_i)$; less the cost of concealment, $C(\pi_i(\mathbf{x}) - d)$. We will discuss each of these in turn.

The firm's gross profit, $\pi_i(\mathbf{x})$, represents the firm's payoff from the first stage of the game. It is a function of the vector of actions taken by firms in the first stage, $\mathbf{x} = (x_1, x_2, ..., x_i, ..., x_n)$. For example, in Bayer and Cowell (2009), where the first stage is a Cournot oligopoly, $\mathbf{x}$ represents the vector of quantities chosen by firms. Similarly in Bertrand competition, $\mathbf{x}$ represents the vector of prices chosen by each firm. This vector of prices or quantities determines the gross profit each firm receives.

The second term in the firm's payoff is the firm's tax bill, td_i. It is simply the tax rate, t, multiplied by the firm's profit declaration, d_i. We assume that firms face a simple constant tax rate which is levied on gross profit. Firms can choose to declare less than what they actually earn from the first stage (that is, choose $d_i < \pi_i$), in an effort to evade tax and reduce their tax bill. However there is a chance that they will be selected for an audit, which means that any undeclared profit will be found out. This brings us to the third term.

The third term is the expected penalty if the firm is selected for an audit and found to have undeclared profit. We assume that if a firm is audited, the audit will be successful in finding any concealed profit and the firm will be forced to pay the penalty. The expected penalty is a function of the following terms. It depends on the audit probability, $\beta_i(\mathbf{d})$, the fine, f, the tax rate (again), t, and the amount of profit that the firm has failed to declare, $(\pi_i(\mathbf{x}) - d_i)$. The functional form of the audit probability, β_i is what we are interested in in this thesis. The audit probability may be a function of the vector of profit declarations, $\mathbf{d} = (d_1, d_2, ..., d_i, ..., d_n)$, made by the firms, but it does not have to be. We are interested in how the functional form affects the way firms behave. There are three different functional forms for β_i that we examine. We will describe each of them quite loosely for the moment—but they will be more formally defined in the coming Chapters. In addition to the audit probability, the expected penalty depends on: the fine, f, which is a constant; the amount of profit concealed, $(\pi_i(\mathbf{x}) - d_i)$; as well as the amount of tax still owed by the firm, $t(\pi_i(\mathbf{x}) - d_i)$.

The last term is the cost of concealing profit, $C(\pi_i(\mathbf{x}) - d)$, and this too is a function of the amount of profit concealed. This might represent for example, the cost of hiring ("creative") accountants to help conceal profit, or the cost of rearranging business structures in order to conceal profit in a way such that it is not immediately obvious to the tax authority.

1.3.2 Audit rules

We will be studying the effect of different audit rules on firm behaviour. The audit probability, $\beta_i(\mathbf{d})$, will take on one of three different functional forms and these different functional forms, we call *audit rules.* The first audit rule is where β_i is simply a constant. As has been mentioned, this is probably the most common way to model audit selection in the tax evasion literature and we refer to it as a *fixed audit rule.* The other two types of audit rules we study are functions of the declarations made by the firms. The first is the *relative audit rule.* The relative rule is where the probability that the firm gets selected for an audit is higher for firms that declare less relative to their competitors. This means that if one firm declares a low profit and their competitors declare a high profit, there is a greater chance that this first firm will receive a tax audit. The reverse is true also: if one firm declares a high profit and their competitors declare low profit, there is a lower chance that this firm will receive a tax audit. The relative rule may be a good policy tool for the tax authority when it believes that higher profit declarations are associated with a lower chance that the firm is evading tax. We examine the relative rule, first in Chapter 2, in relation to how it affects firms competing under different market structures; and then again in Chapter 4, using data from our laboratory experiment. The third audit rule we examine is the rank order rule. The rank order rule is where the firm that declares the lowest amount is the only firm that is audited. We look at this in Chapter 3 and 4. Now that we have our model, we are ready to summarize the remaining chapters of this thesis.

1.4 Program of work

Chapter 2

In Chapter 2 we will examine the effects of audit policy on firm behaviour for different types of markets. In the context of the model, this means we examine the effects of the relative and fixed audit rules for different types of underlying market games (i.e. when $\pi_i(\mathbf{x})$ is different depending on the type of market firms are competing in). Our main result says that the effect of a relative audit rule compared with a fixed audit rule on market behaviour, depends on the types of externalities that firms' actions have on each other and how these actions affect market efficiency.

Firstly, if in the first stage competition is not strategic, and the actions of firms' do not affect their competitors, then the choice of audit policy has no effect on market decisions and the market remains unaffected, whether a relative rule is in place or not. Secondly, in markets where a firm's actions have a negative externality on the profits of their competitors, firms choose more of this action under a relative rule. The converse is also true. In markets where firms' actions have a positive externality on their competitors' profits, the relative rule means that firms choose a lower action than if a fixed rule were in place. Depending on how the audit policy affects firms' market behaviour and depending on whether their actions have a positive or negative effect on the efficiency of the market, we find that welfare could be higher or lower under a relative rule.

Since these results are quite abstract, we present some examples to illustrate. For example, monopoly competition satisfies the requirements of the first result which means that monopolists behave the same in the market under either audit rule. Cournot and Bertrand with substitute goods, as Bayer and Cowell (2009) demonstrate, are markets where the relative rule encourages competitiveness and these markets are more efficient than under a fixed rule. However, the relative rule can create incentives for destructive or uncompetitive behaviour too. For example, in markets where one firm has the option to "sabotage" or reduce the profits of other firms, even at some small cost to their own profits, firms behave less efficiently under the relative rule than under the fixed rule. Also, examples of markets that are less competitive and more inefficient under a relative rule include markets where firms operate vertically in a supply chain or markets where firms produce complement goods.

Chapter 3

In Chapter 3 we examine an audit policy where the firm that declares the least profit is selected for audit. If the tax declarations carry at least some information about which firms are more likely to be evading tax, then the best audit policy may simply be to audit the most "at-risk" firm—the firm that signals the weakest information and reports the lowest declaration. This is the rank-order audit rule.

We discuss how the rank-order audit rule is the limiting case of the relative rule. When the probability of a particular firm being selected for audit is more sensitive to changes in declaration, the relative audit rule seems to have a larger effect on market behaviour. This was shown by Bayer and Cowell (2010) in simulations for a simple linear Cournot model. However, it is not clear a priori, what the effect of the limiting case would be, since at the limit the audit rule becomes discrete. We examine the model with the rank-order rule and see that there are more equilibria than under a relative rule. In particular, even though firms are completely symmetric, there is asymmetry in the pure strategy equilibrium, since one firm produces more compared with under a fixed rule, while the other firm produces less. There are also mixed strategy equilibria where firms randomize between these high and low amounts.

Chapter 4

In Chapter 4 we take the model to the data. We use data from a laboratory experiment that looks at the behaviour of subjects when they face a relative audit rule or a fixed audit rule.

We have three treatments—two with a relative rule and one with a fixed audit rule. The two relative rule treatments include one where subjects are each put in charge of their own firm (owner-controlled firms) and one where the tax decision and market decision is delegated to different subjects (CEO/CFO firms). So the CEO/CFO firms are each comprised of 2 subjects whereas the owner-controlled firms are comprised of 1 subject each. The aim of comparing variation between these two treatments is to see whether there is any difference in tax behaviour and market behaviour when the internal organisation of the firm differs. The third treatment is where a fixed audit rule is in place and firms are comprised of single decision makers (owner-controlled firms). Comparing variation between the fixed and relative rule treatments where the firms are owner-controlled allows us to identify any differences in behaviour due to the different audit rules.

We find from our data that tax evasion is significantly lower under the relative rule than under the fixed rule. In both the relative rule treatments, the amount of tax evaded is significantly less than in the fixed treatment. The results for market behaviour are not so clear. Of the two treatments with a relative rule, in one we observe greater output (meaning the market is more efficient) than in the fixed rule treatment; and in the other we observe less output (a less efficient market) under a relative rule. Finally, in the CEO/CFO treatment with shared decision making, since the incentives for each subject are perfectly aligned, we should see no difference between this treatment and owner-manager firms. However, in the CEO/CFO treatment, firms produce significantly *more* than in the equivalent treatment with owner-managers. This is a surprising result since it would be easier for the single decision maker to see how their choice in the market stage affects their ability to compete in the tax stage. Despite this the owner-managers are significantly better at coordinating on lower output.

Chapter 5

In Chapter 5 we conclude. We summarize the work and results of Chapters 2-4 and discuss some practical implications that might be drawn from our work. In addition, we discuss some of the limitations of the analysis in order to clarify the interpretation of our results. We finish with a discussion of some possible directions for future work.

Chapter 2

Market Structure and Relative Tax Auditing

2.1 Introduction

In this chapter we examine tax audit policy in the context of different types of markets. It seems reasonable that firms might react to various audit policies differently depending on what type of market they are in. Bayer & Cowell (2006, 2009 & 2010) show that audit policy can affect firms' behaviour in two dimensions. It can affect (1) how firms report income for tax purposes, and (2) it can affect how firms make production or pricing decisions in the market. However, they consider audit policy in the context of only two types of markets. One is a 'regular' Bertrand oligopoly with differentiated products and the other is a 'regular' Cournot oligopoly[1].

The aim of this chapter is first to abstract from the exact market structure

[1] These are described in Vives (1999) in Chapter 4 and 5 respectively. The regular Cournot oligopoliy is based on assumptions used by Novshek (1985) in his existence result; we discuss these in more detail below.

used in the model. Rather than assuming a specific market structure, we show how audit policy affects the market depending on properties of the underlying market game. This allows us to demonstrate how the compliance problem works in a more general sense and we see that the effect of audit policy can be understood in terms of the different types of interaction that firms might face. Bayer and Cowell (2006, 2009 & 2010) show that certain audit policies have a positive effect on market efficiency. However, we see that when these audit policies are applied to other types of markets, they may actually have a negative effect on market efficiency (and in some cases—no effect). To illustrate how the results work, in addition to the Cournot and Bertrand cases, we present examples of other types of markets. These examples help to develop intuition for why we see different effects, since they show how the general results apply in different contexts.

We consider two different audit policies. One is where a fixed audit rule is imposed on the market, and the other is where a relative audit rule is imposed on the market . The *fixed audit rule* is where firms are selected for an audit based on a fixed probability. This is a simple way to model audit selection because the audit probability is simply a constant. This is very common in the literature. A *relative audit rule* on the other hand, is endogenous since it depends on the declaration made each firm. A relative rule is where the probability that a particular firm is selected for audit depends on: (1) the information the firm themselves report in their tax declaration; and (2) how this information compares with information that other firms report. When a profit tax is in place, a relative rule means that the probability of audit is conditioned on how much profit the firm reports relative to how much profit other firms report. If one particular firm reports less than the others, it might be viewed as an indication that they are trying to evade tax. So the tax authority would like to impose a greater audit probability on this low declaring firm. Firms declaring weaker information (lower profit) then, under a relative rule, are targeted over the firms declaring stronger information (higher profit) and they are penalized with a greater chance of detection. It is a more sophisticated way to model audit selection than simply assuming a constant probability, as we argued in Chapter 1 since it captures the fact that information, reported directly to the tax authority, can quite easily be used by the tax authority to improve their chances of detecting unreported profit. We contrast the relative audit rule with the fixed rule because we want to understand how the two compare when they are imposed in different types of markets.

Our main result is that the effect of a relative rule compared with a fixed rule is different in different types of markets and the effect depends on the externalities that firms' actions have on the profits of their competitors. Under the relative rule, firms gain an advantage in the tax stage if they can declare more profit relative to their competitors because it means they will face a lower audit probability. To be able to declare relatively more profit, they must *earn* relatively more profit. Obviously, if a firm wishes to increase the relative size of their profit, they would have to either: (1) try to earn more profit themselves; or (2) try to reduce their competitors' profits. The incentive for (1) already exists without a relative rule because even when a relative rule is not in place, a higher profit still means a higher payoff. The incentive for (2) however—is new. Without a relative rule firms have no incentive to try to reduce their competitors' profits. But under a relative rule, by choosing an action that reduces their competitors' profits, the relative size of their own profit increases and they can declare relatively more than their competition which reduces their audit probability. This second incentive can lead to firms choosing different actions to what they would under a fixed rule, and, as we will see, it does not always lead to socially desirable outcomes.

Once we have our general results, we present some examples. Applying our results to different models of competition demonstrates three important points. Firstly, although Bayer and Cowell find that firms are more competitive under a relative rule (in the Cournot and differentiated Bertrand case), this result is not robust to other models of competition. We present examples where firms interact differently, and their actions have different externalities on their competitors. In markets where being more competitive hurts other firms and has a negative effect on their profits (as in the differentiated Bertrand or Cournot case), the relative rule increases market efficiency—which means there is a positive welfare effect. However in other markets, where firms must behave less competitively to reduce the profits of their competitors, (examples include, markets with vertical competition, advertising, or markets for complementary products), the relative rule reduces market efficiency—a negative welfare effect. We also see the separation result that is so common in the literature on corporate tax evasion—but only for cases where firms do not interact strategically. This means that when firms do not interact strategically in the market, there is a separation between the market decisions and tax decisions that firms make. This separation result says that choices firms have about their tax declaration do not affect their choices about what to do in the market. We see that even when the audit rule is endogenous and depends on the reporting behaviour of all firms, the separation result still holds in markets where there is no strategic interaction, such as monopoly or perfect competition.

Secondly, the relative rule may create an incentive for firms to "sabotage" each other to try to reduce each other's profits. A relative audit rule might be a good policy choice in an environment in which firms' only means of reducing other firms' profit is to behave more competitively. However, in an environment where firms can more easily hurt their competition by choosing actions that are uncompetitive or even “malicious”, a relative audit rule might not be such a good policy choice. We present an example where firms have the ability to choose a costly action that negatively affects the profits of competing firms, but otherwise offers no benefit to the firm undertaking the action. Under a fixed rule, firms do not choose this action. Under a relative rule however—they do.

Finally, under a relative rule, as Bayer and Cowell show, the fraction of tax evaded is lower than under a fixed rule. Nevertheless, under a relative rule, the amount of tax revenue collected, may actually be lower *in absolute terms.* This is because under a relative rule, if firms make less profit than they would under a fixed rule, it means their tax base is lower. So if firms are paying a higher fraction of profit but on a lower tax base, the effect on the level of tax revenue is ambiguous—it may be higher or lower (or the same). In many of our examples, it is indeed lower. This last point highlights a possible incentive problem that a tax authority might face. If they are required to meet revenue targets, the tax authority may lack the incentive to keep a relative audit policy in place (or adopt one if they have not already) since it may yield them less revenue[2].

The plan of the rest of this chapter is as follows. In Section 2 we describe the model and define our two audit rules. In Section 3 we characterize the interior solution for the model. In Section 4 we present our main results that explain how the effect of the relative audit rule depends on general properties of the type of interaction firms face. In Section 5 we present some examples to help illustrate the effects of the relative rule when it is applied to more specific models of competition, and in Section 6 we conclude.

[2]In practice the most prominent measure of performance made public is the “tax gap”. It is possible that internally, a revenue target might be an objective of the tax authority however.

2.2 Setup

This section is in two parts. In Section 2.1 we describe the model. We describe the two stages of the game—the market stage and then the tax stage. Then, we state the firm's payoffs and discuss each of its components. In the second part, Section 2.2, we formally define our two audit rules and then explain how they affect the firm's problem. Following that, Section 3 looks at first, the timing of the model and the firm's problem at each stage, and then we solve the model[3].

2.2.1 The Model

The game is played by $i = 1, ..., n$ firms in two stages. The first of the two stages, is called the *market stage*. Firms in the market stage compete by choosing a generic choice variable we call their "action", x_i. For a Bertrand market, x_i could be thought of as a price. Similarly, for a Cournot market, x_i could be a quantity. The vector of actions chosen by the firms is denoted $\mathbf{x} = (x_1, x_2, ..., x_i, ..., x_n)$ and this vector maps into a function describing the gross profit made by each firm, $\pi_i^g(\mathbf{x})$. For now, we keep the gross profit function as general as possible. Precise functional forms for $\pi_i^g(\mathbf{x})$ are introduced in Section 5 where we present our examples to show how the results for the general case work.

The second stage is called the *tax stage*. Firms in this stage decide on how much profit to declare to the tax authority. Each firm has an opportunity to evade tax by declaring less profit than what they truly earned from the market stage. Alternatively, firms might choose to be honest and declare all of their gross profit. The firm's declared profit is denoted d_i which we assume to be no larger than their actual profit. That is, $d_i \leq \pi_i^g(\mathbf{x})$.[4] Firms face a linear tax

[3]This section and the next follows, basically, Bayer and Cowell (2009)

[4]In practice, many individuals, when filing their personal income tax declarations, actually overstate their income. For the case of corporate tax payers this may be less likely, but in any case, we are modeling their behaviour as rational decision makers. This means we rule out the possibility of firms declaring more than they earn, $d_i > \pi_i^g(\mathbf{x})$. No rational firm would

function, t, which is levied on the tax base, which we take to be gross profit. They also pay a fine, f, if the tax authority chooses to conduct an audit and they have failed to declare all of their gross profit earnings. The fine is proportional to the amount of undeclared gross profit. To keep the exposition simple, assume that the gross profits and the declarations are non-negative.[5]

We now explain the firm's payoffs. The representative firm's payoffs are as follows and they are made up of four components:

$$E\pi_i = \beta_i(\mathbf{d})\underline{\pi}_i(d_i, \mathbf{x}) + [1 - \beta_i(\mathbf{d})]\,\overline{\pi}_i(d_i, \mathbf{x}) - C(\pi_i^g, -d)$$

The payoffs consist of: (1) the firm's payoff if no audit takes place, $\overline{\pi}_i(d_i, \mathbf{x})$, (2) the payoff if an audit does take place, $\underline{\pi}_i(d_i, \mathbf{x})$, (3) the evasion cost, $C(\pi_i^g, -d)$, and (4) the probability of an audit, $\beta_i(\mathbf{d})$. We discuss each in turn.

Firstly, the firm's payoff if no audit takes place. The amount of tax paid by the firm depends on whether the firm receives an audit. If there is no audit, the representative firm pays tax only on the amount of profit declared and there is no penalty for evading tax since the firms actions go undetected. The firms tax bill in this case is td_i. The firm's net profit if an audit does not take place then, is simply gross profit less the tax bill:

$$\overline{\pi}_i(d_i, \mathbf{x}) = \pi_i^g(\mathbf{x}) - td_i$$

Secondly, the firm's payoff if an audit does take place. If an audit does take place, and if the firm has declared less than $\pi_i^g(\mathbf{x})$, it must pay a fine in addition to their tax liability. The fine, is proportional to the amount of profit that the

choose to pay tax on more income than they need to.

[5]Without this restriction, but assuming still that there are no concessions to firms that make a loss or declare a negative profit, the tax owed would become $\max\{0, t\pi_i\}$, and tax paid would become $\max\{0, td_i\}$. This would make the exposition a bit more cumbersome but does not change anything in what follows.

firm did not declare—so the fine is $f\left[\pi_i^g(\mathbf{x}) - d_i\right]$. This time, tax is payable on the full gross profit amount since the true amount of profit is known once an audit is conducted. The firm is liable for tax owed on this whole amount—the tax bill in this case is $t\pi_i^g(\mathbf{x})$. So if an audit does take place, the firm's profit net of tax, $t\pi_i^g(\mathbf{x})$, and the fine, $f\left[\pi_i^g(\mathbf{x}) - d_i\right]$ is:

$$\underline{\pi}_i(d_i, \mathbf{x}) = \pi_i^g(\mathbf{x}) - t\pi_i^g(\mathbf{x}) - f\left[\pi_i^g(\mathbf{x}) - d_i\right]$$

Thirdly, the evasion cost. In addition to the profit tax and the fine, firms face a cost, C, if they wish to conceal profit. This cost might include the cost of hiring accountants to help conceal profit, the cost of researching the best ways to evade tax or the cost of restructuring activities that might be needed to facilitate evasion. We do not impose a specific functional form on C however, but we do impose some restrictions on C. We have: $C' > 0$; $C'' > 0$; $C(\pi_i^g - d_i) \geq 0$, (where the primes denote derivatives); and also $C(\pi_i^g - d_i) = 0$ for $d_i = \pi_i^g$. We explain what these restrictions mean in order. The first two conditions mean that the cost of concealing profit increases with the amount of profit concealed, and that the marginal cost of concealing profit increases with the amount of profit concealed. This means that it is more costly to conceal more profit and, also, the cost of concealing profit increases as more profit is concealed. The third restriction says that the evasion cost is non-negative (that is, the cost is always positive and for no values does it increase the firm's payoff). The last restriction says that there is no cost to the firm if they choose to declare truthfully so that their profit declaration is equal to their actual profit, $d_i = \pi_i^g$.

Finally, whether the firm is audited by the tax authority depends on an *audit rule*. An audit rule is denoted $\beta_i(\mathbf{d})$, and it is the probability that firm i is audited by the tax authority. It is a function of the vector of profit declarations, $\mathbf{d} = (d_1, d_2, ..., d_i, ..., d_n)$ made by firms in the industry. Much of the literature assumes that the audit probability is simply a constant, or when it is made conditional, it is conditional only on the firm's own actions (see the discussion in Chapter 1, 2.1). A conditional audit rule, as discussed in Chapter 1, seems to be a richer and more realistic way to model audit selection. It captures the fact that the tax authority has the ability to use information from taxpayer reports to target taxpayers that are more likely to be trying to evade tax. The information declared can be thought of as a signal of the strength of the firm's true profit which can be compared with the information declared by other firms. We therefore consider the case where audit selection is a function of both the firm's own declaration as well as the declarations of all other firms in the market. This is a *relative audit rule*. We compare the relative audit rule with the case where firms face a *fixed audit rule*. A *fixed audit rule* is where the probability of an audit is constant for each firm. We will define these formally below.

We are now ready to interpret the payoff function as a whole. Combining the terms just described, for convenience, we restate the firm's payoff function:

$$E\pi_i = \beta_i(\mathbf{d})\underline{\pi}_i(d_i, \mathbf{x}) + [1 - \beta_i(\mathbf{d})]\,\overline{\pi}_i(d_i, \mathbf{x}) - C(\pi_i^g, -d)$$

In words, it says that the expected payoff is a function of the expected value of the firm's profit if they receive an audit, plus the expected value of the firm's profit if they do not receive an audit, less the cost of concealment. For ease of exposition, assume that firms are risk neutral. The assumption of risk neutral preferences is strong, but since the qualitative results hold through for the risk averse case it is not crucial for anything that follows [6]. We set risk preferences aside for the moment. In the remainder of this section we define our two audit rules and describe them in a little more detail. The fixed rule is first and the relative rule is second.

[6] Risk preferences are discussed in the concluding chapter, see Chapter 5.

2.2.2 Audit Rules

The two audit rules used in this chapter are defined as follows:

Definition 1. A *fixed audit rule*, β_i^F is a constant $\in [0,1]$ for any i and for all d_i.

Under a fixed audit rule the second stage decision will be non-strategic in that firms need not consider the behaviour of rival firms. We will see that their payoffs are independent of their rivals' declarations. The firm simply weighs up the benefit of evading tax (that is, $t(\pi_i^g - d_i)$) with the prospect of having to pay the fine $(\beta_i^F f(\Pi_i - d_i))$. Under a relative rule this will not be the case. Let $\mathbf{d} \equiv (d_1, ..., d_n)$ be the vector of firm declarations.

Definition 2. A *relative audit rule*, $\beta_i^R(\mathbf{d})$, is a function that satisfies the following conditions whenever $0 < \beta_i^R(\mathbf{d}) < 1$ for all $i \epsilon N$

- $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i} < 0 \ \forall i \epsilon N$
- $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j} > 0 \ \forall j \neq i \epsilon N$
- $\sum_{i=1}^n \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j} = 0 \ \forall j \epsilon N$

The first two conditions of the relative rule say that the probability that the tax authority will audit a particular firm is decreasing in its own declaration and increasing as rival firms increase their declarations. A firm that declares a higher profit is rewarded with a lower probability of being audited. Also, if a particular firm declares a higher profit it means that the other firms have a higher probability of being audited. The last condition says that for any given change in a particular firm's declaration, the sum of all the changes in audit probabilities is zero. This condition is so that the resources required to implement the audit rule are kept constant; and where the tax authority has a fixed budget, it is not in danger of exceeding this as declarations change. Since the relative rule is conditional on the declarations of all firms, the tax decision will be strategic. A firm's declaration decision will affect not only its own payoff, but also the payoffs of other firms through the change in audit probabilities.

2.3 Equilibrium

In this section, we describe the timing of the model and present the firm's problem. We solve the model by looking at the second stage first (the tax stage) and then we work backwards and solve the first stage (the market stage). Since the solution to the fixed rule case is nested within the relative rule case, we solve the more complicated relative rule case first. The fixed rule case can be found by setting all of the derivatives of β_i to zero, since β_i then is just a constant.

2.3.1 Timing

The timing of the model is as follows:

- Firms learn what the tax policy is including the tax rate, the penalty for under reporting and what the audit rule will be.
- Firms compete in the market stage. They choose their action (for example, quantity or price). Gross profits are realized.
- Firms observe the gross profits of all firms in the industry and then choose how much to declare.

- Nature selects which firms to audit and the final payoffs are received.

2.3.2 Tax Stage

The firm's problem in the second stage is:

$$\max_{d_i}\{\beta_i^R(\mathbf{d})\underline{\pi}_i(d_i,\mathbf{x}) + \left[1-\beta_i^R(\mathbf{d})\right]\overline{\pi}_i(d_i,\mathbf{x}) - C(\pi_i^g - d)\}$$

Since there is no reward for over-compliance no firm will declare more than the gross profit earned in the first stage. Also assume there are no concessions for declaring a loss, so no firm would choose to declare a negative amount. We have, $d_i \leq \pi_i^g(\mathbf{x})$ and $d_i \geq 0$. Solving the firm's second stage problem, the first order condition for an interior solution for firm i is:

$$\frac{\partial E\pi_i}{\partial d_i} = \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i}\left[\underline{\pi}_i - \overline{\pi}_i\right] + \beta_i^R(\mathbf{d})\left[\frac{\partial \underline{\pi}_i}{\partial d_i} - \frac{\partial \overline{\pi}_i}{\partial d_i}\right] + \frac{\partial \overline{\pi}_i}{\partial d_i} - \frac{\partial C_i}{\partial d_i} = 0$$

Differentiating again, the second order condition is:

$$\frac{\partial^2 E\pi_i}{\partial d_i^2} = \frac{\partial^2 \beta_i^R(\mathbf{d})}{\partial d_i^2}\left[\underline{\pi}_i - \overline{\pi}_i\right] + 2\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i}\left[\frac{\partial \underline{\pi}_i}{\partial d_i} - \frac{\partial \overline{\pi}_i}{\partial d_i}\right] + \frac{\partial^2 C_i}{\partial d_i^2} < 0$$

Since we know that $\left[\underline{\pi}_i - \overline{\pi}_i\right] < 0$, $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i} < 0$, $\left[\frac{\partial \underline{\pi}_i}{\partial d_i} - \frac{\partial \overline{\pi}_i}{\partial d_i}\right] = f + t > 0$, and $\frac{\partial^2 C_i}{\partial d_i^2} = C''>0$, the second order condition holds when:

$$\frac{\partial^2 \beta_i^R(\mathbf{d})}{\partial d_i^2} \geq 0$$

To interpret the second-stage first-order condition notice that $\frac{\partial \underline{\pi}_i}{\partial d_i} = -t$, $\frac{\partial \overline{\pi}_i}{\partial d_i} =$

f, and $\frac{\partial C_i}{\partial d_i} = -C_i'$. With some rearranging we get:

$$\underbrace{t - \beta_i^R(\mathbf{d})\,[f+t]}_{Return(>0)} + \underbrace{\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i}\,[\overline{\pi}_i - \underline{\pi}_i]}_{Shading(<0)} = \underbrace{C_i'(\pi_i^g - d)}_{Cost(>0)}$$

The marginal return to declaring slightly less profit is given by the first term. This is the tax saving less the potential penalty and tax if the firm is audited. The marginal cost of declaring slightly less profit is represented by the term on the right hand side. The shading term in the middle represents the reduction in the marginal return due to the change in audit probability that occurs with a slight increase in the amount of profit concealed.

The first order and second order conditions for the fixed rule case can be found by setting the derivative of β_i with respect to d_i to zero since under a fixed rule, β_i is constant. They are, respectively:

$$t - \beta_i^F\,[f+t] = C_i'(\pi_i^g - d)$$

and,

$$\frac{\partial^2 E\pi_i}{\partial d_i^2} = \frac{\partial^2 C_i}{\partial d_i^2} < 0$$

2.3.3 Market Stage

In the first stage the firms choose an action, which can be thought of as, for example, a price or a quantity of goods to produce. The choice of x_i will affect the audit probability in the second stage because the vector of actions, $\mathbf{x} = (x_1, x_2, ..., x_i, ..., x_n)$, determines all of the gross profit amounts and therefore how much each firm can declare to the tax authority. We will also refer later to the vector of all actions except for the action of firm i, which will be denoted $\mathbf{x}_{-i}$. The firm must choose x_i to maximize the expected payoff that now depends, not only on gross profit, but also on the declaration decision from the second stage. The firm's first stage problem becomes:

$$\max_{q_i}\{\beta_i^R(\mathbf{d}^*)\underline{\pi}_i(d_i^*,\mathbf{x}) + \left[1-\beta_i^R(\mathbf{d}^*)\right]\overline{\pi}_i(d_i^*,\mathbf{x}) - C(\pi_i^g, d_i^*)\}$$

The first order condition for the firm's first stage problem is:

$$\frac{\partial E\pi_i}{\partial x_i} = \frac{\partial \beta_i^R(\mathbf{d})}{\partial x_i}\left[\underline{\pi}_i - \overline{\pi}_i\right] + \beta_i^R(\mathbf{d}^*)\left[\frac{\partial \underline{\pi}_i^*}{\partial x_i} - \frac{\partial \overline{\pi}_i^*}{\partial x_i}\right] + \frac{\partial \overline{\pi}_i^*}{\partial x_i} - \frac{\partial C_i^*}{\partial x_i} = 0$$

With a bit of manipulation, we can rearrange this by solving the firm's first order condition from the second stage for C', and by noticing that:

$$\frac{\partial C_i^*(\pi_i^g(\mathbf{x}) - d_i^*(\mathbf{x})}{\partial x_i} = C'(\pi_i^g - d_i^*(x_i))\left[\frac{\partial \pi_i^g(\mathbf{x})}{\partial x_i} - \frac{\partial d_i^*(\mathbf{x})}{\partial x_i}\right]$$

Also, the change in the audit probability from a change in firm i's quantity includes the sum of changes due to all of the firms' optimal declarations changing in the second stage. That is:

$$\frac{\partial \beta_i^R(\mathbf{d})}{\partial x_i} = \sum_{j\neq i}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j^*}\frac{\partial d_j^*}{\partial x_i} + \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*}\frac{\partial d_i^*}{\partial x_i}$$

Substituting in the first order condition from the first stage we get:

$$\sum_{j\neq i}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j^*}\frac{\partial d_j^*}{\partial x_i}\left[\underline{\pi}_i - \overline{\pi}_i\right] + \frac{\partial \pi_i^g}{\partial x_i}\left[1 - t + \left[\underline{\pi}_i - \overline{\pi}_i\right]\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*}\right] = 0$$

This, together with the firms' first order condition for the declaration stage, characterizes the interior equilibrium. In next section we show, by signing the terms in this first order condition, what effect the relative rule has on competition in the market stage for different assumptions about the underlying market structure that firms face. For the fixed rule case, again setting the derivatives of β_i to zero, the first order condition is simply:

$$\frac{\partial \pi_i^g}{\partial x_i}\left[1 - t\right] = 0$$

2.4 The Market Effect

We argue in this section that the effect of the relative rule compared with the fixed rule on market behaviour, depends on the way firms interact in the market. More specifically, it depends on the type of externality their actions have on the profits of their competitors. We formalize this idea, first for individual firms and then for the market as a whole. We also explain the intuition behind these two results and then examine how this affects welfare. The effect of the two audit policies on market behaviour, as well as the implications of this behaviour for welfare, determine whether the relative rule has a positive effect in real terms, or a negative effect or no effect. We call this effect on welfare the "market effect".

The "market effect", we define as the difference in welfare in equilibrium when firms face a relative rule compared with when they face a fixed rule. In a "regular" Cournot model, if firms face a relative audit rule, their best response is to produce a quantity greater than the Cournot quantity (Bayer & Cowell 2006, 2009 and 2010). Similarly, in a "regular" Bertrand model with differentiated products, equilibrium prices under a relative rule are less than the Bertrand price (Bayer & Cowell 2006). A greater quantity in Cournot competition or lower prices in Bertrand means the market is more efficient (positive market effect). Markets are not always more efficient under a relative rule however. There are other types of markets where the effect is negative or neutral—it depends on what type of market the rule is applied to (as we will see!).

Firstly, we compare firm *behaviour* under the two audit rules. Consider the effect of a fixed audit rule. Under a fixed rule, firms always face the same probability of audit. They choose an action that maximizes gross profit in the market stage and this action does not affect the declaration decision since they have the same probability of receiving an audit regardless of how much profit they can report or how much other firms can report. The first order condition for the market stage in the fixed case is $\frac{\partial \pi_i^g}{\partial x_i}[1-t]=0$. This means firms choose x_i such that $\frac{\partial \pi_i^g(\mathbf{x})}{\partial x_i}=0$. Notice that this is the same action they would choose if there were no tax stage. Without a tax stage, firms simply maximize gross profit and also choose an action x_i such that $\frac{\partial \pi_i^g(\mathbf{x})}{\partial x_i}=0$. So in the fixed rule case, the tax decision does not affect the market decision—the two decisions are separate and the firm's market behaviour is the same. We have the familiar separation result between tax behaviour and market behaviour.

If the relative audit rule is in place however, firms may behave differently than they would under a fixed rule (or if there were no tax stage). They may have an incentive to choose an action which does not maximize gross profit if it gives them an advantage in the tax stage. This is because they can reduce their audit probability if they declare a higher amount relative to the amount that other firms declare. To achieve this they need to increase the relative size of their profit, which they can do by choosing actions that reduce their compeitors' profits. Firms therefore have an incentive to choose an action that "hurts" the profits of their competitors even if it means they sacrifice some of their own profit, as long as their competitors end up worse off than they do[7].

In order to understand the effect of the audit rules however, first we must establish that there exists an equilibrium in the underlying market game and that it is unique. This gives us a unique point of reference so that we can compare the game once we add on the tax stage. Assume that the underlying market

[7]Note though a marginal deviation from the optimal gross profit action does not actually sacrifice any profit since $\frac{\partial \pi_{\mathbf{i}}^{\mathbf{g}}(\mathbf{x})}{\partial \mathbf{x_i}}=0$.

game can be described by $(X_i, \pi_i^g, i\epsilon N)$, where: X_i is firm i's strategy space of the underlying game; π_i^g (again) is firm i's gross profit and payoff from the underlying game; and N, is again the set of players (firms). The best reply map of this underlying game will be denoted by $BR(\mathbf{x}) \equiv (BR_1(\cdot), BR_2(\cdot), \ldots, BR_n(\cdot))$ and assume we have the usual normed vector space, (X, d), where $X \equiv \prod_{i=1}^{n} X_i$ and d, is the Euclidean norm. Lemma 1 provides sufficient conditions for the existence of a unique equilibrium.

Lemma 1. *Assume, for the underlying market game, $(X_i, \pi_i^g, i\epsilon N)$, (X, d) is a non-empty, complete metric space, and the best reply map, $BR : X \to X$, is a contraction of this space. The underlying market game has a unique equilibrium.*

Proof. The result follows immediately from the contraction mapping principle (Banach (1922)). $BR : X \to X$ is a contraction of (X, d) if for some constant $c\epsilon[0, 1)$, $d(BR(\mathbf{x}), BR(\mathbf{y})) \leq cd(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x}, \mathbf{y}\epsilon X$. Consider for any point $\mathbf{x}_0 \epsilon X$, there is a sequence of points, $\{\mathbf{x}_k\}$ defined by $\mathbf{x}_{k+1} = BR(\mathbf{x}_k)$. First, we show first that, by iterating BR to get each new point, the distance between each point contracts so that c is raised to higher and higher powers. Secondly, since the distance between each point gets smaller for each iteration, the sequence is Cauchy and by the completeness of (X, d), its limit must be contained in the space. Finally, because BR is a contraction mapping, it is continuous and therefore both $\mathbf{x_{k+1}} \to \mathbf{x}^*$(where $\mathbf{x}^*$ is our equilibrium) and $BR(\mathbf{x}_k) \to BR(\mathbf{x})$. This gives us a fixed point. Uniqueness follows because a limit is unique.

By definition and applying the contraction inequality:

$$d(\mathbf{x}_k, \mathbf{x}_{k+1}) = d(BR(\mathbf{x}_{k-1}), BR(\mathbf{x}_k)) \leq cd(\mathbf{x}_{k-1}, \mathbf{x}_k)$$

Doing this over and over:

$$d(\mathbf{x}_k, \mathbf{x}_{k+1}) \leq cd(\mathbf{x}_{k-1}, \mathbf{x}_k) \leq c(cd(\mathbf{x}_{k-2}, \mathbf{x}_{k-1})) \leq c(c^2 d(\mathbf{x}_{k-3}, \mathbf{x}_{k-2}) \ldots$$

We get $d(\mathbf{x}_k, \mathbf{x}_{k+1}) \leq c^k d(\mathbf{x}_0, \mathbf{x}_1)$. Using this and the triangle inequality, and then the sum of a geometric series, for some $j > l$:

$$\begin{aligned}
d(\mathbf{x}_j, \mathbf{x}_l) &\leq d(\mathbf{x}_l, \mathbf{x}_{l+1}) + d(\mathbf{x}_{l+1}, \mathbf{x}_{l+2}) + \ldots + d(\mathbf{x}_{j-1}, \mathbf{x}_j) \\
&\leq c^l d(\mathbf{x}_0, \mathbf{x}_1) + c^{l+1} d(\mathbf{x}_0, \mathbf{x}_1) + \ldots + c^{j-1} d(\mathbf{x}_0, \mathbf{x}_1) \\
&= (c^l + c^{l+1} + \ldots + c^{j-1}) d(\mathbf{x}_0, \mathbf{x}_1) \\
&\leq (c^l + c^{l+1} + \ldots) d(\mathbf{x}_0, \mathbf{x}_1) \\
&= c^l \left(\frac{1}{1-c} \right) d(\mathbf{x}_0, \mathbf{x}_1)
\end{aligned}$$

For any $\varepsilon > 0$ then we can find a sufficiently large J, where $j > l \geq J$, such that:

$$d(\mathbf{x}_j, \mathbf{x}_l) \leq c^l \left(\frac{1}{1-c} \right) d(\mathbf{x}_0, \mathbf{x}_1) \leq c^J \left(\frac{1}{1-c} \right) d(\mathbf{x}_0, \mathbf{x}_1) < \varepsilon$$

This means our sequence is Cauchy. Since our (X, d) is complete, it converges within the space. Suppose $\mathbf{x}*$is the limit of our sequence ($\lim_{j \to \infty} \mathbf{x}_j = \mathbf{x}*$) then by continuity of the contraction mapping, $\lim_{j \to \infty} BR(\mathbf{x}_j) = BR(\mathbf{x}*)$. Since $BR(\mathbf{x}_j) = \mathbf{x}_{j+1}$ for all j, the $\lim_{j \to \infty} BR(\mathbf{x}_j) = \mathbf{x}*$ also. Uniqueness follows from the uniqueness of limits. We have a unique fixed point of the best response mapping and thus a unique equilibrium, $\mathbf{x}*$. $\square$

Lemma 1 is sufficient for our purposes; both for our results below and for our examples. The theory that follows can be applied to a larger class of games however. For example, in Cournot, the contraction condition is quite restrictive. It implies for example, when X_i is one dimensional, that $\frac{\partial^2 \pi_i^g}{\partial x_i^2} + \sum_{j \neq i} \left| \frac{\partial^2 \pi_i^g}{\partial x_i \partial x_j} \right| < 0$. This means that the firm's problem is concave and the cross effects are such that firms do not react too strongly to their competitors changing their actions. For n-player games with decreasing best replies, the Bamon & Frayssé (1985)-Novshek (1985) existence result means the contraction condition can be dropped and requires only that the best replies have a slope > -1[8]. The Bamon/Fraysse/Novshek result exploits Selten's (1970) idea that if each player's best response is a function only of the aggregate action of the other players', a fixed point can be found for a cumulative best response that maps the aggregate action of all players into an optimal action for each player. Existence can be established assuming the strategy space is compact and the best replies are upper hemicontinuous and strongly decreasing[9]. Uniqueness then, is achieved by imposing the weaker condition that the best replies have a slope > -1, since they are only a function of the aggregate action. The contraction assumption on the other hand (together with completeness), allows for the best responses to be upward or downward sloping and does not require them to be a function of the aggregate action[10].

Our next result is an re-statement of Bayer-Cowell's (2009) Lemma 3 but for all possible cases. Rather than stating the effect of the relative rule in terms of the properties of a specific Cournot or Bertrand market, we derive the result based on our generic underlying market game and show how the effect depends on the type of interaction that firms face. We now obtain conditions for which the best response under the relative rule, $BR_i^R(\mathbf{x_{-i}}^*)$, is greater than, equal to, or less than the value that is chosen in the fixed rule case, $BR_i^F(\mathbf{x_{-i}}^*)$.

Lemma 2. *For the underlying market game $(X_i, \pi_i^g, i\epsilon N)$ stated in Lemma 1:*

if $\frac{\partial \pi_j^g}{\partial x_i} > 0 \ \forall j \ i \neq j$, *then* $BR_i^R(\mathbf{x_{-i}}) < BR_i^F(\mathbf{x_{-i}})$;

if $\frac{\partial \pi_j^g}{\partial x_i} = 0 \ \forall j \ i \neq j$, *then* $BR_i^R(\mathbf{x_{-i}}) = BR_i^F(\mathbf{x_{-i}})$; *and*

if $\frac{\partial \pi_j^g}{\partial x_i} < 0 \ \forall j \ i \neq j$, *then* $BR_i^R(\mathbf{x_{-i}}) > BR_i^F(\mathbf{x_{-i}})$.

[8] See Vives (2001), page 42 for a discussion on this

[9] A correspondence is strongly decreasing iff all selections are decreasing

[10] Other relevant existence arguments include: Nash's original existence result Nash et al. (1950); those based on supermodularity (Tarski et al. (1955)); and for symmetric games where best replies never jump down and depend on the aggregate action of rivals, Tarski's intersection point theorem (see Tarski et al. (1955)) Theorem 3 and Vives (1990)). Other uniqueness arguments include: Gale & Nikaido (1965) "univalence approach" Theorem 4; or the "index theory approach" based on the Poincaré–Hopf theorem (see Vives (2001) 2.5 for a discussion of this).

Proof. Rearranging the firm's first order condition we have:

$$\frac{\partial \pi_i^g}{\partial x_i} = -\frac{\sum_{j\neq i}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j^*}\frac{\partial d_j^*}{\partial x_i}\left[\underline{\pi}_i - \overline{\pi}_i\right]}{\left[1 - t + \left[\underline{\pi}_i - \overline{\pi}_i\right]\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*}\right]}$$

We know that the denominator must be positive since $1-t > 0$, $\underline{\pi}_i - \overline{\pi}_i < 0$, and from the definition of the relative audit rule, it must be the case that $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*} < 0$. The sign of $\frac{\partial \pi_i^g}{\partial x_i}$, must have the same sign as

$$\sum_{j\neq i}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*}\frac{\partial d_j^*}{\partial x_i}$$

If it is greater than zero, $\frac{\partial \pi_i^g}{\partial x_i}$ is positive; if it is less than zero, $\frac{\partial \pi_i^g}{\partial x_i}$ is negative; or equal to zero if zero. Again, from the definition of the relative audit rule we know that $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i^*} < 0$. Using the implicit function theorem, j's optimal declaration given a change in x_i can be written as:

$$\frac{\partial d_j^*}{\partial \pi_j^g}\frac{\partial \pi_j^g}{\partial x_i} = -\left[\frac{\partial \frac{\partial E\pi_j}{\partial d_j}}{\partial \pi_j^g} / \frac{\partial^2 E\pi_j}{\partial d_j^2}\right]\frac{\partial \pi_j^g}{\partial x_i}$$

or

$$\frac{\partial d_j^*}{\partial \pi_j^g}\frac{\partial \pi_j^g}{\partial x_i} = -\frac{1}{\frac{\partial^2 E\pi_j}{\partial d_j^2}}\left[C'' - (f+t)\frac{\partial \beta_j^*(\mathbf{d})}{\partial d_j}\right]\frac{\partial \pi_j^g}{\partial x_i}$$

We know that $C'' > 0$, $(f+t) > 0$, $\frac{\partial \beta_j^*(\mathbf{d})}{\partial d_j} < 0$, and $\frac{\partial^2 E\pi_j}{\partial d_j^2} < 0$. The sign of this term therefore depends on the sign of $\frac{\partial \pi_j^g}{\partial x_i}$. If $\frac{\partial \pi_j^g}{\partial x_i} > 0$ then $BR_i^R(\mathbf{x_{-i}}^*) < BR_i^F(\mathbf{x_{-i}}^*)$; if $\frac{\partial \pi_j^g}{\partial x_i} < 0$ then $BR_i^R(\mathbf{x_{-i}}^*) > BR_i^F(\mathbf{x_{-i}}^*)$; and if $\frac{\partial \pi_j^g}{\partial x_i} = 0$ then $BR_i^R(\mathbf{x_{-i}}^*) = BR_i^F(\mathbf{x_{-i}}^*)$.

$\square$

This result says that when a firm's action imposes a positive externality on competitors' profits, the firm will choose less of this action under a relative rule than under a fixed rule. The reason for this is because, choosing a lower action means that competitors make less profit, and they cannot declare such a large profit. If competitors cannot declare as much, the firm will face a lower audit probability. When the externality is in the opposite direction the reverse is true. When the firm's action has a negative effect on the profits on competing firms, the firm will choose a higher action under the relative rule than under a fixed rule. In this case, choosing a higher action means that competitors' profits are lower, and they cannot declare as much, which means that the firm will likely face a lower probability of receiving an audit. The intermediate case is simple: if a firm's actions have no effect on competing firms, there is no advantage in choosing anything other than the gross profit maximising action since a firm cannot affect the ability of the competing firms to declare profit and therefore cannot affect the audit probability in this way.

Under the fixed rule, there is no incentive to deviate from the gross profit maximizing action, since it would have no beneficial effect on the audit probability. Under the relative rule however, if firms know that their actions affect the profits of other firms, then they can use this to affect declarations and thus their audit probability. They can reduce the audit probability by appropriately adjusting their actions, "hurting" the other firms' profits, and making it harder for them to declare more in the tax stage. The relative rule therefore creates a new incentive that does not exist under the fixed rule: to try to reduce the profits of competing firms.

When firms are all fairly similar, we can say something as well about the aggregate behaviour of the firms in equilibrium. In equilibrium, the aggregate behaviour of firms under a relative rule depends on all of the externalities of firms' actions. Let $X^* = \sum_{i=1}^{n} x_i^*$ denote the sum of actions of all firms in equilibrium. We have the following result:

Proposition 1. *For the underlying market game* $(X_i, \pi_i^g, i\epsilon N)$*, compared with the equilibrium outcome under a fixed rule, any interior equilibrium outcome under a relative rule yields:*

A higher aggregate equilibrium action, $X_R^* > X_F^*$*, if* $\frac{\partial \pi_j^g}{\partial x_i} < 0\ \forall i, j\ i \neq j$;

An equal aggregate equilibrium action, $X_R^* = X_F^*$*, if* $\frac{\partial \pi_j^g}{\partial x_i} = 0 \forall i, j\ i \neq j$; or

A lower aggregate equilibrium action, $X_R^* < X_F^*$*, if* $\frac{\partial \pi_j^g}{\partial x_i} > 0\ \forall i, j\ i \neq j$.

Proof. Firstly, we would like to show that if $\frac{\partial \pi_j^g}{\partial x_i} < 0\ \forall i, j\ i \neq j$, then $X_R^* > X_F^*$. Consider any vector of actions, $\mathbf{x}$, less than the vector of actions chosen in equilibrium under a fixed rule, $\mathbf{x}^*$ (so $\mathbf{x} < \mathbf{x}*$) . The point $\mathbf{x}$ is clearly not a fixed point of BR^F by Lemma 1 and since BR is a contraction, $BR^F(\mathbf{x}) > \mathbf{x}$. This is because:

$$\left|BR^F(\mathbf{x}^*) - BR^F(\mathbf{x})\right| < c\,|\mathbf{x}^* - \mathbf{x}| \quad c\epsilon[0,1)$$

By Lemma 2, we have $BR^R(\mathbf{x}) > BR_i^F(\mathbf{x}) > \mathbf{x}$. So $\mathbf{x} < \mathbf{x}*$ is not a fixed point of $BR^R(\mathbf{x})$ either. Therefore any equilibrium under a relative rule must be $\geq \mathbf{x}^*$. Notice by Lemma 2 that $BR^R(\mathbf{x}^*) > BR^F(\mathbf{x}^*) = \mathbf{x}^*$. So $\mathbf{x} = \mathbf{x}^*$ is not a fixed point of BR^R and therefore not an equilibrium under the relative rule. We are left with only one case. Any equilibrium under a relative rule must be strictly greater than $\mathbf{x}^*$ and therefore $X_R^* > X_F^*$.

Secondly, we would like to show that if $\frac{\partial \pi_j^g}{\partial x_i} = 0\ \forall j\ i \neq j$, then $X_R^* = X_F^*$. This follows immediately from Lemma 2 since if $BR^R(\mathbf{x}) = BR^F(\mathbf{x})$, then if $\mathbf{x}^*$ is the unique fixed point of BR^F then BR^R must also have the same unique fixed point, $\mathbf{x}^*$. It follow that $X_R^* = X_F^*$.

Finally we would like to show that if $\frac{\partial \pi_j^g}{\partial x_i} > 0\ \forall i, j\ i \neq j$, then $X_R^* < X_F^*$. The same argument holds as for the first case. Consider any vector of actions, $\mathbf{x}$, greater than the vector of actions chosen in equilibrium under a fixed rule, $\mathbf{x}^*$ (so $\mathbf{x} > \mathbf{x}^*$) . The point $\mathbf{x}$ is clearly not a fixed point of BR^F by Lemma 1 and since BR is a contraction, $BR^F(\mathbf{x}) < \mathbf{x}$. This is because:

$$\left|BR^F(\mathbf{x}) - BR^F(\mathbf{x}*)\right| \leq c\left|\mathbf{x} - \mathbf{x}^*\right| \quad c\epsilon[0,1)$$

By Lemma 2, we have $BR^R(\mathbf{x}) < BR_i^F(\mathbf{x}) < \mathbf{x}$. So $\mathbf{x} > \mathbf{x}^*$ is not a fixed point of $BR^R(\mathbf{x})$. Therefore any equilibrium under a relative rule must be $\leq \mathbf{x}*$. Notice by Lemma 2 that $BR^R(\mathbf{x}^*) < BR^F(\mathbf{x}^*) = \mathbf{x}^*$. So $\mathbf{x} = \mathbf{x}^*$ is not a fixed point of BR^R and therefore not an equilibrium under the relative rule. Again we are left with only one case. Any equilibrium under a relative rule must be strictly less than $\mathbf{x}^*$ and therefore $X_R^* < X_F^*$. □

Proposition 1 says that, in the underlying market, if firms' actions have a positive externality on the profits of their competitors, then in equilibrium, under a relative rule, firms choose a lower action in the aggregate than they would under a fixed rule. Similarly, if firms' actions have a negative externality on the profits of their competitors, then under a relative rule, firms choose a greater action in the aggregate, than they would under a fixed rule. When firms actions have no effect on their competitors, the aggregate action is the same under the relative rule in equilibrium as it is under the fixed rule.

The intuition is an extension of the intuition for the first result. When firms' actions have a negative externality on each other, they individually choose more of that action under a relative rule since it reduces the profits of other firms and reduces their ability to declare a higher profit. Since all firms do the same thing, the same is true in the aggregate also. Every firm tries to reduce the profits of every other firm by choosing a higher action so the aggregate of all actions is also greater. Similarly, when firms' actions have a positive externality. Individual firms all choose less of their action and so the aggregate action is less also. If firms' actions have no externality on their competitors, then under a relative

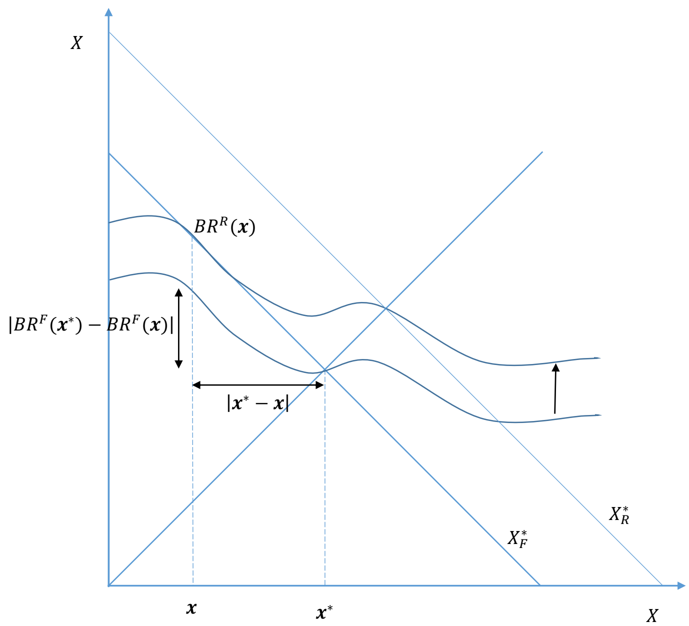

Figure 2.1: An illustration of Proposition 1 where $X_R^* > X_F^*$

rule, no individual firm chooses an action different to the gross profit maximising action and the aggregate action in equilibrium is the same too. Interestingly, we can see that the relative rule creates a prisoners dilemma type situation. Each firm tries to choose an action that will improve their relative standing in the tax stage by reducing the profits of other firms. However, since all firms do this, in equilibrium, no firm gains an advantage and they all lose profit in the process.

We turn, secondly, to the welfare implications of the change in firm behaviour. We are interested not just in how firms change their behaviour in response to a relative rule, but also in how this behaviour affects the surplus generated in the market. Since from Proposition 1, we see that there are three possible outcomes for how firms react to a change in audit policy (they might choose a higher action, a lower action, or the same action), and since this change, if there is one, could lead to an improvement or a reduction in efficiency, the effect of relative auditing might fall into one of the following five categories[11]: (1) an increase in action and an increase in surplus; (2) an increase in action and a reduction in surplus; (3) a decrease in action and an increase in surplus; (4) a decrease in action and a reduction in surplus; or (5) no change. The five categories are summarized in Table 2.1 below. In the next section we use some simple models to demonstrate the different possible outcomes when we apply the relative rule to different types of markets.

2.5 Market Structure

In this section we present some examples to show how the results of the previous section apply to more specific models of competition[12]. Our two results in Section 2.4 are fairly abstract; but our examples put a bit more structure on the problem. The purpose of these examples is to develop some intuition about the way the compliance problem works in specific settings and to develop some intuition that could be applied also to more complicated cases. In the previous section we identified five possible outcomes under a relative rule depending on what externalities firms' actions have on their competitors and how their actions affect market efficiency. We present an example of a market that fits each of

[11] In reality, these categories are not exhaustive for a few reasons. Welfare may not be monotonic in the action variable globally for the action of every firm. Also, the best response functions may not be monotonic or the relationships between firms may not be so symmetric that all of their behaviour changes in the same direction. Obviously, the model is an abstraction. Our aim is only to demonstrate the intuition and the effects of a relative audit rule for simple markets comprised of similar firms behaving as rational, optimizing decision makers. The intuition may be useful in more complicated cases, but this is beyond the scope of the analysis.

[12] We present only a few examples. We are not trying to demonstrate the breadth or range of possible models our results apply to. The reader is invited to think of his or her own examples. However, we have tried to use examples that demonstrate how the results work. The examples also reflect how the paper has developed—first, understanding the standard Cournot case, and then noticing the negative market effect in a differentiated Bertrand market, and then generalizing and applying the argument to other types of markets.

The Market effect (Examples)

	x_i is welfare improving	x_i is welfare reducing
$\frac{\partial \pi_j^g}{\partial x_i} < 0$	Market effect is positive Cournot with substitute goods Common pool resource	Market effect is negative Bertrand with complement goods Vertical competition Sabotage
$\frac{\partial \pi_j^g}{\partial x_i} = 0$	Market effect is neutral Perfect competition Monopoly	
$\frac{\partial \pi_j^g}{\partial x_i} > 0$	Market effect is negative Cournot with complement goods Public goods game	Market effect is positive Bertrand with substitute goods Advertising

Table 2.1: The market effect

these five categories.

From applying the results of section 4 we learn the following: (1) that there exists markets where the relative rule has a negative effect or no effect on market efficiency; (2) how in some settings, the relative rule creates an incentive for firms to behave in a way that is not competitive or socially beneficial; and (3) that even though the fraction of income evaded is lower under a relative rule (demonstrated by Bayer & Cowell 2009), this does not mean that the amount of tax revenue collected will be higher under a relative audit rule. The amount of tax revenue collected under a relative rule can actually be less than the amount collected under a fixed rule in some cases. We work through the examples in the same order that they are set out in the table (from top left to bottom right), starting with the case where firms' actions have a negative externality on each other's profits, and have a positive effect on welfare (Section 5.1). We finish with the case where firms' actions have a positive externality and a negative effect on welfare (Section 5.5).

In each of the following examples, there are just two firms, $i = 1, 2$. All of the functional forms are very simple (either linear or quadratic) and they are the same for each example—with the exception of gross profit. The gross profit function may take on various quadratic forms depending on the type of market we are discussing. Using simple functional forms allows us to impose enough structure on the problem to discuss different applications and explain how the model works without the model becoming unnecessarily complicated.

Since the functional forms for the evasion technology and the two audit rules are the same for all of the examples, we present these now and explain what each of the parameters represent. The evasion cost is first and then the two audit rules.

Firstly, the evasion cost is quadratic in the amount of income not declared with a scaling parameter, k. We have:

$$C(\pi_i^g - d_i) := \frac{(\pi_i^g - d_i)^2}{k}$$

For positive k, this function satisfies our restrictions since it is increasing in the undeclared amount, convex, non-negative and satisfies $C(\pi_i^g - d_i) = 0$ for $d_i = \pi_i^g$.

Secondly, the relative audit rule is linear in both the firm's own declaration and the declaration of the competing firm. The functional form for the relative rule is:

$$\beta_i^R(\mathbf{d}) := \begin{cases} 0 & \text{if } d_i > \frac{\alpha}{\gamma} + d_j \\ \alpha - \gamma(d_i - d_j) & \text{if } \frac{\alpha-1}{\gamma} + d_j < d_i < \frac{\alpha}{\gamma} + d_j \\ 1 & \text{if } d_i < \frac{\alpha-1}{\gamma} + d_j \end{cases}$$

The parameter α represents the average audit probability. In practice this parameter would be depend on the size of the budget set aside for auditing—since the more resources that are available for auditing, the greater the average number of audits that can be conducted. The parameter γ is a sensitivity parameter that represents how strongly the tax authority reacts to differences in declarations. If γ is large, then the difference in the audit probability faced by each firm will be very large. The high declaring firm will face a much lower audit probability and the low declaring firm will face a much higher probability. When γ is small on the other hand, the difference in the audit probabilities will be quite small, which means the relative rule will not be very sensitive to differences in declarations. This linear form satisfies the relative rule assumptions that we imposed above since it is decreasing in d_i, increasing in d_j and the sum of all changes in the audit probabilities for any given change in one of the declarations is zero ($\sum_{i=1}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j} = 0 \ \forall j$). In order that the relative rule and the fixed rule are comparable, we would like the average audit probability to be the same for both rules. Therefore we set the fixed rule probability to α, the same as the average audit probability of the relative rule. The amount of resources required to implement each rule then, will also be the same[13]. The fixed rule is therefore:

$$\beta_i^F := \alpha$$

We are now ready to present some examples.

2.5.1 Firms' actions impose negative externalities on competitors and their actions are welfare improving

We begin with the case where firms' actions have a negative externality on each other's profits and their actions have a positive effect on welfare. According to

[13] $n\alpha$ times the cost of carrying out one audit

our results in the previous section, models of competition that fit this description should have a positive market effect. This is the case, for example, when the model is applied to a "regular" Cournot market (see Bayer & Cowell (2010)).

Bayer and Cowell (2010) present a simulation of a "regular" linear Cournot model. We reproduce this simulation now, in this subsection, and then in Section 5.2, we argue that by changing the definition of the choice variables from quantity to price, we can reinterpret the model as a Bertrand market with complementary goods. There is a positive market effect when a relative audit rule is applied to a Cournot model with substitutes. However, in the Bertrand market with complement goods, we see there is a negative market effect. The Bertrand market with complements is our first example of a market where firms behave more inefficiently under a relative rule than they do under a fixed rule—contrary to the regular Cournot case.

Firstly, the Bayer and Cowell example. In this example, firms compete in a simple linear Cournot duopoly. There are two firms and they each produce substitute goods. They compete by choosing how much of these goods to produce. The gross profit function is given by:

$$\pi_i^g(q_i, q_j) = (a - bq_i - dq_j)q_i - cq_i$$

Each firm faces a linear inverse demand function and constant marginal costs. Their gross profit is a function of the quantities produced by each firm, q_i and q_j; and the parameters a, b and d, that describe the firm's inverse demand function (which is the term in brackets); and the parameter c, which is the constant marginal cost of production. This profit function satisfies the assumptions of a "regular" Cournot model (as well as our assumptions that the best responses are a contraction): inverse demand is decreasing; it is logconcave in $q_i + q_j$; and production costs, c, satisfy $c_i''(q_i) - p'(Q) > 0$. The two audit rules and the evasion costs in the Bayer-Cowell example are the same as the two that we will use for our examples; they are given by the equations above, at the beginning of this section. We will choose the parameter values so that they are also the same as the Bayer-Cowell example. The parameter values are: $a = 0.9$, $b = 1$, $c = 0$, $d = 1$, $\alpha = 0.25$, $\gamma = 2.5$, $k = 1$, $f = 0.5$ and $t = 0.3$. Let the vector of quantities chosen be denoted by $\mathbf{q} = (q_1, q_2)$ and given the functional forms and parameter values just mentioned, the firm's expected payoff function can be simplified to:

$$E\pi_i = \pi_i^g(\mathbf{q}) - td_i - \beta_i(\mathbf{d})(f+t)[\pi_i^g(\mathbf{q}) - d_i] - C(\pi_i^g(\mathbf{q}) - d_i)$$

where

$$\pi_i^g(\mathbf{q}) = (1 - q_i - q_j)q_i - 0.1q_i,$$

$$C(\pi_i^g - d_i) := (\pi_i^g - d_i)^2,$$

$$\beta_i^R(\mathbf{d}) := \begin{cases} 0 & \text{if } d_i > 1 + d_j \\ 0.25 - 2.5(d_i - d_j) & \text{if } d_j - 4 < d_i < 1 + d_j \\ 1 & \text{if } d_i < d_j - 4 \end{cases}$$

and

$$\beta_i^F := 0.2$$

We can now generate numerical solutions to the model. For the relative rule and fixed rule cases respectively, the results can be seen in Table 2.2. The quantity produced in equilibrium under a relative rule is greater than it is under a fixed rule. Firms are more competitive and they declare more of their profit. Also, price is lower and each firm earns less profit under the relative rule.

The intuition, used in Section 2.4 to explain the general results, can help explain this case. Notice that by differentiating profit with respect to q_j. The

Equilibrium Solutions

	Cournot with substitutes	
	Relative Rule	Fixed Rule
$d^*_{1,2}$	0.0633948	0.04
p^*	0.389655	0.4
$q^*_{1,2}$	0.305172	0.3
Q^*	0.610344	0.6
$\pi^{g*}_{1,2}$	0.0883948	0.09
$E\pi_{1,2}$	0.637513	0.6555
$t(d^*_1 + d^*_2)$	0.0380369	0.024
Expected Penalty	0.01	0.02
Total Revenue	0.0480369	0.044

Table 2.2: Numeric example: Cournot with substitutes

derivative is negative (for positive q_i): $\frac{\partial \pi^g_i}{\partial q_j} = -q_i \leq 0$. From Proposition 1, we know that when this is the case, firms will choose a higher value of their choice variable if they face a relative rule. The chain of reasoning is as follows. If a particular firm chooses a higher quantity then this has a negative externality on the profits of the firm's competitor. A higher quantity means there is more of the good to be sold; price falls; and the firm's competitor makes less profit. If the firm's competitor has a lower profit, then the representative firm's profit is relatively greater in comparison. A relatively higher profit means the representative firm has the potential to declare relatively more in the tax stage, and achieve a lower audit probability. A lower audit probability gives the firm a better chance to successfully evade tax, and thus, a higher expected payoff.

Also notice, since welfare is increasing in the choice variable (total surplus is higher when more is produced), and firms choose more of this choice variable, the market effect is positive. We have the Bayer-Cowell result.

2.5.2 Firms' actions impose negative externalities on competitors and their actions are welfare reducing

Suppose instead, we re-define the choice variables so they are prices instead of quantities. But everything else from the previous example remains the same. We have:

$$\pi_i^g(p_i, p_j) = (a - bp_i - dp_j)q_i - cp_i$$

Keeping the parameters the same also, so that again, $a = 0.9$, $b = 1$, $c = 0$, $d = 1$, means that:

$$\pi_i^g(\mathbf{q}) = (0.9 - p_i - p_j)p_i$$

The profit function can now be interpreted as the profit function of a firm competing in a differentiated Bertrand duopoly where firms produce complementary products. The term in brackets now represents demand, rather than inverse demand. Notice that it is a decreasing function of both the firm's own price and the price of the competitor. In an ordinary Bertrand market with substitute products, d would be negative since in this case demand increases when a competitor increases their price. In a Bertrand market with complementary products, d is positive. This is because profit decreases when a competing firm increases their price since demand for the good is lower if its complement is more expensive. We choose again, $d = 1$.

The reason why we are able to reuse this same model is because analytically, a linear Bertrand duopoly with complements is the dual of a linear Cournot duopoly with substitutes. There exists a duality between some simple types of markets. Sonnenschein (1968) first noticed this for a Cournot duopoly with perfect substitutes and a monopoly selling two perfectly complementary goods. Later Singh & Vives (1984) demonstrated the duality for differentiated Cournot and Bertrand duopolies. Bertrand duopoly with complements is analytically

Equilibrium Solutions		
	Bertrand with complements	
	Relative Rule	Fixed Rule
$d^*_{1,2}$	0.063	0.04
p^*	0.305	0.3
$q^*_{1,2}$	0.390	0.4
Q^*	0.779	0.8
$\pi^{g*}_{1,2}$	0.088	0.09
$E\pi_{1,2}$	0.637	0.6555
$t(d^*_1 + d^*_2)$	0.038	0.024
Expected Penalty	0.01	0.02
Total Revenue	0.048	0.044

Table 2.3: Numeric example: Bertrand with complements

isomorphic to the Cournot with substitutes example in the previous subsection and in Bayer and Cowell (2010)[14].

Running the simulation again, and swapping price and quantity, we get a slightly different result (again, see Table 2.3). Notice that everything is the same as in the previous example except the values for the price and quantities. These have swapped around. Rather than quantities being larger under the relative rule, the prices are now larger. This is not desirable from an efficiency perspective since it means that firms are actually *less* competitive and we have a negative market effect.

The intuition is again an extension of the intuition of our general results in Section 2.4. Notice that, since firms' actions (price) have a negative effect on their opponent's profit, ($\frac{\partial \pi^g_j}{\partial p_i} = -p_i < 0$), according to our general results, we should see that firms choose higher prices under a relative rule. This is exactly what they do. Except, unlike the Cournot case where higher quantities were desirable, higher prices are not desirable (in terms of surplus and therefore welfare). Firms are less competitive and the market effect is negative. The reason why we get this result is that the two firms choose higher prices because they know this has a negative effect on their opponent's gross profit. They

[14] which is what we are exploiting here. Similarly, the duality works for the reverse case too: a Cournot duopoly with complementary goods is the dual of a Bertrand duopoly with (differentiated) substitutable goods.

wish to increase their relative size of their profits so they can choose a relatively higher declaration in the second stage and achieve a lower audit probability.

One special case where the goods produced by firms are complements, is where firms interact vertically along a supply chain. We can re-interpret the results in this way too. Take the simplest case of two firms that compete vertically, where each has a monopoly over their own link in the supply chain. Let the two firms be: (1) a manufacturer (the upstream firm); and (2) a retailer (the downstream firm). Let, the manufacturer and the retailer compete vertically in price and assume that they decide on their prices simultaneously[15]. Simple profit functions for these two firms can be written as follows:

$$\pi_m := (w - c)(Q(w + m))$$

$$\pi_r := m(Q(w + m))$$

The gross profit function for the retailer, π_r, depends on the markup, m, they choose over the wholesale price, w; and on the final demand which is given by $Q(\cdot)$. Final demand for the retail good is simply a decreasing function of the retail price, $p = w + m$. The gross profit function for the manufacturer on the other hand, is the product of the wholesale price less the production costs, c, and derived demand. Since there is only one retailer and one manufacturer, the manufacturer's derived demand is equivalent to the retailer's demand, $Q(w+m)$.

Adding on the tax stage, we can see that this is another case where the relative rule will reduce the competitiveness of the firms. The reason for this is that an increase in either mark-up has a negative externality on the other firm's profit ($\frac{\partial \pi_m^g}{\partial m} = (w - c)\frac{\partial Q(w+m)}{\partial m} < 0$ and $\frac{\partial \pi_r^g}{\partial w} = (m)\frac{\partial Q(w+m)}{\partial w} < 0$). Given that each firm's action has a negative externality on the other, and applying

[15] Alternatively firms could compete by choosing prices sequentially, making it a three stage game. Even with very simple functional forms however, it becomes hard to find numerical solutions. See Ahmad, Anders and Marcoul (2013).

Equilibrium Solutions

	Vertical Competition		Sabotage	
	Relative Rule	Fixed Rule	Relative Rule	Fixed Rule
$d^*_{1,2}$	0.063	0.04	0.849	0.85
$a^*_{1,2}$	-	-	0.025	0.0
p^*	0.610	0.6	-	-
$q^*_{1,2}$	0.390	0.4	-	-
Q^*	0.390	0.4	-	-
$\pi^{g*}_{1,2}$	0.088	0.09	0.873	0.9
$E\pi_{1,2}$	0.637	0.655	0.614	0.633
$t(d^*_1 + d^*_2)$	0.038	0.024	0.509	0.51
Expected Penalty	0.01	0.02	0.1	0.2
Total Revenue	0.048	0.044	0.519	0.53

Table 2.4: Numeric example: Vertical Competition and Sabotage

Proposition 1, the mark-up that each firm chooses will be greater under the relative rule than it is under a fixed rule. This increase in price creates a further inefficiency in the market[16]. If $c = 0$, and we choose the same linear demand function as the Bertrand with complements case, we get the same results except the final retail price is twice as large since it is the sum of the two markups and the total quantity of the good sold is half as much since both firms are selling necessary elements of the same good. The market is extremely inefficient! The numerical solutions for equilibrium are given in Table 2.4.

We have one final example for this subsection: a model of “sabotage”. This example is designed to highlight the possibility that even when firms are otherwise completely independent, firms have an incentive to undertake actions that “hurt” the profits of other firms if they face a relative rule. Even at some cost to their own payoff. This is potentially a dangerous incentive since it not only encourages firms to engage in uncompetitive behaviour (as in the previous example where firms increase their prices) but also in spiteful behaviours like restricting access to inputs, imposing unnecessary costs or even illegal activities like vandalism.

[16]The market is already inefficient and there is also inefficiency from double marginalization—from having two monopolistic firms apply their own mark-ups in the same supply chain—which is why the final retail price is so high.

To illustrate we take the simplest possible model. Again, assume there are just two firms and suppose each firm earns a fixed profit each period, Y_i. However, each firm has the ability to "sabotage" the other by choosing some action, a_i, that is costly. Denote the cost of this action for the firm undertaking it as $C(a_i)$, where $C'(a_i) > 0$, so the cost is increasing with the extent of the action undertaken. Denote the cost of sustaining the sabotage imposed by another firm as $S(a_j)$. The firm's gross profit function is:

$$\pi_i^g(\mathbf{a}) = Y_i - C(a_i) - S(a_j)$$

For the sake of argument, let the cost for the firm imposing the action be quadratic and convex so that it is increasingly costly to undertake more sabotage, and suppose the cost of being subjected to sabotage is linear. We can write: $C(a_i) = a_i^2$ and $S(a_j) = a_j$. Given these restrictions:

$$\pi_i^g(\mathbf{a}) = Y_i - a_i^2 - a_j$$

With $Y_i = Y_j = 0.9$, and again, using the same evasion costs and audit rules, the equilibirum solutions under the fixed and relative rule policies are given in Table 2.4. Under a fixed rule, the amount of sabotage undertaken is 0. Under the relative rule however, the amount of sabotage undertaken is positive. As a result, both firms earn a lower profit and receive a lower payoff.

Notice that under a fixed rule, without the tax stage, there is no reason for a firm to engage in this action. Undertaking the action is costly, and there is no benefit to the firm. On the other hand, if there is a relative rule in place, firms gain an advantage if they can declare more profit relative to that which other firms can declare, since it means they can reduce their audit probability. To achieve this, firms must earn relatively more profit. They have an incentive to try to reduce the other firm's profit by sabotaging them, as long as the cost of doing so is less than the benefit derived from the lower audit probability. This is exactly what we see.

2.5.3 Firms' actions impose no externalities on competitors

Let us return to the Bertrand example again. The parameter d in the firm's demand function represents the type of interdependency they face. In the previous section d was positive, which meant firms were producing complementary goods. A change in the price of either good affected firms in the same way—if the price of the firm's own good increased or the price of its complement increased, demand went down. On the other hand, if d is negative, we have substitute goods. A price increase for one good increases the demand for the other good because consumers substitute away from the relatively more expensive good and choose the cheaper one.

In Figure 2.2, equilibrium prices are plotted against different values of d under a relative rule and under a fixed rule. Notice that the difference between price in equilibrium under the relative rule, p^{R*}, and price under a fixed rule, p^{F*} increases with d. When $d > 0$, we have complement goods, and $p^{R*} > p^{F*}$, as we saw in the previous subsection. When $d < 0$ however, we have substitute goods, and $p^{R*} < p^{F*}$. The reason we see this is that when $d < 0$, firms' actions have a positive externality on their competitors, ($\frac{\partial \pi_j^g}{\partial p_i} = -dp_i > 0$). Applying Proposition 1, we should see that firms choose a lower price in equilibrium, and the market is more efficient.

The case when $d = 0$ however, is special. When $d = 0$ firms' actions do not affect each other since then the gross profit functions depend only on the firm's own price. When each firm is producing a good that is completely unrelated to the other firm's good, and when each firm's demand is affected only by their

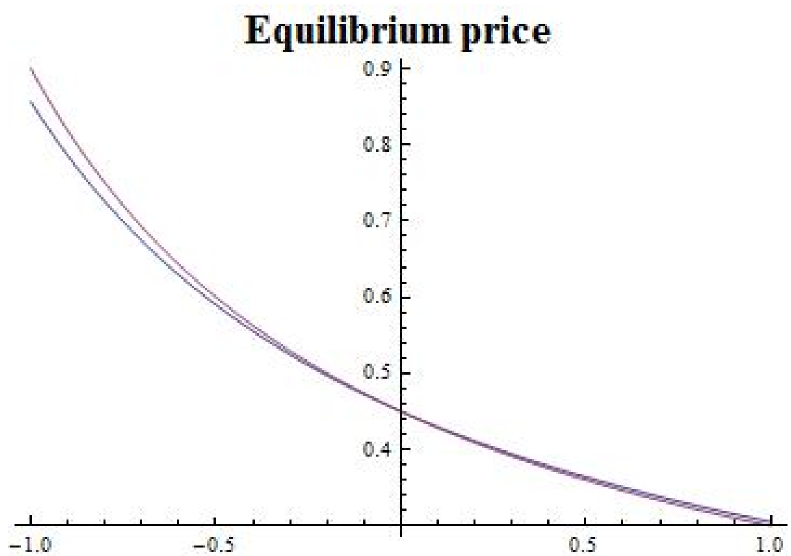

Figure 2.2: Equilibrium price for values of $d\epsilon[-1,1]$ under the relative rule (purple) and fixed rule (blue).

own price, each firm has a monopoly over their own good. The externality of a firm's price decision in this case does not exist; the derivative, $\frac{\partial \pi_j^g}{\partial p_i} = -dp_i$, is zero. We now have an example that fits the second case in both Lemma 2 and Proposition 1. Unsurprisingly, if we keep everything else the same as it is in the Bertrand-with-complements example, the numerical results when $d = 0$ are the same under a relative rule as they are under a fixed rule (see Table 2.5). This is the standard separation result between audit policy and market behaviour that is seem so often in the literature.

2.5.4 Firms' actions impose positive externalities on competitors and their actions are welfare improving

One example of a game where firm's actions have a positive externalitiy on each other's profit and where their actions are welfare improving, is a public goods

Equilibrium Solutions		
	2 Monopolies	
	Relative Rule	Fixed Rule
$d^*_{1,2}$	0.178	0.153
p^*	0.45	0.45
$q^*_{1,2}$	0.55	0.55
Q^*	1.1	1.1
$\pi^{g*}_{1,2}$	0.203	0.203
$E\pi_{1,2}$	0.144	0.144
$t(d^*_1 + d^*_2)$	0.107	0.092
Expected Penalty	0.01	0.02
Total Revenue	0.117	0.112

Table 2.5: Numeric example: 2 Monopolies. The equilibrium solutions are the same under either rule

game. The payoffs for a simple public goods game can be written as follows:

$$\pi_i^g(\mathbf{g}) = Y_i - g_i^2 + \theta\left(\sum_{i=1}^{2} g_i\right)$$

The firm's contribution to the public good is g_i. Their payoff, is a function of their income, Y_i, less the cost of contributing to the public good, g_i^2, plus the value of the public good, $\theta\left(\sum_{i=1}^{2} g_i\right)$. A profit function of this type could be interpreted as a market where firms must turn out their product from a common resource that is non-rivalrous and non-excludable. Examples of this might include situations where firms benefit from jointly funded research or from investing in some common infrastructure.

The externalities of each firm's action from contributing to the public good are positive since the derivative of gross profit with respect to the another firm's contribution is positive, $\frac{\partial \pi_j^g}{\partial g_i} = \theta > 0$. Also, the welfare effect of firms choosing a higher contribution is also positive if firms are (presumably) able to produce more or are more efficient as more is invested into the public good. According to Proposition 1 however, since $\frac{\partial \pi_j^g}{\partial g_i} > 0$, firms will choose a lower contribution under a relative audit rule than they would under a fixed rule. This lower contribution means, in turn, that welfare is lower. The public goods game is thus another case where the relative rule has a negative effect on market

Equilibrium Solutions

	Public Goods	
	Relative Rule	Fixed Rule
$d^*_{1,2}$	1.599	1.6
$g^*_{1,2}$	0.475	0.5
$\pi^{g*}_{1,2}$	1.624	1.65
$E\pi_{1,2}$	1.139	1.158
$t(d^*_1 + d^*_2)$	0.959	0.96
Expected Penalty	0.01	0.02
Total Revenue	0.969	0.98

Table 2.6: Numeric example: Public Goods

behaviour. Numeric results for the case where $Y_1 = Y_2 = 0.9$ and $\theta = 1$ are given in Table 2.6.

2.5.5 Firms' actions impose positive externalities on competitors and their actions are welfare reducing

In our final example we consider the effect of a relative rule on a market where firms compete in terms of advertising. It is not always clear how advertising by one firm affects the demand for a competing firm's product. There are two possible effects. On the one hand, advertising by one firm may draw attention to all of the goods in that particular market and affect the demand for all goods positively. On the other hand, perhaps only demand for the advertised good increases and consumers substitute away from the other goods—this would imply a negative effect on a competitor's demand.

The effects of advertising on welfare are also not always clear. On the one hand, advertising may inform consumers about what products are available and help them achieve more optimal spending decisions. This information might otherwise be quite costly to obtain. On the other hand, advertising might impose costs on the consumer, such as inconvenience or loss of time. We present the model first and then discuss the effect of different audit policies in each case.

To keep the model simple, we focus only on the firm's advertising decision. Consider again just two firms. Assume that the quantity and quality of a firm's advertising efforts can be summarized by a single parameter, a_i. For example, suppose that each firm's demand is a function of the prices, p_i, p_j, and advertising efforts, a_i, a_j of each firm. That is:

$$q_i = q_i(p_i, p_j, a_i, a_j)$$

Demand then is: decreasing in the firm's own price; increasing in the price of firm j; increasing in own advertising effort; and either increasing or decreasing in the advertising effort of the other firm. In order to concentrate on the effect of advertising, take prices to be constant, $\bar{p_i}, \bar{p_j}$, and consider that there is a production cost, $C(q_i)$ which is increasing in q_i, and an advertising cost, $A_i(a_i)$ that increases in a_i. We have

$$\pi_i(p_i, p_j, a_i, a_j) = \bar{p}q_i(\bar{p_i}, \bar{p_j}, a_i, a_j) - C_i(q_i(\bar{p_i}, \bar{p_j}, a_i, a_j)) - A_i(a_i)$$

Solving the firm's problem, maximizing this profit function with respect to a_i, we get

$$\bar{p}\frac{\partial q_i(\bar{p_i}, \bar{p_j}, a_i, a_j)}{\partial a_i} - \frac{\partial C_i(q_i(\bar{p_i}, \bar{p_j}, a_i, a_j))}{\partial a_i} - \frac{\partial A_i(a_i)}{\partial a_i} = 0$$

The necessary second order conditions for an interior solution is

$$\bar{p}\frac{\partial^2 q_i(\bar{p_i}, \bar{p_j}, a_i, a_j)}{\partial a_i^2} - \frac{\partial^2 C_i(q_i(\bar{p_i}, \bar{p_j}, a_i, a_j))}{\partial a_i^2} - \frac{\partial^2 A_i(a_i)}{\partial a_i^2} < 0$$

The firm's gross profit is globally concave, assuming diminishing gains in demand for increased advertising effort and convex advertising costs, when

$$\frac{\partial A_i^2(a_i)}{\partial a_i^2} - \bar{p}\frac{\partial^2 q_i(\bar{p_i}, \bar{p_j}, a_i, a_j)}{\partial a_i^2} > \frac{\partial C_i^2(q_i(\bar{p_i}, \bar{p_j}, a_i, a_j))}{\partial a_i^2}$$

With concave profit functions, an interior solution exists when

$$\lim_{a\to 0}\frac{\partial \pi_i}{\partial a_i} > 0 > \lim_{a\to\infty}\frac{\partial \pi_i}{\partial a_i}$$

The externality of one firm's advertising efforts on the competing firm is given by the sign of the cross derivative which is given by:

$$\frac{\partial \pi_j^g}{\partial a_i} = \bar{p}\frac{\partial q_j(\bar{p_j}, \bar{p_i}, a_j, a_i)}{\partial a_i} - \frac{\partial C_j(q_j(\bar{p_j}, \bar{p_i}, a_j, a_i))}{\partial a_i}$$

which could plausibly be positive or negative. If advertising by one firm serves to increase the demand for the product of a competing firm also, we have a case where firm's actions have a positive externality. If the firm increases their advertising efforts and demand for a competitor's product is greater, then $\frac{\partial \pi_j^g}{\partial a_i} > 0$. Assume that advertising by one firm has positive externalities on the demand for the other firm's product and advertising has a net negative impact on social welfare. These restrictions means that we have an example that fits the final category. For numerical solutions consider a demand function that is linear in advertising effort, $q_i(\bar{p_i}, \bar{p_j}, a_i, a_j) = (1 + a_i + a_j)$; has a constant marginal production cost, $C_i(q_i) = cq_i$; and a quadratic and convex advertising cost, $A_i(a_i) = -a_i^2$. The gross profit function becomes:

$$\pi_i^g(p_i, p_j, a_i, a_j) = \bar{p}(1 + a_i + a_j) - c(1 + a_i + a_j) - a_i^2$$

Given the following parameter values: $\bar{p} = 1$, $c = 0.1$, $\alpha = 0.25$, $\gamma = 2.5$, $k = 1$, $f = 0.5$ and $t = 0.3$, the results for the equilibrium advertising efforts under a relative audit rule and under a fixed audit rule are as given in Table 2.7.

Since, we have assumed that the effect of advertising effort on the demand for a competitor's product is positive, $\frac{\partial \pi_j^g}{\partial a_i} = (\bar{p} - c)a_j > 0$, application of Lemma 1 and Proposition 1 imply that firms choose less advertising effort under a relative

Equilibrium Solutions

	Advertising	
	Relative Rule	Fixed Rule
$d^*_{1,2}$	1.461	1.458
$a^*_{1,2}$	0.427	0.45
$\pi^{g*}_{1,2}$	1.486	1.508
$E\pi_{1,2}$	1.042	1.058
$t(d^*_1 + d^*_2)$	0.877	0.875
Expected Penalty	0.01	0.02
Total Revenue	0.887	0.895

Table 2.7: Numeric example: Advertising

rule compared with the fixed rule, and the aggregate equilibrium effort spent on advertising is also less under the relative rule. Indeed, the results in Table 5 support this: advertising effort is lower under a relative rule than under a fixed rule; and firm's profits are lower also. Since we have assumed that advertising is detrimental to social welfare, and since the relative rule reduces the amount of advertising, there is a positive market effect. Notice however, that this example could easily be rearranged so that the assumption about the effect of advertising on a competitor's product is to reduce demand. In this case, we would have yet another example where the relative rule leads to a *less* efficient market outcome.

A final point. Many of the examples show also that tax revenue is lower under the relative rule than it is under the fixed rule. This is puzzling since under a relative audit rule, both firm's declare a higher fraction of their profits in every example. Indeed, Bayer and Cowell (2009) show that there always exists a relative audit rule, such that firms declare a higher fraction of their profits than they do under a fixed rule. But a higher fraction of profit declared does not necessarily mean that more tax revenue is received by the tax authority. Since firms earn less profit under a relative rule, there is less profit owed to the tax authority under a relative rule. The net result of these two competing effects is ambiguous. It could be that the absolute amount of tax revenue is lower

Equilibrium Solutions

	Bertrand with Substitutes		Cournot with Complements	
	Relative Rule	Fixed Rule	Relative Rule	Fixed Rule
$d^*_{1,2}$	0.746	0.76	0.746	0.76
p^*	0.856	0.9	0.9	0.9
$q^*_{1,2}$	0.9	0.9	0.856	0.9
Q^*	1.8	1.8	1.713	1.8
$\pi^{g*}_{1,2}$	0.771	0.81	0.771	0.81
$E\pi_{1,2}$	0.541	0.570	0.541	0.570
$t(d^*_1 + d^*_2)$	0.447	0.456	0.447	0.456
Expected Penalty	0.01	0.02	0.01	0.02
Total Revenue	0.457	0.476	0.457	0.476

Table 2.8: Numeric examples: Bertrand with Substitutes and Cournot with Complements

even though firms are declaring a higher fraction. There are possible incentive problems then for the tax authority to use a relative rule if they are facing revenue targets. If their objectives however, are to meet targets based on the proportion of unpaid tax revenues, then this is not such a problem.

2.6 Conclusion

This chapter has looked at how audit policy might affect the way firms behave when they face different forms of market interaction. We considered an audit policy—the relative rule—that we have argued is a closer approximation to the way audit selection is usually carried out. The relative rule captures the fact that the tax authority can cheaply and easily condition audit selection on information provided by the firms in their tax reports, in order to increase the chances of recovering unpaid taxes.

Bayer and Cowell (2009) have shown that this audit rule increases the fraction of income reported by firms and it increases the competitiveness of firms in a Cournot oligopoly or in a Bertrand oligopoly with substitutable but differentiated goods. However, we see that if we apply a relative audit rule to other types of markets, it yields different outcomes in terms of how firms behave in the market and how efficient the market is. In particular, in markets where firms do not interact strategically, such as a monopoly or perfect competition, the relative rule has no effect on the way firms behave in the market. This is another form of the separation result that is prominent in the literature. More importantly, however, the relative rule may actually result in firms behaving *less* competitively and therefore in a less efficient market.

If the actions of each firm in the market have a negative externality on the profits of their competitors, then firms choose more of their action under a relative rule. If these actions are socially undesirable, then the relative rule makes the market less efficient. Examples of markets that fit this description include markets where firms produce complementary goods, or markets where firms interact vertically in a supply chain. Another example is where firms have the opportunity to sabotage a competing firm by choosing an action that carries only a small cost to their own profit but damages the profits of the competing firm.

On the other hand, if the actions of each firm in the market have a positive externality on the profits of their competitors then firms choose less of their action under a relative rule. If these actions are socially desirable, in this case, then the relative rule makes the market less efficient. One example of a market that fits this description is a public goods game where, for example, profits depend on some common resource that each firm can contribute to. This maybe the case in industries where firms are mostly independent except that they require access to some infrastructure (think of research or software perhaps), that can easily be shared. Another market that fits these criteria would be a model of advertising where a firm's advertising provides useful information to the consumer and is therefore socially desirable, but also draws attention to the products of other firms in the market and benefits these firms also. A relative rule in either of these two markets would, again, be detrimental to the efficiency of these markets.

Chapter 3

A Rank-Order Tax Audit Rule

3.1 Introduction

In this chapter we examine the effects of a rank order audit rule on two firms competing in a Cournot duopoly. The two firms are ranked according to their profit declarations. The firm that reports the highest declaration does not receive a tax audit, whereas the firm that reports the lowest declaration receives an audit with probability 1. This audit rule is just one of many ways a tax authority could select which firms to audit. However, it is an extreme case of the tax authority utilizing information from the firms' income declarations to help them better detect tax evasion. It is extreme in the sense that any variation in information declared by firms results in the full weight of audit resources falling on the firm with the weaker information. We aim to address the question: How do rational firms behave if a tax authority decides to target only the firm with the lowest declared profit?

There are already audit rules that take into account information declared by the tax-payer. We compare the effects of the rank-order rule with one other, the "relative audit rule". The relative rule is examined in work by Bayer and Cowell (2006, 2009, 2010) as well as Chapter 1 of this thesis. Firms in a Cournot market, as Bayer and Cowell demonstrate, behave more competitively when the relative rule is imposed on them. The rationale behind the relative rule is that a low declaration of profit by a particular firm might be viewed by the tax authority as a signal that this firm is trying to conceal profit. The tax authority then, would regard the lowest declaring firm as the most "at-risk" and allocate more audit resources towards them.

The question that arises is how much more should the tax authority allocate to these low declaring firms? This chapter aims to take the first step in addressing this question. The rank-order rule is a rule that can be thought of as the limit when the relative rule becomes extremely sensitive to differences in declarations. In some cases, it may be in the best interests of the tax authority to direct all audit resources to this most "at-risk" firm. A policy like this might be optimal if it is only feasible for the tax authority to audit one firm in the industry and firms realise that the tax authority cannot commit to choosing randomly so that randomizing is not a credible threat. This is our first reason for examining the rank-order rule: the tax authority may need to be *very* sensitive to the firms' reports and put all of the audit weight on the firm declaring the weakest information.

The second reason that we would like to study this audit policy is that Bayer and Cowell (2009) show that audit policy can have real effects—it can effect the way firms behave in the market—and these effects appear to be greater when audit selection is more sensitive to reported information. Since the rank order rule, is the *most* sensitive policy to differences in declarations we would like to know whether the effect on market behaviour that we see under a relative rule holds for the rank-order rule also. At the other extreme is an audit policy that is completely insensitive to tax declarations. The fixed rule has received the a lot of attention in the literature and audit selection is often modeled by simply assuming the probability of an audit is a fixed constant. Since in the fixed case where audit selection is not conditional on any of the information available to the tax authority, the audit probability is completely *in*sensitive to taxpayer reports. As a consequence, under a fixed audit regime, there is no effect on the way firms behave in the market. This a result that has been very well documented in the literature and has been shown to be robust to a variety of assumptions about market structure and firm objectives[1]. Firms make the same market decisions under a fixed rule as they would if there were no tax decision to make. Tax behaviour and market behaviour are completely separable because the firm's tax decision is a simple decision theory problem and does not depend on the behaviour of other firms. Each firm simply trades off the prospect of successfully evading tax with the prospect of being caught and paying the penalty.

Tax decisions and market decisions are no longer separable when audit selection depends on the information reported by firms however. Bayer & Cowell (2006, 2009, 2010) show that when audit selection depends on differences in the information reported by firms, the firms behave more competitively in the market. Furthermore, Bayer & Cowell (2009) simulate the effects of this types of audit policy for a 'regular' Cournot market. The equilibrium quantities are greater and the market is more efficient under a relative audit rule that takes into account reported information than under a fixed audit policy. However, the equilibrium quantities are higher still when the audit policy becomes more sensitive to differences in the firms' reports.

[1] Again, see Chapter 1 for a discussion of this.

In this chapter we examine the Bayer-Cowell model but substitute in the rank-order policy in place of the relative rule policy. We find, unlike in the relative auditing case, that the model has multiple equilibria—two pure strategy equilibria and one mixed. The set of rationalizable (pure) strategies contains two elements—one where a high quantity is chosen (which is higher than the Cournot quantity) and another where a low quantity is chosen (which is lower than the Cournot quantity). The qualitative results of the relative rule for a Cournot market thus fail to carry over to the limiting rank-order rule case. We cannot say that an individual firm's quantity choice will be more competitive under this rule. In addition to these two equilibria in pure strategies there is an equilibrium in mixed strategies where each firm randomizes between the high and low quantities. We argue the predictions of the model are therefore much less clear under the rank order rule given there are 3 equilibria. It is uncertain how firms might coordinate on any one in particular. Also the presence of a mixed strategy equilibrium means it is ambiguous which pure strategy will be realized and qualitatively, firms may move in either direction—they might supply more output or they may supply less output compared with the usual Cournot outcome depending on the equilibrium and the firm.

The structure of this chapter is as follows. We describe the model in Section 2 and formally define each of the audit rules. In Section 3 we solve for the equilibria of the model. Then in Section 4 we compare the outcomes of the rank order rule case with the outcomes of the model under a fixed rule and a relative rule. Section 5 concludes.

3.2 The Model

The tax compliance game in strategic form is $\Gamma := (N, (C_i)_{i\epsilon N}, (U_i)_{i\epsilon N})$ where N is the set of players, $(C_i)_{i\epsilon N}$ is a (finite) strategy space and $(U_i)_{i\epsilon N}$ is the set of payoffs. Each player is referred to as a "firm" and there are $n = 2$ of them $(i = 1, 2)$. Firms play a 2 stage game.

In the first stage or "market stage", they compete in a simple Cournot duopoly by choosing a quantity of goods to produce. In the second stage, the "tax stage", firms choose a declaration of profit to write on their tax report. The strategy set for a firm, C_i, is therefore the set of all pairs, $(q_i, d_i)\epsilon C_i$ where q_i is the quantity chosen in the first stage, and d_i is the declaration of profit made in the second stage.

We assume the second stage of the game is played in discrete space so that when we analyse the firms' tax behaviour, the best response functions are well defined. The rank-order rule creates problems in continuous space because firms will try to out-bid each other by infinitesimally small amounts. To avoid this, we assume that their choice variables are chosen from finite sets rather than intervals of the real line. The smallest increment will be ε in declaration space. Denote by G^d a finite grid such that $G^d := \{0, \varepsilon, 2\varepsilon, 3\varepsilon, \ldots\}$. We can use this to restrict the strategy space to subsets of a continuous interval. Let the values of d_i for $i = 1, 2$ be restricted to the set $G^d \bigcap [0, d^{max}]$. The values 0 and d^{max} are the minimum and maximum elements of this set respectively and assume they are located on the grid. The natural interpretation of ε is that it represents the smallest possible monetary unit available. The quantity space on the other hand is defined as the continuous interval $[0, q^{max}]$ of Euclidean space. In what follows we describe the payoff functions, then we explain each of the audit rules and finally, the timing of events in the game.

The set of payoffs, $(U_i)_{i\epsilon N}$, are defined as follows. After applying the expectations operator, the payoff for firm i is:

$$EU_i := \pi_i(\mathbf{q}) - td_i - \beta_i(\mathbf{d})(f + t)[\pi_i(\mathbf{q}) - d_i] \;\; for \quad i = 1, 2$$

In the first stage firms play a simple linear Cournot duopoly game. They produce a homogenous good and choose a quantity q_i of this good to sell in the market. The vector of quantities for each firm is given by $\mathbf{q} = (q_1, q_2)$. Each firm is identical and they each face the same gross profit function. Assume that firms face a simple linear inverse demand function with constant marginal costs. Without loss of generality gross profit can be written, $\pi_i(\mathbf{q}) := (\phi - q_i - q_j)q_i$.

In the tax stage, firms decide on an amount of gross profit to declare for tax purposes. Denote this profit declaration for the firm, by d_i. Again, these choice variables for each firm $i = 1, 2$, can be arranged in a vector, $\mathbf{d} = (d_1, d_2)$. In the firm's payoff function, in addition to the firm's gross profit, we also have to subtract the firm's tax bill, td_i and a penalty if they are caught evading tax, $\beta_i(\mathbf{d})(f + t)[\pi_i(\mathbf{q}) - d_i]$. The tax bill td_i is the product of the firm's own tax declaration d_i and the tax rate t. To keep things simple assume that t is constant. Finally, we have the expected penalty for evading tax is $\beta_i(\mathbf{d})(f + t)[\pi_i(\mathbf{q}) - d_i]$. The first term is $\beta_i(\cdot)$, which is the probability that the firm receives an audit. If a firm does receive an audit their evasion is necessarily found out—the audit is always assumed to be successful in recovering any unpaid tax. Two different functional forms for $\beta_i(\cdot)$ are called "audit rules". We formally define each of these audit rules below and describe each in more detail. Briefly though, the "relative audit rule" is where $\beta_i(\cdot)$ is a function of the firms' profit declarations and it depends on the relative strength of the firm's declaration, and importantly, the "rank-order audit rule" is where an audit is only awarded to the lowest declaring firm. The remaining terms in the expected penalty are described as follows. If an audit is realised, the firm must pay the rest of the tax they are liable for on the undeclared profit, $t[\pi_i(\mathbf{q}) - d_i]$. Also, they must pay a fine, $f[\pi_i(\mathbf{q}) - d_i]$. The fine rate f is constant so the fine paid is proportional to the undeclared amount. We impose $f > t$ to make the fine a good deterrent for evasion.

To summarise the second stage then, there are three possible outcomes for a firm. The first case ($d_i < \pi_i(\mathbf{q})$ and $\beta_i(\cdot) = 1$), is where the firm tries to evade tax but an audit is conducted and they are found out. This means the firm must pay tax on the full amount of profit, $td_i + t[\pi_i(\mathbf{q}) - d_i]$ and they receive a fine proportional to the amount they failed to declare $f[\pi_i(\mathbf{q}) - d_i]$. Their payoff is $\pi_i(\mathbf{q}) - td_i - (f + t)[\pi_i(\mathbf{q}) - d_i]$. The second case ($d_i < \pi_i(\mathbf{q})$,$\beta_i(\cdot) = 0$), is where the firm tries to evade and escapes audit. It means that the firm receives gross profit less the tax paid on only what was declared. The payoff is then $\pi_i(\mathbf{q}) - td_i$. Finally the third case ($d_i = \pi_i(\mathbf{q})$, and any $\beta_i(\cdot)$) is where the firm pays tax on all profit voluntarily. In this case they receive no penalty. Their payoff is $(1 - t)\pi_i(\mathbf{q})$.

The timing of the model is as follows[2]:

- Firms learn what the tax policy is including the tax rate, the penalty for under reporting and what the audit rule will be.

- Firms compete in the market stage. They choose a quantity.

- Gross profits are realized.

- Firms observe the gross profit for each firm in the industry and then choose how much to declare.

- The audit rule determines which firms receive an audit and the final payoffs are received.

We now describe the three different audit rules: the fixed audit rule, the relative audit rule, and the rank-order rule. We define each rule and then discuss each in turn.

[2]This is the same as in Chapter 1 still.

Relative audit rule

Definition 3. A *relative audit rule*, $\beta_i^R(\mathbf{d})$, is a function that satisfies the following conditions whenever $0 < \beta_i^R(\mathbf{d}) < 1$ for all $i = 1, 2$ and $j \neq i$:

- $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_i} < 0$
- $\frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j} > 0$
- $\sum_{i=1}^{n} \frac{\partial \beta_i^R(\mathbf{d})}{\partial d_j} = 0$

Rank-order rule

Definition 4. A *rank-order audit rule*, $\beta_i^{RO_k}(\mathbf{d})$, for two players $i = 1, 2$, and $j \neq i$ is such that:

$$\beta_i^{RO}(\mathbf{d}) = \begin{cases} 2\alpha & if \quad d_i < d_j \\ \alpha & if \quad d_i = d_j \\ 0 & if \quad d_i > d_j \end{cases}$$

The first two audit rules were discussed in Chapter 2 (see Section 2.2). The first rule—relative audit rule—says that the probability of receiving an audit is a decreasing function of the firm's own profit declaration and an increasing function of the rival firm's profit declaration. From simulations in Bayer and Cowell 2010, it appears that increasing the sensitivity of this rule to the differences in declarations between firms means that the effect of the relative rule on equilibrium quantities gets larger (at least for the linear Cournot duopoly case). More specifically, as the rule becomes more reactive to the difference in the declaration made by one firm and the declaration of the other, the equilibrium quantities become larger and the market becomes more competitive. This motivates our second rule.

The rank-order rule can be thought of as an extreme case of the relative rule where the sensitivity of the rule to differences in declarations is at its limit. The rank-order rule says simply, that the firm that declares the least profit receives an audit with certainty and sustains the weight of all audit resources available. Assume the average resources the tax authority has for each firm is α. Then with two firms in the market, the lowest declaring firm receives audit weight 2α. The firm that declares the most profit receives no audit with probability 1. In the case of a tie, the audit weight is split evenly between each firm—so each receives an audit with only half the total available resources. The rank-order rule is perfectly sensitive to any difference between the profit declaration of firm 1 and the profit declaration of firm 2 since, *any* difference will result in the low declaring firm bearing the full weight of audit. One might conjecture that the equilibrium outcome of the behaviour under the limit of the relative rule (in other words, behaviour under the rank-order rule), results in the most competitive outcome of the three policies. Before comparing the three rules, we first solve for the equilibria assuming a rank-order rule is in place.

3.3 Equilibria

The cases where a relative or fixed rule is in place have already been characterised in Bayer and Cowell (2006, 2009, 2010) as well as in Chapter 1 here. For a rank-order rule, we find the equilibium of the tax stage and then we work backwards to the Cournot stage.

3.3.1 Tax Stage.

Combining the payoff function with rank-order rule, the expected payoffs can be re-written:

$$EU_i = \begin{cases} \pi_i(\mathbf{q}) - td_i - 2\alpha(f+t)[\pi_i(\mathbf{q}) - d_i] & if \quad d_i < d_j \\ \pi_i(\mathbf{q}) - td_i - \alpha(f+t)[\pi_i(\mathbf{q}) - d_i] & if \quad d_i = d_j \\ \pi_i(\mathbf{q}) - td_i & if \quad d_i > d_j \end{cases}$$

The second stage works like Bertrand competition. Each firm tries to outbid their rival by making a slightly higher declaration. To see this consider firm i's behaviour given the following senarios:

1. $d_i < d_j - \varepsilon$
2. $d_i = d_j - \varepsilon$
3. $d_i = d_j$
4. $d_i > d_j$

1. If $d_i < d_j - \varepsilon$ and $d_i < \pi_i(\mathbf{q})$ then firm i can always do better by increasing d_i since the change in expected utility is positive. For an increase in d_i by an amount ε the corresponding change in EU_i is:

$$-t\varepsilon + 2\alpha(f+t)\varepsilon > 0$$

2. If $d_i = d_j - \varepsilon$ and $d_i < \pi_i(\mathbf{q})$ then increasing d_i by one unit ε, means each firm ties on the same amount. The change in EU_i is positive again:

$$-t\varepsilon + \alpha(f+t)\varepsilon + \alpha(f+t)[\pi_i(\mathbf{q}) - d_i] > 0$$

3. If $d_i = d_j$ and $d_i < \pi_i(\mathbf{q})$ the firm can increase their declaration by ε again

and gain still:

$$-t\varepsilon + \alpha(f+t)[\pi_i(\mathbf{q}) - d_i] > 0$$

4. Finally however, if $d_i > d_j$ and $d_i < \pi_i(\mathbf{q})$ then there is no benefit to increasing d_i since firm i already escapes the audit. Increasing d_i by ε results in a loss:

$$-t\varepsilon < 0$$

Notice we have assumed that $d_i < \pi_i(\mathbf{q})$ in all four cases. If $d_i = \pi_i(\mathbf{q})$ then assume firm i cannot increase their declaration any more since this would mean the firm is declaring more than they actually earn. Practically, it does not make a lot of sense for a firm to declare more than it earns because they would have to pay more tax than they have to. We will rule this possibility out by imposing that the firm's declaration is no larger than their profit[3]: $d_i \in [0, \pi_i(\mathbf{q})] \bigcap G^d$.

The information above for each of the four cases, as well as the fact that $d_i \leq \pi_i(\mathbf{q})$ means that we are ready to construct a best response function for each firm. Firms will increase their declaration until $d_i = \pi_i(\mathbf{q})$ or $d_i = d_j + \varepsilon$, depending on which is smaller. Thus, the best response function for firm i to the declaration of firm j, $j \neq i$ is:

$$d_i^{BR}(d_j) = \min\{\pi_i(\mathbf{q}), d_j + \varepsilon\}$$

The best response functions are depicted in Figure 1. Figure 1 illustrates how the firms try to outbid their rival in an effort to avoid an audit. In a dynamic context it is possible to imagine firms successively trying to coordinate on low

[3]We could have re-written the payoffs to make it explicit that the penalty only applies when firms are trying to evade: $EU_i = \begin{cases} \pi_i(\mathbf{q}) - td_i - \beta_i(\mathbf{d})(f+t)[\pi_i(\mathbf{q}) - d_i] & if \quad d_i < \pi_i(\mathbf{q}) \\ \pi_i(\mathbf{q}) - td_i & otherwise \end{cases}$. Then the payoff function makes it clear that increasing d_i above $\pi_i(\mathbf{q})$ results in a loss (of $-t\varepsilon$) and it is obvious that we can rule this possibility out. However, the exposition would be unnecessarily cluttered.

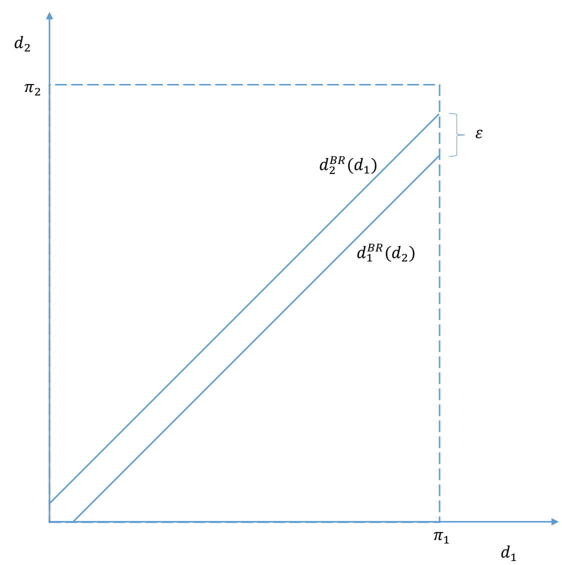

Figure 3.1: Best response declarations for Firm 1 and Firm 2

declarations but also bidding higher and higher amounts in order to secure the gains from evasion. There is only one equilibrium however because if one firm declares a low amount, the other always has an incentive to declare marginally higher. It means that for the firm with the lowest profit, all gains from evasion are lost in the bidding process since the high profit firm can always declare more. In equilibrium, the high profit firm declares an amount which is marginally more than the greatest amount the low profit firm can declare. The low profit firm declares everything. If both firms make the same profit however, then each has an incentive to outbid the other, which means both firms declare all of their profit. All of the gains from evasion are lost and each firm declares honestly.

The subgame perfect continuation declarations are:

$$d_i^* = \begin{cases} \pi_i(\mathbf{q}) & if \quad \pi_i(\mathbf{q}) \leq \pi_j(\mathbf{q}) \\ \pi_j(\mathbf{q}) + \varepsilon & \quad\;\; \pi_i(\mathbf{q}) > \pi_j(\mathbf{q}) \end{cases} \qquad i = 1, 2$$

Substituting d_i^* into EU_i the continuation payoffs become:

$$EU_i = \begin{cases} (1-t)\pi_i(\mathbf{q}) & if \quad \pi_i(\mathbf{q}) \leq \pi_j(\mathbf{q}) \\ \pi_i(\mathbf{q}) - t\pi_j(\mathbf{q}) - t\varepsilon & \quad\;\; \pi_i(\mathbf{q}) > \pi_j(\mathbf{q}) \end{cases} \qquad i = 1, 2$$

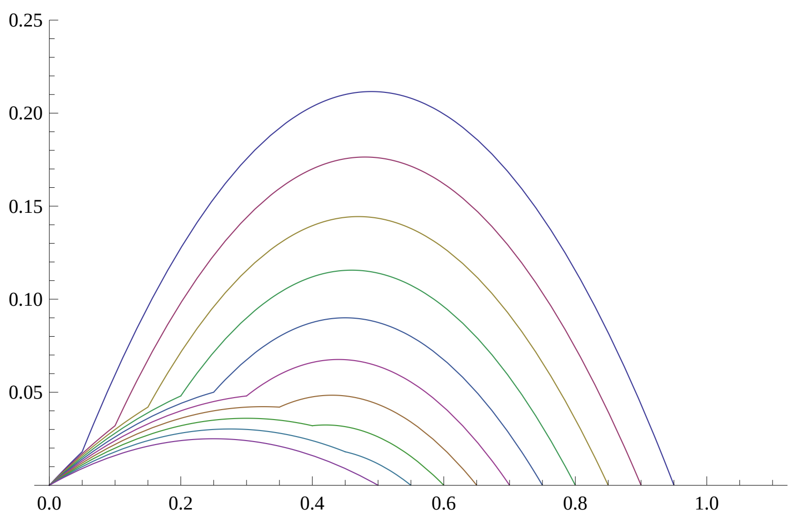

Figure 3.2: The expected utility functions for various quantities of the rival firm

3.3.2 Cournot Stage.

To understand the Cournot stage, we proceed as follows. Firstly, we find each firm's best response quantity to quantities chosen by the rival firm. This is done in a few steps. The continuation payoffs are rewritten so that the discontinuities are written in terms of the choice variables q_1 and q_2 rather than $\pi_1(\mathbf{q})$ and

$\pi_2(\mathbf{q})$. Then the firm's "optimal response" to each particular branch of the payoff function is identified. That is, we find the maximisers of each branch of the payoff function ignoring whether $\pi_i(\mathbf{q}) \leq \pi_j(\mathbf{q})$ and/or $\pi_i(\mathbf{q}) > \pi_j(\mathbf{q})$ actually holds for the firm.[4] With these steps completed, the firms' best response quantities can then be derived. First for "low" quantities of the rival firm, then for "high" quantities of the rival firm, and then finally for "intermediate" quantities. Once the best response functions are organised, we turn to solving for the equilibria. We first narrow the set of equililbrium candidates by solving for the set of rationalizable strategies. This is done through a process of Iterated Elimination of Never Best Responses (IENBR). From the set of rationalizable strategies we then solve the model for the pure strategy equilibria and then for the mixed strategy equilibrium.

3.3.2.1 The Best Response Correspondences

As mentioned, in order to characterize the best response functions we need to establish a some pieces of information. The "to do" list is as follows:

1. In quantity space, show where each branch of the firm's payoff function applies in order to identify where the discontinuities in the best response functions are.

2. State the "optimal response" of each firm to the rival's quantity for each branch of the payoff function.

3. Combine this information to find the best response quantities for each firm:

    - to "low" q_j

[4] "Optimal response" to distinguish these functions from the best response functions. They are maximisers of each branch, not neccessarily the best response. The best response will depend on which branch is applicable.

- to "high" q_j
- to "intermediate" q_j.

1. Notice that the points in quantity space where firms switch from being the high profit firm to the low profit firm are divided by the 45 degree line. For $q_i > q_j$ we have $\pi_i(\mathbf{q}) > \pi_j(\mathbf{q})$ for points where both firms are making nonnegative profit. Also, if $q_i = q_j$ we have $\pi_i(\mathbf{q}) = \pi_j(\mathbf{q})$. Figure 2 illustrates the regions of the quantity space where the following hold: $\pi_1(\mathbf{q}) > \pi_2(\mathbf{q})$, $\pi_1(\mathbf{q}) > \pi_2(\mathbf{q})$ and $\pi_1(\mathbf{q}), \pi_2(\mathbf{q}) < 0$. Since $q_i \lesseqgtr q_j$ implies $\pi_i(\mathbf{q}) \lesseqgtr \pi_j(\mathbf{q})$, the subgame perfect continuation payoffs can be rewritten (for profitable values of q_1 and q_2):

$$EU_i = \begin{cases} (1-t)\pi_i(\mathbf{q}) & if \quad q_i \leq q_j \\ \pi_i(\mathbf{q}) - t\pi_j(\mathbf{q}) - t\varepsilon & \quad q_i > q_j \end{cases} \qquad i = 1, 2$$

2. The second item we need is the optimal response to each branch of the payoffs. Define:

$$\begin{aligned} q^l(q_j) &= \arg\max_{q_i}(1-t)\pi_i(\mathbf{q}) \\ &= \frac{\phi - q_j}{2} \end{aligned}$$

$$\begin{aligned} q^h(q_j) &= \arg\max_{q_i}(\pi_i(\mathbf{q}) - t\pi_j(\mathbf{q}) - t\varepsilon) \\ &= \frac{\phi - (1-t)q_j}{2} \end{aligned}$$

These two functions are illustrated in Figure 3 for firm 2.

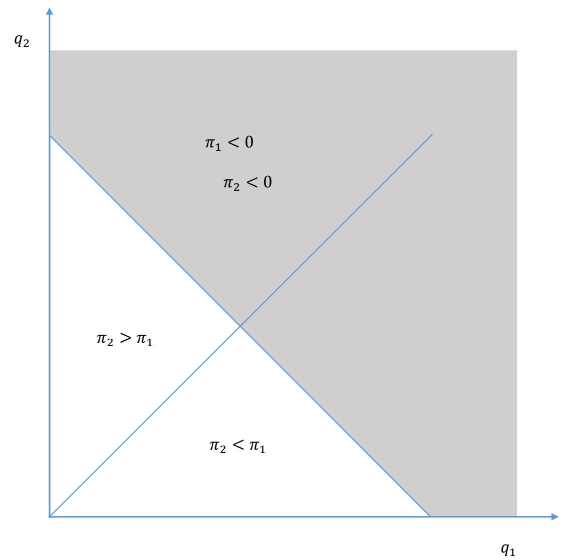

Figure 3.3: Profitable regions of quantity space

3. Now the best response quantities. We find the best response function in three steps. Define the Cournot quantity as the quantity where both firms choose $q^l(\cdot)$. Also let $\tilde{q}$ be defined as the quantity where both firms choose $q^h(\cdot)$. Then:

$$q^c = \frac{\phi}{3}$$

$$\tilde{q} = \frac{\phi}{3-t}$$

We find the best response for firm i firstly for low quantities, $q_j < q^c$, then for high quantities, $q_j \geq \tilde{q}$ and then for intermediate quantities, $q_j \epsilon [q^c, \tilde{q})$. Notice the slope of $q_i^h(q_j)$ is greater than the slope of $q_i^l(q_j)$. Also $q_i^h(q_j)$ lies above $q_i^l(q_j)$ for $q_j > 0$. We are now ready to show the how the firm responds to small

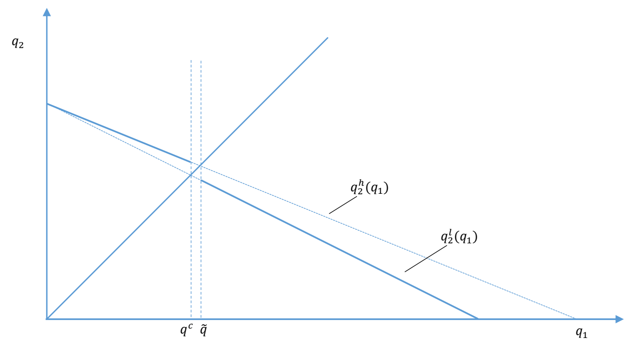

Figure 3.4: The 2 candidates for the best response of firm 2 to quantities of firm 1: $q_2^h(q_1)$ and $q_1^l(q_2)$

or large values of q_j.

Best response to small q_j

Lemma 3. *Let $\underline{q_j}$ be such that $\underline{q_j} < q^c$. Then for sufficiently small ε,*

$$\pi_i(q_i^h(\underline{q_j}), \underline{q_j}) - t\pi_j(\underline{q_j}, q_i^h(\underline{q_j})) - t\varepsilon) > (1-t)\pi_i(q_i^l(\underline{q_j}), \underline{q_j}).$$

Proof. Since $\underline{q_j} < q^c$ and since $q_i^l(\cdot)$ is downward sloping, $q_i^l(\underline{q_j}) > q^c > \underline{q_j}$. We know also that if firm i is on the low branch of the payoffs, the best they can do is choose $q_i^l(\underline{q_j})$. But since $q_i^l(\underline{q_j}) > \underline{q_j}$ and therefore $\pi_i(q_i^l(\underline{q_j}), \underline{q_j}) > \pi_j(\underline{q_j}, q_i^l(\underline{q_j}))$, we have

$$\pi_i(q_i^l(\underline{q_j}), \underline{q_j}) - t\pi_j(\underline{q_j}, q_i^l(\underline{q_j})) - t\varepsilon) > (1-t)\pi_i(q_i^l(\underline{q_j}), \underline{q_j})$$

By definition of $q_i^h(\cdot)$, it is at least as good for firm i to choose $q_i^h(\underline{q_j})$, so

$$\pi_i(q_i^h(\underline{q_j}), \underline{q_j}) - t\pi_j(\underline{q_j}, q_i^h(\underline{q_j})) - t\varepsilon) \geqslant \pi_i(q_i^l(\underline{q_j}), \underline{q_j}) - t\pi_j(\underline{q_j}, q_i^l(\underline{q_j})) - t\varepsilon) > (1-t)\pi_i(q_i^l(\underline{q_j}), \underline{q_j})$$

□

Best response to large q_j

Lemma 4. *Let $\bar{q}_j$ be such that $\bar{q}_j \geq \tilde{q}$. Then for sufficiently small ε,*

$$(1-t)\pi_i(q_i^l(\bar{q}_j), \bar{q}_j) > \pi_i(q_i^h(\bar{q}_j), \bar{q}_j) - t\pi_j(\bar{q}_j, q_i^h(\bar{q}_j)) - t\varepsilon.$$

Proof. The argument is similar to the argument for Lemma 2. Since $\bar{q}_j \geq \tilde{q}$ and since $q_i^h(\cdot)$ is downward sloping, $q_i^h(\bar{q}_j) \leq \tilde{q} \leq \bar{q}_j$. We know also that if firm i is on the high branch of the payoffs, the best they can do is choose $q_i^h(\bar{q}_j)$. But since $q_i^h(\bar{q}_j) \leq \bar{q}_j$ and therefore $\pi_i(q_i^h(\bar{q}_j), \bar{q}_j) \leq \pi_j(\bar{q}_j, q_i^h(\bar{q}_j))$, we have

$$(1-t)\pi_i(q_i^h(\bar{q}_j), \bar{q}_j) > \pi_i(q_i^h(\bar{q}_j), \bar{q}_j) - t\pi_j(\bar{q}_j, q_i^h(\bar{q}_j)) - t\varepsilon$$

By definition of $q_i^l(\cdot)$, it is at least as good for firm i to choose $q_i^l(\bar{q}_j)$, so

$$(1-t)\pi_i(q_i^l(\bar{q}_j), \bar{q}_j) \geq (1-t)\pi_i(q_i^h(\bar{q}_j), \bar{q}_j) > \pi_i(q_i^h(\bar{q}_j), \bar{q}_j) - t\pi_j(\bar{q}_j, q_i^h(\bar{q}_j)) - t\varepsilon$$

□

So far Lemma 1 and 2 imply the following:

$$q_i^{BR}(q_j)\big|_{q_j < q^c} = q_i^h(q_j)\big|_{q_j < q^c}$$

$$q_i^{BR}(q_j)\big|_{q_j \geq \tilde{q}} = q_i^l(q_j)\big|_{q_j \geq \tilde{q}}$$

Best response to intermediate q_j

For intermediate quantities $q_j \epsilon [q^c, \tilde{q})$, both optimal response curves satisfy the restrictions in the payoff function. That is, $q_i^h(q_j)\big|_{q_j \epsilon [q^c, \tilde{q})} > q_j$ and $q_i^l(q_j)\big|_{q_j \epsilon [q^c, \tilde{q})} \leq q_j$. So it is not immediately clear which branch is applicable. To see which branch yields a higher payoff consider the difference in firm i's payoff when they choose $q_i^h(q_j)$ compared with if they choose $q_i^l(q_j)$.

$$\begin{aligned} EU_i^h(q_i^h(q_j), q_j) - EU_i^l(q_i^l(q_j), q_j) &= [\pi_i(q_i^h(q_j), q_j) - t\pi_j(q_j, q_i^h(q_j)) - t\varepsilon] - [(1-t)\pi_i(q_i^l(q_j), q_j)] \\ &= t\left[\frac{(3+t)}{4}q_j^2 - \phi q_j + (\frac{\phi}{4} - \varepsilon)\right] \end{aligned}$$

Solving the zeros of this quadratic yields the q_j at which firm i would like to switch from being on the high payoff branch, playing the high response $q_i^h(q_j)$ to being on the low payoff branch and choosing $q_i^l(q_j)$. Denote this quantity[5] $\hat{q}$:

$$\hat{q} = \frac{2\phi - \sqrt{(1-t)\phi^2 + 4(3+t)\varepsilon}}{3+t}$$

The corresponding values for $q^h(\hat{q})$ and $q^l(\hat{q})$ are:

$$q^h(\hat{q}) = \frac{(1+3t)\phi + (1-t)\sqrt{(1-t)\phi^2 + 4(3+t)\varepsilon}}{2(3+t)}$$

$$q^l(\hat{q}) = \frac{(1+t)\phi + \sqrt{(1-t)\phi^2 + 4(3+t)\varepsilon}}{2(3+t)}$$

The best response quantities are:

$$q_i^{BR} = \begin{cases} q_i^h(q_j) & if \quad q_j \leq \hat{q} \\ q_i^l(q_j) & \quad\quad q_j \geq \hat{q} \end{cases} \qquad i = 1, 2$$

Where $q_i^h(q_j) = \frac{\phi-(1-t)q_j}{2}$, $q_i^l(q_j) = \frac{\phi - q_j}{2}$ and $\hat{q} = \frac{2\phi - \sqrt{(1-t)\phi^2+4(3+t)\varepsilon}}{3+t}$. Figure 4 illustrates the best response curves for $\phi = 1$ and $t = 0.3$.

[5] The other zero, $\frac{2\phi+\sqrt{(1-t)\phi^2+4(3+t)\varepsilon}}{3+t}$ is the point at which firm i would switch back to the high branch if making a negative profit meant that firms pay a negative tax (receive a transfer). This does not make much sense in our model—assume there is no benefit in making a loss.

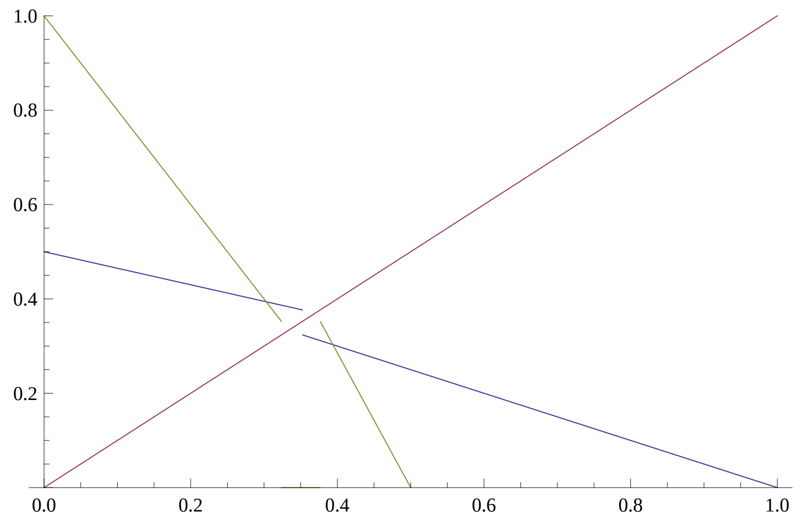

Figure 3.5: The best response curves for $\phi = 1$ and $t = 0.3$

3.3.2.2 Rationalizable Quantities

To narrow the set of candidates for equilibria, we use a process of iterated elimination of never best responses. An action is never a best response if it is not a best response to at least one belief of firm i about the rival firm. First of all, for firm i, there are two gaps in their quantity interval, $[0, q^{max}]$ which contain quantities that are never a best response to any q_j of the rival firm. These gaps are the intervals where q_i is greater than the monopoly quantity, $(\frac{\phi}{2}, q^{max}]$, and also the interval where the best response function jumps down, $(q^l(\hat{q}), q^h(\hat{q}))$. After one round of elimination the set of quantities is therefore narrowed to:

$$A_i^1 = \{q_i \epsilon [0, q^{max}] \setminus \{(\frac{\phi}{2}, q^{max}], (q^l(\hat{q}), q^h(\hat{q})\}\}$$

In the second round of elimination, we can rule out any quantities that are a best response to the equivalent quantities of firm j. After the second round of

	q_2^l	q_2^h
q_1^l	$EU_1^l(q_1^l,q_2^l)$,$EU_2^l(q_2^l,q_1^l)$	$EU_1^l(q_1^l,q_2^h)$,$EU_2^l(q_2^l,q_1^h)$
q_1^h	$EU_1^h(q_1^h,q_2^l)$,$EU_2^l(q_2^h,q_1^l)$	$EU_1^l(q_1^h,q_2^h)$,$EU_2^l(q_2^h,q_1^h)$

Table 3.1: Rationalizable quantities and their corresponding payoffs

elimination the quantity space is narrowed to:

$$A_i^2 = \{q_i \epsilon A_i^1 \setminus \{(q^l(\frac{\phi}{2}),0],(q^l(q^h(\hat{q}),q^h(q^l(\hat{q})))\}\}$$

Continuing to define the restricted sets recursively:

$$\begin{aligned}
A_i^3 &= \{q_i \epsilon A_i^2 \setminus \{(q^h(q^l(\frac{\phi}{2})),q^{max}],(q^l(q^h(q^l(\hat{q}))),q^h(q^l(q^h(\hat{q})))\}\} \\
A_i^4 &= \{q_i \epsilon A_i^3 \setminus \{(q^l(q^h(q^l(\frac{\phi}{2}))),0],(q^l(q^h(q^l(q^h(\hat{q}))),q^h(q^l(q^h(q^l(\hat{q})))))\}\} \\
\vdots & \quad \vdots
\end{aligned}$$

The limiting set of this series is the pair $\{q^l,q^h\}$ where $q^l = \frac{\phi}{3+t}$ and $q^h = \frac{\phi(1+t)}{3+t}$ (sum of a geometric series). The aggregate quantity under the rank order rule is thus $Q^{RO} = \frac{\phi(2+t)}{3+t} > Q^C = \frac{2\phi}{3}$, greater than the Cournot quantity (and therefore greater than under a fixed rule).

3.3.2.3 Equilibria

Since the set of rationalizable strategies contains only two elements, these can be presented in a 2 by 2 matrix. When firm 2 plays the low quantity, it is clearly better for firm 1 to play the high quantity since $q^l = \frac{\phi}{3+t} < q^c = \frac{\phi}{3}$. Similarly, if firm 2 plays the high quantity then it is better for firm 1 to play the low quantity—this time since we are on the low branch, $q^h = \frac{\phi(1+t)}{3+t} > \tilde{q} = \frac{\phi}{3-t}$. We have a game of Chicken. Each firm would like to play a different quantity to their rival. The set of equilibrium quantities in pure strategies are: $\{(q^l,q^h),(q^h,q^l)\}$.

Also we must consider equilibrium in mixed strategies. Let θ be the probability that firm i chooses q^l and let $(1-\theta)$ be the probability that firm i chooses q^h. Firm i chooses σ such that:

$$\sigma EU_j^l(q^l, q^l) + (1-\sigma)EU_j^l(q^l, q^h) = \sigma EU_j^h(q^h, q^l) + (1-\sigma)EU_j^l(q^h, q^h)$$

Solving for θ yields the weight attributed by each firm to the low quantity and $(1-\theta)$ gives the weight on the high quantity in equilibrium. We have[6] as $\varepsilon \to 0$:

$$\sigma = \frac{1-t}{2-t}$$

$$1-\sigma = \frac{1}{2-t}$$

3.4 Comparison with the Relative Rule

We now consider a relative rule that is comparable to the rank order rule in the previous section. Assume the expected payoff function and the gross profit function are the same as in the previous section. Also assume that the relative rule is linear. That is:

$$EU_i = \pi_i(\mathbf{q}) - td_i - \beta^R(d_i, d_j)(f+t)[\pi_i(\mathbf{q}) - d_i]$$

$$\pi_i(\mathbf{q}) := (\phi - q_i - q_j)q_i$$

$$\beta^R(d_i, d_j) = \begin{cases} 0 & \text{if } d_i > \frac{\alpha}{\gamma} + d_j \\ \alpha - \gamma(d_i - d_j) & \text{if } \frac{\alpha-1}{\gamma} + d_j < d_i < \frac{\alpha}{\gamma} + d_j \\ 1 & \text{if } d_i < \frac{\alpha-1}{\gamma} + d_j \end{cases}$$

[6] Or for positive ε, $\sigma = \frac{\phi^2(1-t)t}{\phi^2(2-t)t-(3-t)^2\varepsilon} 1-\sigma = \frac{\phi^2 t-(3-t)^2\varepsilon}{\phi^2(2-t)t-(3-t)^2\varepsilon}$

For the relative rule and the rank order rule to be comparable, the amount of audit resources required to implement each must be the same. In the rank order rule only one firm is audited and all of the audit resouces are allocated to that firm. The equivalent in the relative rule case is to allocate split the audit resources evenly amongst each firm and allow their declarations to determine which firm receives more or less than this amount. In other words, we need to keep the average audit resources, α, the same for each case. In the remainder of this section we solve the model with the linear relative rule in place, then, present some numerical results to compare the two.

Begin with the tax stage first. For an interior solution each firm chooses a declaration, $d_i^{BR}(d_j)$:

$$\begin{aligned} d_i^{BR}(d_j) &= \arg\max_{d_i}(\pi_i(\mathbf{q}) - td_i - \beta^R(d_i, d_j)(f+t)[\pi_i(\mathbf{q}) - d_i]) \\ &= \frac{\alpha(f+t) - t}{2\gamma(f+t)} + \frac{\pi_i(\mathbf{q})}{2} + \frac{d_j}{2} \end{aligned}$$

The equilibrium declaration for each firm is then:

$$d_i^* = \frac{\alpha f - (1-\alpha)t}{\gamma(f+t)} + \frac{2\pi_i(\mathbf{q}) + \pi_j(\mathbf{q})}{3}$$

Working backwards to the market stage the subgame perfect continuation is:

$$EU_i = \pi_i(\mathbf{q}) - td_i^* - \beta^R(d_i^*, d_j^*)(f+t)[\pi_i(\mathbf{q}) - d_i^*]$$

The best response quantities and equilibrium respectively are:

$$q_i^{BR}(q_j) = \arg\max_{q_i}(\pi_i(\mathbf{q}) - td_i^* - \beta^R(d_i^*, d_j^*)(f+t)[\pi_i(\mathbf{q}) - d_i^*])$$

$$q_i^* = \frac{\phi[(3-t) - 2\alpha(f+t)]}{(9-5t) - 4\alpha(f+t)}$$

Equilibrium Solutions

	Linear Cournot Duopoly	
	Relative Rule	Rank Order Rule
$d^*_{1,2}$	0.037	0.074
p^*	0.282	0.273
$q^*_{1,2}$	0.309	(0.355, 0.273)
Q^*	0.618	0.627
$\pi^{g*}_{1,2}$	0.0872	(0.097, 0.074)
$EU_{1,2}$	0.066	(0.074, 0.052)
$t(d^*_1 + d^*_2)$	0.022	0.045
Expected Penalty	0.02	0.00
Total Revenue	0.042	0.045

Table 3.2: Comparison of the rank order rule and the relative rule

Notice, again the aggregate quantity produced in the market is larger than Q^C assuming the parameter values give an interior solution and $t > \frac{\alpha f}{1-\alpha}$. The aggregate quantity is $Q^R = \frac{2\phi[(3-t)-2\alpha(f+t)]}{(9-5t)-4\alpha(f+t)}$. The rank order rule produces a greater quantity still for reasonable parameter values—if $f > \frac{t(1+3t-4\phi)}{4\phi}$. A numerical comparison of the two audit rules is provided in Table 3.2 for the parameter values $(\phi = 0.9, \alpha = 0.25, \gamma = 2.5, t = 0.3. f = 0.5)$. The values given for the rank-order rule are for the pure strategy equilibrium.

3.5 Conclusion

The rank order audit rule is an extreme that lies at the opposite end of the spectrum to a very common audit rule used in modelling tax evasion, the fixed rule. The rank order rule means that the tax authority is not ignoring the information that firms report but that they are starkly sensitive to such information. We have shown that in a symmetric linear duopoly with no evasion costs, the rank order rule is superior to the relative audit rule in that the expected aggregate quantity produced in any of the equilibia is greater than it is under a relative rule. Furthermore, firms declare more in equilibrium under this rule than they do under the relative rule. However, the predictions of the firm behaviour are far less clear cut. There are two choices of quantity that can be rationalized in the rank-order case and it is not clear how firms might coordinate their quantities to reach an equilibrium. One firm is required to produce a much lower quantity, make a lower profit and receive a lower payoff net of tax. Even less satisfying is the presence of a mixed strategy equilibrium. Mixed-strategy equilibria exhibit a regret property in that once one firm knows what quantity its rival is producing, it immediately regrets its choice of mixed strategy and would have preferred to respond optimally to the quantity observed (Vives 2001 pg. 45). The policy implication for a tax authority that is tempted to push relative auditing to the limit and focus all resources on the firm it judges to be the most likely evaders, is that the behavioural predictions in even the simplest case are not straight forward. Indeed, for our linear case with rational firms, the market is more competitive in the aggregate. But the for individual firms we cannot claim they each behave more competitively like they do under a relative rule. This phenomena fails to carry over to the rank-order case.

Chapter 4

Tax Audit Policy: Experimental Evidence

4.1 Introduction

In this chapter we present the results of a laboratory experiment on corporate tax evasion. Our aim is to observe the market behaviour and tax reporting behaviour of subjects under two different audit rules. The first is a "fixed" audit rule, which means that a tax audit is conducted on each firm with some fixed probability. The fixed audit rule has been prominent in the theoretical literature since the seminal work of Allingham & Sandmo (1972)[1]. The second audit rule is a "relative audit rule" which means that the probability of an audit is conditional on the relative income that each firm reports. A firm that reports lower (higher) income relative to other firms in the industry, faces a higher (lower) probability of being audited. In a Cournot market, under the relative audit rule a rational firm should not only evade less but they also behave more competitively in the

[1]See Chapter 1 2.1

market (Bayer and Cowell 2009). We test this prediction empirically using data from an experiment. Subjects compete in a two-stage game. In the first stage they compete in a Cournot duopoly. In the second stage, they choose a fraction of their profit to declare for tax purposes. Subjects can try to evade tax by declaring less profit than they actually make, but if audited, they must pay a penalty[2].

That we are aware of, there are no experimental studies that look at how tax policy might affect market behaviour. The bulk of the experimental economics literature on tax evasion focuses on decisions that individual taxpayers face, rather than those that firms face[3]. The reason for this is that in theoretical work, it is commonly assumed that firms face a fixed audit rule. Under the fixed rule, a very robust result is that tax decisions and market decisions are independent, which means there are no good theoretical reasons why the tax behaviour and market behaviour of firms should be related. For example, perhaps the most commonly used framework for thinking about corporate tax evasion is the Taxpayer-As-Gambler model. This treats the tax evasion decision as a gamble and assumes a fixed audit probability. Firms simply trade-off the prospect of evading tax with the prospect of being caught and the behaviour of other firms, including market behaviour, has no effect (see Cowell (2004) for a review). The lack of theory about how market behaviour might depend on tax policy, has made it difficult to examine empirically.

Bayer and Cowell (2009) show however, how under a relative rule firms' tax decisions are interdependent. The relative rule is an audit rule where the

[2]They lose all earnings from that period in our setup, which is quite extreme. We do this to try to maximise the potential variation in predicted market outcomes under the two audit rules.

[3]Exceptions to this include studies like Robben et al. (1990) and Webley (1987) where subjects face decisions framed in a business context. Subjects are required to make many decisions for their business of which the reporting decision is just one. The focus of these experiments is on the tax behaviour of the subjects; the other decisions subjects must make are included simply to add realism and make the purpose of the experiment less obvious to the subjects.

probability that a firm is audited depends on the firm's declared income relative to the income declared by other firms in the industry. The probability of being audited is decreasing in the firm's own declared income and increasing in the amount of income declared by the other firms. In this framework, firms have two incentives: firstly, to report higher income; and secondly, to reduce the amount of income their competitors can declare. The first incentive means that firms evade a smaller fraction of their profit than they would under a fixed rule. The second incentive means that rational firms behave more competitively in the market to try to reduce the relative profit of their rivals. Firms' decisions under a relative rule, thus depend on both the tax decisions and market decisions of their competitors—they are strategically interdependent.

Although there have been no experiments that examine the effect of audit rules on market behaviour, there have been many laboratory experiments that examine the effects of different audit rules on compliance behaviour[4]. In most of these experiments the audit rules do not import strategic tension between taxpayers. Exceptions include experiments by Alm & McKee (2004) and Tan & Yim (2014).

Alm & McKee (2004) examine a relative audit rule designed to mimic the Discriminant Index Function (DIF) used by the Internal Revenue Service. A taxpayer declaring the lowest declaration in the group (of 5 subjects) is audited with certainty—except if they all declare the same amount—then no one is audited. There are some treatments that add a random component as well—if all taxpayers manage to coordinate on declaring the same amount, one is audited at random. This makes it harder for subjects to coordinate on the same declaration. Also there are "cheap talk" treatments where subjects were allowed to communicate briefly prior to the experiment. They find that cheap talk al-

[4]See for example: Mittone (2006); Cadsby et al. (2006); Clark et al. (2004); Alm et al. (1993)Alm et al. (1993); Alm, Jackson & McKee (1992); Alm, McClelland & Schulze (1992); and Becker et al. (1987); as well as the other references mentioned here.

lows the subjects to coordinate their declarations and evade more, but when the random component is added, the effect is undone. Unfortunately there is no treatment with a fixed audit rule to compare with the DIF rule, but they observe an average compliance rate which is quite high and similar to what we observe in the data from our relative treatments[5].

Tan & Yim (2014) also look at an experiment where subjects interact strategically. They look at a "bounded audit rule". Taxpayers receive either a high income or a low income according to some fixed probability. Taxpayers can then declare either high or low income irrespective of what they actually receive, but the bounded rule means that there is an upper bound on the number of audits carried out on those that declare low. Since there is no need to audit those that declare high income, the audit probability depends on how many taxpayers declare low income. They find that the bounded rule is effective in reducing tax evasion and that tax evasion can be reduced further by increasing the level of strategic uncertainty[6].

We look at a slightly different audit rule in this paper. We examine a "relative audit rule" where probability that a particular firm receives an audit is a linear function, increasing in the firm's own declaration and decreasing in the declaration of its competitor. We use a linear rule because it is easy to understand and it allows us to derive predictions for the rational decision maker while still capturing the strategic nature of the tax compliance game. As in Alm & McKee (2004) and Tan & Yim (2014) just described, having an audit rule that is conditional on the declarations made by each firm should encourage subjects to report more truthfully. This is our first hypothesis.

---

[5]The average compliance rate in the DIF + random treatment is 71%, which is similar to what we find in our 2 relative treatments (74% and 66%).

[6]The policy recommendation is that increasing the level of strategic uncertainty could be effective in deterring tax evasion. In the Tan and Yin framework this means having more people receive higher income so more people can potentially evade. (A policy I think most would support!)

Importantly however, we would also like to see whether audit policy has an effect on behaviour in the market stage. The empirical literature is very quiet on this but our theory provides a good explanation of how audit policy might affect market behaviour and gives us a framework to understand the data. The nature of tax evasion and the problems with identification and measurement means it is hard to study with field data. Although the behaviour of subjects in a laboratory cannot be assumed to be perfectly representative of behaviour in more complex situations, the laboratory setting allows us to identify differences in behaviour due specifically to policy. Our theoretical framework is parameterised so that firms should declare the same amount in equilibrium whether they face a relative rule or a fixed rule. This is our second hypothesis.

We also examine the separation of decision-making roles between the tax manager (we will call them CFOs) and the market operations manager (CEOs) within a firm. We contrast this with the case where a firm has one agent that is responsible for both the tax decisions and market decisions. We will call these firms owner-manager firms. When a fixed audit rule is in place, the behavioural analysis implied by the theory of the firm's decision making is very simple. The department in charge of market decisions can operate without consulting those in charge of taxation because of the separation result—they can work independently and need not communicate. The literature on tax evasion up until recently would suggest that this is a reasonable way for firms to operate. Under the fixed rule the firm evades up until the point where the expected marginal cost of evasion equals the expected marginal tax saving. There is no reason why the outcome of the market should affect this. Under a relative rule too, if the incentives of both decision makers are perfectly aligned, as they are in our experiment, the theoretical predictions for the relative rule under the CEO/CFO arrangement are the same as they are under the owner-manager arrangement. Behaviourally however, it is not so clear whether subjects would behave the same under the CEO/CFO separation since subjects are likely be more focussed on their own decisions and CEOs may fail to see how their actions affect the tax decision of the CFO. We examine this separation of duties in our third treatment. Our third hypothesis is that the separation of decision-making does not have an effect on the market behaviour of firms under the relative rule.

We have three treatments therefore. The first is where firms have only one decision maker (an owner-manager) and there is a fixed rule in place. The second is where firms again, are owner-manager firms, but there is a relative rule in place, and the third treatment is where firms comprise of two subjects (a CFO and CEO) and a relative rule is in place.

We find from our data that the relative rule has a significant effect on tax behaviour but not on the market behaviour of firms. There is a significantly lower fraction of subjects evading tax under the relative rule compared with under the fixed rule and the average amount evaded is also significantly lower. In the market stage however, the relative rule causes no significant difference in the subjects' behaviour. For owner-manager firms, interestingly, we see that they are *less* competitive under the relative rule than they are under a fixed rule, although the effect is not significant. Under the relative rule, when firms are comprised of two subjects (CEO/CFO), we see that they are significantly more competitive than they are in the treatment with owner-manager firms.

The remaining sections of this paper are organised as follows. In Section 2 we present the theoretical framework for both the fixed case and the relative case. We then describe the three treatments in more detail, we present the equilibrium predictions of the model and describe the experimental procedure. In Section 3 we present the results of the experiment including a descriptive overview and we estimate the effect of the market outcomes on the tax declarations to determine how the market behaviour affects tax behaviour. Section 4 concludes.

4.2 Experimental Design

4.2.1 Theoretical Framework

The game is played in two stages and there are two players: firm 1 and firm 2. In the first stage, the two firms compete in a Cournot duopoly. They select a quantity of goods to supply to the market and once each firm has chosen their quantity, the market clears. Each firm receives their profit as a function of the quantities chosen. In the second stage, they choose a fraction of their profit to declare for tax purposes. Each may choose to declare an amount less than their true gross profit so as to reduce the amount of tax they have to pay, but they risk having to pay a penalty if found out. Nature chooses whether each firm is found out according to a probability that is either constant, (in the fixed treatment) or depends on the amounts that each firm declares (in the two relative treatments). The penalty a firm faces if it is caught evading tax is that it loses all profit from that period. The timing of the game is as follows:

1. Firms learn what the tax policy is including the tax rate, the penalty for under reporting (they lose everything) and what the audit rule will be.
2. Firms compete in the market stage in Cournot competition. They choose quantities.
3. Gross profits are realised.
4. Firms observe their own gross profit and the gross profit of their rival.
5. They choose how much gross profit to declare.

6. Nature decides whether each firm is found out (if they choose to under report).

7. Firms are informed whether or not they are found out and they receive their ex-post income (which is gross profit less taxes if not found out, or zero if they are).

In the first stage, we use a standard linear Cournot game without cost and with a fixed income component, Z. Let q_i denote firm i's quantity choice and q_{-i} denote the quantity choice of firm i's competitor. A, Z and γ are parameters that we fix below. Firm i's gross profit then is given by:

$$\Pi_i(q_i, q_{-i}) := [(A - q_i - q_{-i})q_i + Z]\gamma$$

The second stage, is different depending on which audit rule is in place. We discuss the fixed rule case first and then the relative rule case.

4.2.1.1 The Fixed Case

Suppose we denote the fixed probability of detection as α. Furthermore, suppose we have a tax function, t, also a constant, so that the tax paid by a firm is directly proportional to reported profit, d_i. We can write the expected net payoff as:

$$EU_i := \begin{cases} (1-\alpha)(\Pi_i(q_i, q_{-i}) - td_i) & if \quad d_i < \Pi_i(q_i, q_{-i}) \\ \Pi_i(q_i, q_{-i})(1-t) & if \quad d_i = \Pi_i(q_i, q_{-i}) \end{cases}$$

A risk-neutral firm either reports honestly or declares zero profit. A firm will evade everything if the expected payoff from doing so, is weakly greater than the payoff from declaring honestly. That is,

$$(1-\alpha)\Pi_i(q_i, q_{-i}) \geqslant \Pi_i(q_i, q_{-i})(1-t)$$

$$\Longrightarrow \alpha \leq t$$

In the first stage, under a fixed audit regime, the firm maximises $\Pi_i(q_i, q_{-i})$. We can see that the tax decision does not affect the firm behaviour in the first stage because maximizing the firm's EU with respect to q_i yields the standard Cournot solution

$$q_i^* = A/3$$

Gross profit from the first stage then is given by

$$\Pi_i(q_i^*, q_{-i}^*) = \left[\left(\frac{A}{3}\right)^2 + Z\right]\gamma$$

4.2.1.2 The Relative Case

Now suppose instead that the probability of detection is a function of d_i and d_{-i}. We impose the following linear rule

$$\alpha_i(d_i, d_{-i}) := \alpha - \beta(d_i - d_{-i})$$

The average audit probability is still α but the probability each individual firm faces depends on how much each declares relative to their competitor. A firm that declares more than its competitor will have a lower chance of getting audited, while a firm that declares less will have a greater chance of getting audited.

Replacing the fixed audit rule with the linear relative rule in the firm's payoff function, the expected payoff of firm i is now:

$$EU_i := \begin{cases} (1-\alpha+\beta(d_i-d_{-i}))(\Pi_i(q_i,q_{-i})-td_i) & if \quad d_i < \Pi_i(q_i,q_{-i}) \\ \Pi_i(q_i,q_{-i})(1-t) & if \quad d_i = \Pi_i(q_i,q_{-i}) \end{cases}$$

Maximising the firm's payoffs with respect to d_i, the first order condition for the firm's second stage best response is:

$$d_i^*(d_{-i}) = \begin{cases} 0 & if & d_{-i} \leq \frac{1-\alpha}{\beta} - \frac{\Pi_i}{t} \\ \frac{\beta\Pi_i - t(1-\alpha)}{2\beta t} + \frac{d_{-i}}{2} & if & \frac{1-\alpha}{\beta} - \frac{\Pi_i}{t} < d_{-i} < \frac{1-\alpha}{\beta} + \Pi_i(2-\frac{1}{t}) \\ \Pi_i & if & \frac{1-\alpha}{\beta} + \Pi_i(2-\frac{1}{t}) \leq d_{-i} \end{cases}$$

A necessary condition for an interior solution is $\alpha_i \epsilon [0,t]$. We take monopoly profit, Π^m, and zero profit as the two bounds for gross profit. We calibrate the model so that:

$$\alpha_i(\Pi^m, 0) = \alpha - \beta(A^2/4 + Z)\gamma = 0$$

$$\alpha_i(0, \Pi^m) = \alpha + \beta(A^2/4 + Z)\gamma = t$$

$$\Longrightarrow \alpha = t/2$$

With the right choice of parameters, the best response reduces to

$$d_i^*(d_{-i}) = \frac{\beta\Pi_i - t(1-\alpha)}{2\beta t} + \frac{d_{-i}}{2}$$

Solving simultaneously, the equilibrium declarations then are

$$d_i^*(d_{-i}) = \frac{2\Pi_i + \Pi_{-i}}{3t} - \frac{1-\alpha}{\beta}$$

for $i = 1, 2$. We need also, that this is part of an interior solution in the market stage. Solving for the quantity stage, a firm maximises the expected net profit taking into account the optimal declarations:

$$EU(d^*) := \frac{[\beta(\Pi_i - \Pi_{-i}) + 3(1-\alpha)t]^2}{9t\beta}$$

$$= \frac{[\beta(q_i - q_{-i})(A - q_i - q_j)\gamma + 3(1-\alpha)t]^2}{9t\beta}$$

It can be shown that a weakly dominant strategy solving for the symmetric equilibrium gives:

$$q_i^* = \frac{A}{2}$$

If this is an interior solution we have the competitive outcome. Alternatively, if we have corner solutions in the second stage and an interior solution in the first stage the equilibrium quantities are again Cournot.

4.2.2 Treatments and Equilibrium Predictions

We conduct three treatments. In the fixed individual treatment ($FIXIND$) subjects are audited according to a fixed probability and each firm is made up of just one subject, owner-managers. In the relative individual treatment ($RELIND$) subjects are audited according to the relative audit rule and again, each firm is made up of just one subject. Finally, in the relative group treatment ($RELGROUP$) each firm faces a relative audit rule but is made up of two subjects; one has the job of being CEO and chooses the quantity produced each round, while the other is CFO and makes the declaration decision each round. This allows us to examine whether the relative rule has any effect on market behaviour given the separation in decision-making.

Setting the tax rate to 30% means that from our $\alpha = t/2$ condition, $\alpha = 0.15$.

In order to get the gross profit values to be round numbers we scale quantities in the payoff function so that $q_i = \frac{q_i^{'}+1}{2}$, where $q_i^{'}$ is the quantity the subjects actually choose. The gross profit matrix can be seen in the instructions for the experiment. We set the remaining parameters as follows. Let:

$$t = 0.3$$

$$\alpha = 0.15$$

$$\beta = 0.001$$

$$Z = 2.55$$

$$A = 6$$

$$\gamma = 100$$

With these parameters, we get equilibrium predictions for a rational, risk neutral firm. The equilibrium outcome of the fixed treatment, is that both firms choose $d = 0$ and $q = 3$. Note $q = 3$ is the Cournot solution; if the game were played without the tax stage, the Cournot-Nash equilibrium is $q = 3$. For the relative treatments, $RELIND$ and $RELGROUP$, the Nash equilibrium is again both firms choose $d = 0$ and $q = 3$. When price equals marginal cost (zero), the competitive output is $q = 5$ given our vector of parameters and the transformation mentioned above.

4.2.3 Experimental Procedure

The experiment was conducted at Adlab, at the University of Adelaide, in November, 2013. There were 8 sessions held with an average of 23 subjects each. Two sessions were held for the $FIXIND$ treatment, two were held for the $RELIND$ treatment, and four were held for the $RELGROUP$ treatment.

	FIXIND	*RELIND*	*RELGROUP*	All treatments
Number of subjects	46	46	92	184
Gender (% female)	41	39	50	45
Age (%)				
16-25	83	61	65	68
26-30	13	24	25	22
31-40	0	11	9	7
41-50	2	2	0	1
51-60	2	2	0	1
>60	0	0	1	1
Maths (% yes)	76	73	78	77
Study level (%):				
Vocational	0	0	1	1
Undergraduate	72	54	66	65
Post grad. Coursework	11	26	17	18
Post grad. Research	17	20	15	17

Table 4.1: Subject characteristics

All treatments were carried out on computers with subjects typing their choices into software programmed using z-tree (Fishbacher, 2007).

Table 1 describes some of the demographics of the subject pool. They were almost all students, most of whom were studying at the undergraduate level (65%). The majority of participants were in the 16-25 category (68%). There were slightly less female subjects (45%) than male subjects and the majority had studied maths at some level (77%).

Once registered for the experiment, subjects were randomly allocated a computer and were asked to read the instructions for the game. After reading the instructions, they then played the game repeatedly for 30 periods, they completed a short questionnaire and then received their payment. The sessions took roughly 90 minutes to complete from the time they arrived until the time they left and subjects earned on average A$19 in the relative treatments and $22.50 in the fixed treatment. 1000 Experimental Currency Units (ECUs) were worth A$3 in the individual treatments and A$1.5 in the relative group treatment (the figures in table 3 are adjusted for the group treatment so they are all comparable with the ECUs of the individual treatment). Those who arrived late and not required to participate received a show-up fee of A$7 and were allowed to leave.

4.3 Results

Table 4 summarizes the results of the experiment. On average subjects declared a much higher fraction of their gross profit than predicted (compared with the equilibrium prediction, 0). The relative rule treatments had much higher declarations than the fixed treatments (443 and 394 compared with 119). In the market stage however, all three treatments had average quantities close to the Cournot solution ($q = 3$). Interestingly, the $RELIND$ treatment was actually less competitive on average than the $FIXIND$ and $RELGROUP$ treatments, contrary to what the theory predicts. In what follows, we discuss

Averages	Fixed		Relative		
	$FIXIND$	Predicted	$RELIND$	$RELGROUP$	Predicted
Quantity, q_i	3.048	3	2.936	3.241	3
Gross Profit (Before Tax), Π_i	610	655	611	603	255
Declaration,d_i	119	0	443	394	0
$decfrac$	0.204	0	0.739	0.659	0
Net Profit, U_i	502	557	420	421	217
Total Net Profit, $\sum_{t=1}^{30} U_i$	15,062	16,703	12,592	12,616	6,503
Total Gross Profit, $\sum_{t=1}^{30} \Pi_i$	18,303	19,650	18,336	18,093	7,650
Total Declared, $\sum_{t=1}^{30} d_i$	3,583	0	13,296	11,835	0
Total tax revenue (including fines)	3,241	2,948	5,745	5,477	1,148
Frequency of audits (%)	11.7	15	10.9	12.3	15
Subjects' pay (A$)	22.59	25.05	18.89	18.92	9.75
Fraction that evade (%)	81.5	100	48.8	66.8	100

Table 4.2: Descriptive overview

the descriptive evidence as it relates to each of our three hypotheses in turn.

4.3.1 Hypothesis 1: Tax evasion is lower under the relative rule than under the fixed rule.

Recall that the predicted declaration in all three treatments is 0. In all treatments the subjects were, on average, much more truthful than predicted and especially so in the two relative treatments. Figure 1 presents the distributions of the results from the tax stage. The variable $decfrac$, is the fraction of total gross profit declared over the whole experiment in each market divided by the by the total gross profit in each market. We can see that although in the fixed treatment, the distribution is highly skewed to the right, which is what we would expect given the predicted declaration is 0, the distributions of the two relative treatments are highly skewed to the left. On average, $decfrac$ was quite high, given our theoretical predictions: in the $RELGROUP$ treatment the mean of $decfrac$ was 0.659, and even higher in the $RELIND$ treatment,

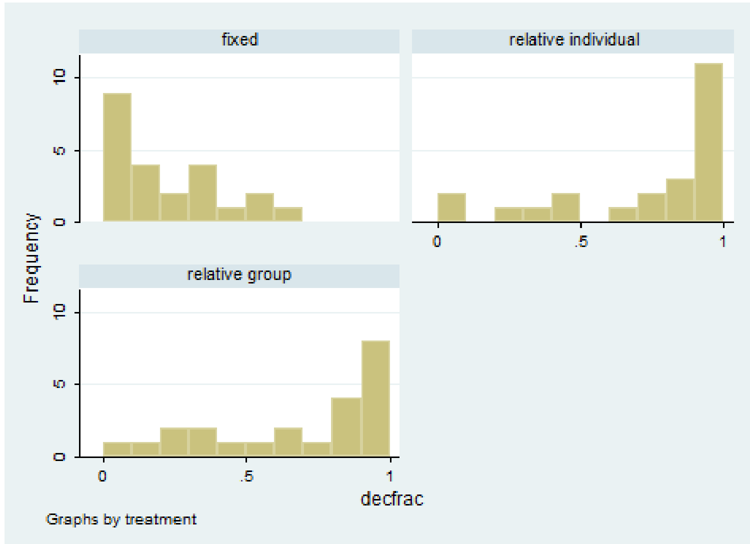

Figure 4.1: The total gross profit declared for each treatment

0.739. If we take into consideration that the gross profits were on average much higher than expected, the amounts declared were a lot lower than what the best response would predict. To get a crude idea, substituting the average gross profit values into the best response equation or the equilibrium condition for the relative treatments, we should see subjects declaring their entire earnings from the first period. However, it appears there was a considerable amount of coordination between subjects in these treatments since we observe values much lower. Although this was the case, subjects did declare more under the relative rule than under the fixed rule. Applying the Mann-Whitney U statistic to compare $decfrac$ for each treatment pairwise, the fractions declared in the two relative treatments are both significantly different to the fixed treatment (p=0.000 and 0.000). We can say that under the relative rule, subjects declared a significantly greater fraction of their income than under the fixed rule with the same expected number of audits. The data support our first hypothesis. The relative rule works well on the first dimension, reducing tax evasion compared

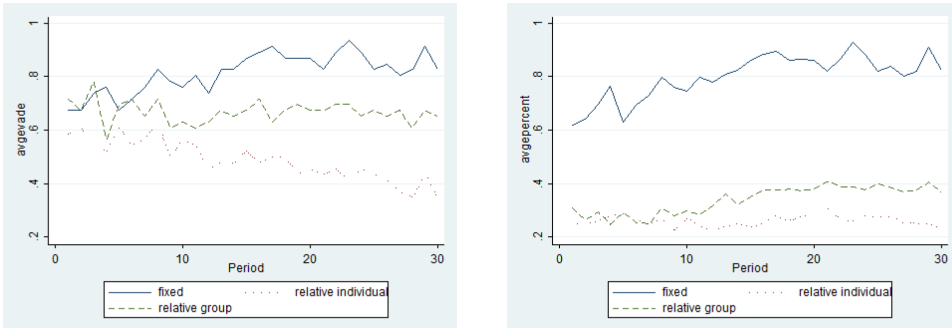

Figure 4.2: The average fraction of subjects evading per period for each treatment (left) and the average fraction of gross profit evaded per period for each treatment (right)

with the fixed rule.

Also, we observe the tax behaviour of subjects over time. Figure 2 (left) is a plot of the fraction of subjects evading, that is, the fraction of subjects declaring less than their true gross profit, for each period. Under the fixed rule, the fraction of evaders tends slightly upwards over time. Under the relative rule however, in the $RELGROUP$ treatment the fraction of evaders stays roughly constant and in the $RELIND$ treatment, it actually decreases over time. The $RELIND$ series and the $RELGROUP$ series are both below the fixed treatment for almost the whole experiment. It appears from a descriptive standpoint, that the data indicate that the fraction of people evading is lower under the relative rule and we see no sign of convergence. In the $RELIND$ treatment, if anything, the dynamics tend away from the $d = 0$ equilibrium. Finally, Figure 2 (right) plots the average amount evaded as a fraction of gross profit for each period. There is a stark contrast between the fixed treatment compared with the two relative treatments. Although the percentage of gross profit evaded lies between 20% and 40% in the relative treatments, in the fixed treatment subjects fail to declare more than 60% of their gross profit. The descriptive evidence and the Mann-Whitney U statistics support our first hypothesis.

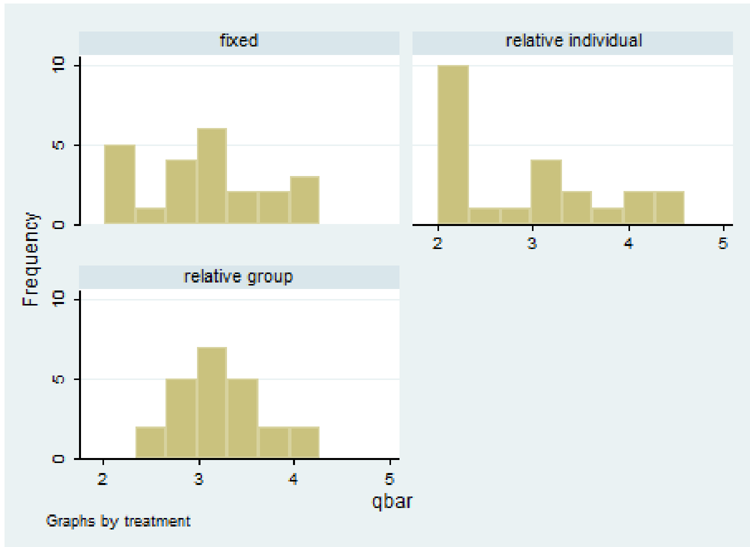

Figure 4.3: The distribution of the average quantity per group for each treatment

4.3.2 Hypothesis 2: Quantities under the relative rule are the same as under the fixed rule.

The variation in quantities chosen in the market stage between treatments is not quite so pronounced. Table 2 also shows the average quantities chosen by subjects in each treatment. Interestingly, $RELIND$ was actually the least competitive on average. The average quantity in the relative individual treatment was 2.936 which was *less* than 3.048 in the $FIXIND$ treatment. The $RELGROUP$ treatment however was higher, 3.241. The histograms of average quantities for each treatment show also that subjects in the $RELIND$ treatment were able to achieve more collusive quantities more frequently. Although the average quantities in the fixed treatment were fairly evenly distributed, the relative $RELIND$ treatment is skewed to the right; there are quite a few periods around the collusive, $q = 2$, mark. The $RELGROUP$ treatment however is bell-shaped around $q = 3$ which is greater than Cournot but not quite the

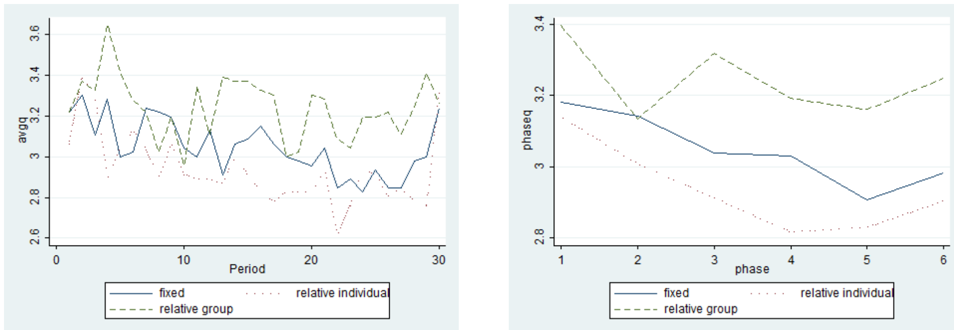

Figure 4.4: Average quantity each period for each treatment (left) and Quantities averaged over 5 period intervals for each treatment (right)

competitive, $q = 5$. Applying the Mann-Whitney U statistic to test the null hypothesis that the average quantity is different for each treatment pairwise, we see that the only two treatments that are significantly different are the two relative rule treatments, $RELIND$ and $RELGROUP$ (p=0.0161). This suggests that under a relative rule, the separation in decision making has an effect on market behaviour. There is no significant difference however, between the average quantity in the fixed treatment compared with the average quantity in either of the relative treatments.

Again, we also observe how behaviour evolved over time. Figure 4 (left) presents the time series of average quantities for each treatment and Figure 4 (right) presents the time series of these quantities for 6 phases, each phase being an average of q over 5 periods. All three treatment look very similar with $RELGROUP$ being the most competitive and $RELIND$ the least. With the quantities averaged over 5 period blocks, in Figure 4 (right), the difference in the treatments is more obvious. The dynamics do not appear to converge towards the predicted equilibria in either of the relative treatments, whereas the $FIXIND$ treatment seems to be fluctuating around the predicted Cournot outcome. We cannot say that the relative rule has an effect on market behaviour, especially since the quantities observed in the $RELIND$ treatment are actually *lower* than in the $FIXIND$ treatment. We have little evidence to reject the second hypothesis.

4.3.3 Hypothesis 3: Quantities chosen are the same under the relative rule for both owner-manager firms and firms with separate decision makers.

Comparing the two relative treatments, $RELIND$ and $RELGROUP$, the average quantities in each treatment are significantly different, contrary to what was predicted by the model. The Mann-Whitney U statistic, comparing the average quantities in the $RELIND$ and $RELGROUP$ treatments, is significant with a p-value of 0.016 and the quantities were actually lower in the $RELIND$ treatment. This seems surprising since it seems more likely that the owner-manager firms would be more able to see the connection between the market decision and the tax decision given that they make both decisions. However, we observe the opposite, the owner-manager firms are less competitive. This result suggests that the organisation of the firm does affect market behaviour. One potential explanation is that owner-manager firms have twice as many interactions with their competitors and this may be conducive to cooperative, collusive behaviour. The declarations however were higher in the individual treatment, so at best this could only be a partial explanation. Interestingly, although subjects in the $RELIND$ treatment were more collusive on average, their gains from the market stage were completely offset by their competitiveness in the tax stage. They were more truthful and evaded less. The resulting payoffs on average for each of the relative treatments were very similar.

4.3.4 Panel Models

The observed quantity choices from the market stage were a lot lower than predicted in the relative treatments. This means gross profits were higher than predicted. We now discuss whether subject behaviour is consistent with the equilibrium conditions for a rational decision-maker given the higher gross profit values we observe. Under a relative audit rule, the condition that describes the tax declaration in equilibrium is a linear function of the firm's own profit and the profit of its competitor. Given our parameterisation, the equilibrium declaration

$$d_i^*(d_{-i}) = \frac{2\Pi_i + \Pi_{-i}}{3t} - \frac{1-\alpha}{\beta}$$

becomes:

$$d_i^*(d_{-i}) = -850 + \frac{20}{9}\Pi_i + \frac{10}{9}\Pi_{-i}$$

We estimate this simple linear equation and test whether the coefficients are as predicted. Even though the predicted declarations are zero in equilibrium, given the amount of gross profit firms are making in the market stage, the optimum declarations should be higher than we observe. To get an idea, if we substitute the average gross profit values into the equilibrium condition, we get $d^* = 1187$ in the $RELIND$ treatment and 1160 in the $RELGROUP$ treatment. Therefore there should be close to perfect compliance. What we observe is much lower: 443 in the $RELIND$ treatment and 395 in the $RELGROUP$ treatment. Although subjects in the relative treatments declare much more than in the fixed treatment, it appears they are still less compliant than what we should observe if subjects behaved as rational, risk-neutral decision makers. The fixed treatment subjects however, are more compliant compared with our theoretical predictions.

The equilibrium condition for rational, risk neutral firms in the fixed treatment, is that firms should declare all or nothing depending on whether the expected value of the gamble (of evading) exceeds the after-tax profit from declaring everything. This is the case in our model when the probability of being caught, α, is less than the tax rate t, which it is ($\alpha = 0.15 < t = 0.3$). The declaration amount in the fixed treatment therefore should not depend on the outcome of the market stage; that is, declarations should not be a function of Π_i or Π_{-i}, assuming risk neutrality.

Relaxing the assumption of risk neutrality, a firm's declaration could depend on own profit but not on the rival firm's profit. Whether or not the firm's declaration is a function of own profit depends on the curvature of firms' utility functions in the aggregate; that is, the firms' risk preferences. It is possible, that declarations could be increasing or decreasing in Π_i or a more complex function of Π_i in the aggregate. Individual firms will always declare all or nothing (or be indifferent between the two) however.

We estimate the equation using a random effects panel model, censoring the dependent variable d_i to values within the interval $[0, \Pi_i]$. Assume there is a latent variable, d_i^*, call it willingness-to-declare, which is not observable for certain values. For any values below zero, we can only observe $d_i = 0$ since it is not possible to declare negative profit. For values greater than Π_i, we observe $d_i = \Pi_i$, since there is no benefit in declaring more that what is earned. The results are presented in Table 3.

The coefficients on own profit are all positive and significant as predicted. In the relative treatments the coefficients on rival's profit are also positive and significant and they are smaller than the coefficients on own profit. However, the magnitudes are significantly smaller than the 2.222 and 1.111 that the equilibrium condition predicts. We cannot reject their relative size, that the coefficient on own profit is twice that of the coefficient on rival's profit, in the $RELIND$ treatment. In the $RELGROUP$ treatment however, the coefficient on rival's profit is *larger* than the coefficient on own profit. Subjects in the $RELGROUP$ treatment appear to be more sensitive to a change in their rival's profit than a

Dependent variable: d_i						
	Fixed Individual Treatment		Relative Individual Treatment		Relative Group Treatment	
gprofit	1.225**	0.373	0.434**	0.691**	0.293**	0.693**
	(0.4094)	(0.3501)	(0.1191)	(0.0920)	(0.0595)	(0.0581)
gprofitother	-0.852*		0.257*		0.400**	
	(0.4159)		(0.1208)		(0.0618)	
profitgap		0.852*		-0.2570*		-0.3996**
		(0.4159)		(0.1208)		(0.0618)
period	-34.058*	-34.058**	-0.4145		-3.962**	-3.962**
	(4.3606)	(4.3606)	(1.0295)		(0.5947)	(0.5947)
maths	-358.032	-358.032	-114.180		86.601	86.601
	(398.2701)	(398.2701)	(104.739)		(104.8154)	(104.8154)
age						
26-30	821.933	821.933	198.723	198.723	70.881	70.881
	(514.4869)	(514.4869)	(172.6124)	(172.6124)	(108.9501)	(108.9501)
31-40			116.622	116.622	226.422	226.422
			(237.8482)	(237.8482)	(168.1242)	(168.1242)
41-50	1093.966	1093.966	2338.542	2338.542		
	(1081.952)	(1081.952)	(46531.21)	(46531.21)		
51-60	8058.954	8058.954	126.6399	126.6399		
	(655580)	(655580)	(445.3257)	(445.3257)		
>60					1761.72	1761.72
					(21226.15)	(21226.15)
Study level						
Post grad. Research	1013.573	1013.573	107.914	107.914	-7.977919	-7.977919
	(713.5725)	(713.5725)	(191.5012)	(191.5012)	(127.8526)	(127.8526)
Undergraduate	281.580	281.580	-52.0264	-52.0264	31.1617	31.1617
	(557.8303)	(557.8303)	(162.509)	(162.509)	(109.3764)	(109.3764)
Vocational					-444.264	-444.264
					(296.3147)	(296.3147)
Male	-563.1005	-563.1005	-110.3271	-110.3271	-82.42062	-82.42062
	(342.1758)	(342.1758)	(138.8031)	(138.8031)	(76.21287)	(76.21287)

Table 4.3: Interval regression. Standard errors are shown in brackets. * (**) significant at 5% (1%)

change in their own profit when making declarations. In the $FIXIND$ treatment, the coefficient on

4.4 Conclusion

The effect of tax audit policy on the market behaviour of firms has received relatively little attention. Bayer and Cowell, and Chapter 2 and 3 of this thesis show theoretically how audit policy might affect market behaviour. We have taken the theoretical model to the data to identify variation in behaviour that might be attributable to the choice of audit rule. We conducted two treatments with a relative audit rule and one with the fixed audit rule. We found that the audit policy had a significant effect on tax declarations. Subjects in the relative rule treatments reported a significantly higher fraction of their income despite all treatments having the same expected frequency of audits. The difference in quantities in the market stage however is actually not significant however, so our second hypothesis stands. Also, the structure of the firm is important. Owner-manager firms were significantly less competitive in the market stage than firms with the CEO/CFO arrangement for market and tax decisions.

To understand the observed behaviour at the market stage in our experiment, we need to move beyond the rational, risk neutral decision maker to a more complex behavioural analysis. This is a possible direction for future work. We have considered only the case of duopoly competition and have allowed firms to interact repeatedly over 30 periods. These two features create an environment which is conducive to cooperative behaviour and collusion. It is possible that more firms and more uncertainty with respect to who firms are having to compete with, would reduce their ability to cooperate, and lead to greater differences between treatments, and possibly, more competitive behaviour in the relative treatments as predicted.

Chapter 5

Conclusion

To conclude, we summarize the three main chapters (2-4), reiterate the main results from each and then discuss some of the practical implications and policy implications that might be drawn from the results. We discuss some miscellaneous issues about the analysis and some common questions. To finish, we briefly comment on possible directions for future work.

5.1 What we have done

Chapter 2

In Chapter 2 we asked: Does tax audit policy have different effects in different types of markets? We looked at two audit policies. Firstly, the relative audit rule. This rule is designed to capture the fact that the tax authority can condition its choice about which firms to audit on information that is available from tax reports. Secondly, we contrasted this audit rule with the fixed audit rule where firms are selected randomly by a fixed probability. We classified different markets by how the actions of each firm affect the profits of other firms in the market. We found that under a relative audit rule, while there are some markets where firms are more efficient, there are others where firms are less efficient than they are when a fixed rule is in place. In addition, the relative rule creates an incentive for firms to sabotage the profits of other firms in the market, which does not exist under a fixed regime. Finally, we said that under the relative audit rule, although the fraction of profit that goes undeclared is lower than it is under a fixed rule, the absolute amount of tax revenue may be higher or lower depending on how the audit policy affects firms in the market stage.

Chapter 3

In Chapter 3 we asked: What happens when the tax authority focusses all of its audit resources on the most 'at risk' firm? We examined an audit rule where an audit is conducted on only the firm that declares the least amount of profit in their tax report. We find that the model has multiple equilibria and some of these equilibria are asymmetric—meaning that although firms are a priori identical, in equilibrium one firm may produce more than the other. We also find that there exists a mixed strategy equilibrium where firms alternate between high and low quantities. Not all firms are more competitive under this rank order rule, but overall, the market is more efficient than under a relative audit rule or a fixed audit rule, at least for our linear duopoly case.

Chapter 4

In Chapter 4 we took the model to the data. We wanted to see how the predictions of the model would compare with the behaviour of subjects in a laboratory. We ran three treatments: one with owner-manager firms under a fixed audit rule, one with owner-manager firms under a relative audit rule and one with CEO/CFO firms under a relative rule. We found that the fraction of firms evading tax was significantly lower in the relative rule treatments than it was under a fixed rule and that the average amount evaded was also significantly lower in the relative rule treatments. In terms of market behaviour, there was no significant difference between the quantities produced in either of the two relative rule treatments compared with the fixed rule treatment. In fact, we found that the owner-manager firms were actually less competitive under the relative rule than they were under the fixed rule, although the difference was not significant. Finally, there was a significant difference between the two relative rule treatments. The CEO/CFO firms were significantly more competitive than the owner-manager firms under a relative rule.

5.2 Policy/Practical Implications

There are a few policy or practical implications that can be drawn from these results.

Chapter 2

Firstly, from Chapter 2 we might suggest the following. Firstly, in contrast to much of the literature, we should re-emphasise the point that is made in work by Bayer-Cowell: there are good theoretical reasons why audit policy can have real effects since it can affect the production activities and market decisions that firms make. The tax authority may be able to take advantage of this when setting policy because if they use reported information to condition the way audit resources are allocated, they can potentially affect not just tax reporting behaviour but market behaviour also. We argue that a relative rule is a closer approximation to the way tax authorities often select firms than the fixed audit rule which appears in much of the literature on corporate tax evasion. Secondly, the effect of conditioning audit selection on information available in tax reports is not always positive for market efficiency and policy makers may need to take this into account. There are some types of markets where conditioning audit selection may lead firms to be less competitively and the market to be more inefficient. The effect that audit policy has on market behaviour depends on how firms' actions affect each other. If a firm can reduce the profits of its competitors by behaving more competitively or in a more socially efficient manner (by lowering prices or increasing production for example), then the relative rule could have a positive effect. However, there are also cases where firms can reduce their competitors' profits by behaving less competitively or even by behaving maliciously. In these cases the relative rule should have a negative market efficiency effect. To understand the effect of a certain audit policy, one must understand the type of market it is being applied to and how firms' actions affect each other within that market.

Chapter 3

Secondly, from Chapter 3 we might draw the following implications. We argued that under certain conditions a rank-order rule may be the optimal policy choice and that this type of policy is a limit case of the relative rule. We see in simulations in Bayer and Cowell (2009), in a Cournot market, that as the sensitivity of audit rule increases, firms become more competitive and the market becomes more efficient. However, pushing the audit rule to the extreme, where all of the audit weight falls on the firm declaring the weakest information, the behavioural predictions of the model in equilibrium are not so obvious. There are multiple equilibria and the model suggests that different firms choose different actions in equilibrium. The market effect is still stronger overall however, at least in the linear version of the game. The policy implication is that in situations where a rank-order rule is optimal in terms of its effectiveness in collecting tax revenue, it may be the best policy option from a market efficiency perspective too (depending on the market it is applied to—Chapter 2). However, since the model has multiple equilibria, we do not have a unique prediction of the equilibrium behaviour of individual firms. If the tax authority would like to understand the effect of a rank-order-type rule on the behaviour of a particular firm, already for a perfectly rational firm, this is not clear. In models with more complex behavioural assumptions or for empirically observed behaviour, it may become even more complicated.

Chapter 4

Finally, from Chapter 4 we suggest the following. Firstly, our result that subjects conceal less income under a relative rule than under a fixed rule, suggests that a relative rule is a better policy option if the only objective is to reduce tax evasion. Secondly, we found no significant difference between the market behaviour of subjects facing a relative audit rule and those facing a fixed audit rule. Our subjects may not have been able to understand the relationship between their tax decisions and their market decisions. Whether or not managers of real firms would be able to understand this relationship is an open question however—we cannot make any conclusions about this. Thirdly, we did see that under the relative rule, when firms had 2 decision makers, they were significantly more competitive than the firms with a single decision maker even though the incentives of subjects were still completely aligned. The implication of this result is that firm structure does play a role, although it is not clear yet exactly how this works.

5.3 Limitations and Caveats

We now address some miscellaneous issues to clarify the interpretation of our analysis and to address some commonly asked questions. We briefly discuss the following issues: partial equilibrium analysis; risk preferences; what happens when the audit rule is not made explicit; collusion and dynamics; other policy tools and considerations; and lastly, tax avoidance.

5.3.1 Partial equilibrium analysis

Firstly, we have modeled the corporate tax problem by examining a single market in isolation—abstracting away the rest of the economy. Partial equilibrium analysis requires that there are no income effects on the market, or if there are, that the income effects are small enough to justify the use of a downward sloping demand function and consumer surplus as a measure of welfare. A consumer problem where preferences are quasi-linear—where the good produced in the market is additively separate—yields a demand function which is independent of income so that all income effects are captured by the numeraire good. Marshall & Marshall (1920) suggests that such a utility function is justified when the good produced makes up only a small share of each consumer's expenditure. Vives (1987) formalizes this idea, providing conditions such that, firstly, demand is downward sloping when the number of goods is large and secondly, that Marshallian consumer surplus is a good approximate measure of welfare as the number of goods increase. Loosely speaking, consumer preferences must not be too asymmetric and any two goods must not be too close to being perfectly substitutable, and a curvature property on the utility function is required so that the income derivatives of demand are not unbounded. See Chapter 3 of Vives (2001).

5.3.2 Risk preferences

Throughout the thesis, we have assumed that firms are risk neutral. This is a common assumption since owners or shareholders can diversify risk and should therefore invest in profit maximising firms. Nevertheless, the analysis does not crucially hinge on risk preferences as long as firms are not overly risk seeking—which seem reasonable. There are papers that consider the effect of tax rates and penalties on tax evasion (for example, the "tax payer as gambler" type models mentioned in Chapter 1) where risk preferences are crucial, but this is not something that we do here. The effect of risk aversion in our model is that the firms' payoffs become more convex. Interior solutions still exist and all of the results carry through. This is important because the assumption that firms behave as if they are risk neutral is strong and there are many reasons why it may not hold. For example, Choudhard & Levine (2009) and Asplund (2002) list reasons including: non-diversified ownership of the firm, liquidity limitations and uncertainties in employment dynamics. Managers' incentives may not be completely aligned with those of the owners either. The managers may behave warily if they are risk averse themselves and if their pay is performance based. In any case, the main qualitative results carry through when firms have other risk preferences as long as firms are not overly risk seeking.

5.3.3 Static vs dynamic settings and collusion

The framework we use in this thesis is static. In a dynamic, repeated game firms may be able to sustain actions that result in higher payoffs. We therefore briefly discuss the implications of collusion. Bayer and Cowell (2009) discuss for the baseline model with standard Cournot competition, three types of collusion. Firms may collude at the tax stage, the market stage or at both the tax and market stages. Firstly, if firms collude (maximise joint payoffs) at only the tax stage, their declarations under a relative rule will be the same as under a fixed rule if firms are symmetric. If firms have different gross profits then their declarations depend on the evasion costs. Secondly, if firms collude only at the market stage, then the relative rule has no effect on market decisions but their declarations in the tax stage will still be higher than under the fixed rule. Finally, if firms are able to collude in both stages of the game, then the relative rule and the fixed rule lead to the same outcome.

5.3.4 No audit rule

One common question is: what happens if the tax authority does not specify an audit rule? Alternatively: what happens when the audit rule is not common knowledge? There are two points that can be made. Firstly, firms may be able to estimate the audit rule depending on how much information they are able to observe, in which case, their behaviour should be the same. Secondly, if firms are unable to estimate the audit rule, then the effectiveness of not specifying an audit rule compared with the case of (for example) a relative rule, depends on firms' preferences over ambiguity. Where firms exhibit a high degree of aversion to ambiguity, a secretive audit selection policy would be preferable to one that is common knowledge, ceteris paribus. Indeed rarely, if ever, would a tax authority reveal the exact function with which it selects firms for audit even though they might reveal, for example, that they use discriminant analysis. This suggests that tax authorities believe that there is some value in not revealing the exact audit rule but there is still value in revealing that selection is based on differences in reports.

A related question is: what if the tax authority cannot pre commit to an audit policy? It may be that, once the tax authority has decided on an audit policy and has observed the information in reports, that it is better for them to select firms differently from what they would under the method of the original audit policy. For example, under the rank order rule, in equilibrium the firm that receives an audit is the firm with the lowest declaration. However, this firm does not evade tax, they declare truthfully. The firm that declares marginally more does not declare truthfully yet they do not receive an audit. If the tax authority were not able to pre commit to their audit policy and firms still behaved as if they could, the tax authority would revise their audit rule. There is a strand of literature (Reinganum & Wilde (1986), Graetz et al. (1986), Melumad & Mookherjee (1989), and Erard & Feinstein (1994)) that assumes the tax authority cannot pre commit to one particular audit policy and analyses the sequential game between taxpayers and the tax authority. Tax payers report their income first and then the tax authority chooses whether to audit second. The audit rule is then endogenous so the probability that the tax payer receives and audit depends on how the players respond to each other. There are different possible equilibria in these models. If we add this extra stage to our model, the results will depend on assumptions about the objective function and constraints of the tax authority. The relative rule and the rank order rule already try to mimic the behaviour of a tax authority behaving strategically but in a way that is similar to that described in Chapter 1. Extending the model to include different assumptions about the tax authority may be task for future work.

5.3.5 Impact of other policy tools

We should mention that we have examined our three audit rules in isolation from many other policy issues. Understanding audit policy in the context of all of these other issues and how they interrelate is important. For example, issues like: the choice of tax instruments, the incidence of tax, the optimal allocation of enforcement resources, the extent of enforcement, and optimal penalties and how all of these issues impact on efficiency and equity need to be taken into consideration. Surveys contain discussions of these issues (see Slemrod & Yitzhaki (2002), Andreoni et al. (1998), and Cowell (2004)).

5.3.6 Tax avoidance

In Chapter 1 we mentioned that the difference between tax evasion and tax avoidance is that tax evasion is illegal. In practice however, the distinction is not always clear cut. Some practices might be classified as evasion if they were to be examined in a court of law. However legal costs or the cost of gathering sufficient evidence might be prohibitively high meaning that illegality cannot be easily established. For the purposes of this thesis, we have only dealt with tax evasion but there are models that study tax avoidance or tax avoidance together with tax evasion. See for example Slemrod & Yitzhaki (2002), Mayshar (1991), Cowell (1990), Alm (1988), and Cross & Shaw (1982). Tax avoidance is often a prominent topic in the media, especially stories about the low effective tax rates achieved by large multinational corporations. It is important to mention that since our model assumes that all firms operate within the same tax jurisdiction, we implicitly assume they do not have the ability to rearrange their operations or employ complicated transfer pricing strategies or other practices in to reduce their tax bill. We assume that firms do not have the option to do this.

5.4 Future Work

Finally, there is more work that can be done. We have recently run three more treatments in the laboratory. We have run a treatment with the fixed audit rule with CEO/CFO firms as well as running two treatments with a rank order rule—one with owner-managers and the other with CEO/CFOs. This gives us 6 treatments in total that vary along two dimensions: 3 different audit rules, and 2 different organizational structures. Once we look at the data from these three treatments we should be able to gain an understanding of how subjects behave under the rank-order rule and get a better idea of how the internal organization of the firm affects decision making. We also wish to eventually run some treatments with business professionals to see whether their behaviour differs from that of tertiary students.

In terms of theory, the tax problem can be thought of as a contest between the firms. From a mechanism design perspective, the natural question to ask is what is the optimal way to design the contest (the audit rule) to elicit the largest declarations or the most competitive market? This is a (very difficult) question for future work.

Finally, applications of the phenomenon studied in this thesis, the idea that the tax policy can affect market behaviour, are undoubtedly more varied and numerous than our work so far would suggest. It is likely that there is scope for future work on different applications also.

Bibliography

Allingham, M. G. & Sandmo, A. (1972), 'Income tax evasion: A'.

Alm, J. (1988), 'Uncertain tax policies, individual behavior, and welfare', *The American Economic Review* pp. 237–245.

Alm, J., Cronshaw, M. B. & McKee, M. (1993), 'Tax compliance with endogenous audit selection rules', *Kyklos* **46**(1), 27–45.

Alm, J., Jackson, B. R. & McKee, M. (1992), 'Estimating the determinants of taxpayer compliance with experimental data', *National Tax Journal* pp. 107–114.

Alm, J., McClelland, G. H. & Schulze, W. D. (1992), 'Why do people pay taxes?', *Journal of Public Economics* **48**(1), 21–38.

Alm, J. & McKee, M. (2004), 'Tax compliance as a coordination game', *Journal of Economic Behavior & Organization* **54**(3), 297–312.

Andreoni, J., Erard, B. & Feinstein, J. (1998), 'Tax compliance', *Journal of economic literature* pp. 818–860.

Asplund, M. (2002), 'Risk-averse firms in oligopoly', *International Journal of Industrial Organization* **20**(7), 995–1012.

Bamon, R. & Frayssé, J. (1985), 'Existence of cournot equilibrium in large markets', *Econometrica: Journal of the Econometric Society* pp. 587–597.

Bayer, R. & Cowell, F. (2006), 'Tax compliance and firms' strategic interdependence', *Distributional Analysis Discussion Paper 81, STICERD, London School of Economics, London WC2A 2AE.* .

Bayer, R. & Cowell, F. (2009), 'Tax compliance and firms' strategic interdependence', *Journal of Public Economics* **93**(11), 1131–1143.

Bayer, R. & Cowell, F. (2010), 'Tax compliance by firms and audit policy', *The University of Adelaide School of Economics Research Paper No. 2010-23* .

Becker, W., Büchner, H.-J. & Sleeking, S. (1987), 'The impact of public transfer expenditures on tax evasion: An experimental approach', *Journal of Public Economics* **34**(2), 243–252.

Cadsby, C. B., Maynes, E. & Trivedi, V. U. (2006), 'Tax compliance and obedience to authority at home and in the lab: A new experimental approach', *Experimental Economics* **9**(4), 343–359.

Christian, C. W. (1994), 'Voluntary compliance with the individual income tax: results from the 1988 tcmp study', *The IRS Research Bulletin* **1500**, 35–42.

Clark, J., Friesen, L. & Muller, A. (2004), 'The good, the bad, and the regulator: An experimental test of two conditional audit schemes', *Economic Inquiry* **42**(1), 69–87.

Cowell, F. (2004), 'Carrots and sticks in enforcement', *The Crisis in Tax Administration* p. 229.

Cowell, F. A. (1990), 'Cheating the government: The economics of evasion', *MIT Press Books* **1**.

Cross, R. & Shaw, G. K. (1982), 'On the economics of tax aversion', *Public Finance= Finances publiques* **37**(1), 36–47.

Erard, B. & Feinstein, J. S. (1994), 'Honesty and evasion in the tax compliance game', *The RAND Journal of Economics* pp. 1–19.

Gale, D. & Nikaido, H. (1965), 'The jacobian matrix and global univalence of mappings', *Mathematische Annalen* **159**(2), 81–93.

Graetz, M. J., Reinganum, J. F. & Wilde, L. L. (1986), 'The tax compliance game: Toward an interactive theory of law enforcement', *Journal of Law, Economics, & Organization* pp. 1–32.

Lee, K. (1998), 'Tax evasion, monopoly, and nonneutral profit taxes', *National Tax Journal* pp. 333–338.

Marrelli, M. (1984), 'On indirect tax evasion', *Journal of Public Economics* **25**(1), 181–196.

Marrelli, M. & Martina, R. (1988), 'Tax evasion and strategic behaviour of the firms', *Journal of Public Economics* **37**(1), 55–69.

Marshall, A. & Marshall, M. P. (1920), *The economics of industry*, Macmillan and Company.

Mayshar, J. (1991), 'Taxation with costly administration', *The Scandinavian Journal of Economics* pp. 75–88.

Melumad, N. D. & Mookherjee, D. (1989), 'Delegation as commitment: the case of income tax audits', *The RAND Journal of Economics* pp. 139–163.

Mittone, L. (2006), 'Dynamic behaviour in tax evasion: An experimental approach', *The Journal of Socio-Economics* **35**(5), 813–835.

Nash, J. F. et al. (1950), 'Equilibrium points in n-person games', *Proc. Nat. Acad. Sci. USA* **36**(1), 48–49.

Novshek, W. (1985), 'On the existence of cournot equilibrium', *The Review of Economic Studies* **52**(1), 85–98.

Office, U. G. A. (1999), 'Tax administration: Irs' return selection process', *Technical Report 99-33, GAO/GGD.* .

Reinganum, J. F. & Wilde, L. L. (1986), 'Settlement, litigation, and the allocation of litigation costs', *The RAND Journal of Economics* pp. 557–566.

Robben, H. S., Webley, P., Elffers, H. & Hessing, D. J. (1990), 'Decision frames, opportunity and tax evasion: An experimental approach', *Journal of Economic Behavior & Organization* **14**(3), 353–361.

Singh, N. & Vives, X. (1984), 'Price and quantity competition in a differentiated duopoly', *The RAND Journal of Economics* pp. 546–554.

Slemrod, J. & Yitzhaki, S. (2002), 'Tax avoidance, evasion, and administration', *Handbook of public economics* **3**, 1423–1470.

Sonnenschein, H. (1968), 'The dual of duopoly is complementary monopoly: or, two of cournot's theories are one', *The Journal of Political Economy* pp. 316–318.

Tan, F. & Yim, A. (2014), 'Can strategic uncertainty help deter tax evasion? an experiment on auditing rules', *Journal of Economic Psychology* **40**, 161–174.

Tarski, A. et al. (1955), 'A lattice-theoretical fixpoint theorem and its applications', *Pacific journal of Mathematics* **5**(2), 285–309.

Virmani, A. (1989), 'Indirect tax evasion and production efficiency', *Journal of Public Economics* **39**(2), 223–237.

Vives, X. (1987), 'Small income effects: A marshallian theory of consumer surplus and downward sloping demand', *The review of economic studies* **54**(1), 87–103.

Vives, X. (1990), 'Nash equilibrium with strategic complementarities', *Journal of Mathematical Economics* **19**(3), 305–321.

Vives, X. (2001), *Oligopoly pricing: old ideas and new tools*, MIT press.

Wang, L. F. & Conant, J. L. (1988), 'Corporate tax evasion and output decisions of the uncertain monopolist', *National Tax Journal* pp. 579–581.

Webley, P. (1987), 'Audit probabilities and tax evasion in a business simulation', *Economics Letters* **25**(3), 267–270.

Yaniv, G. (1988), 'Withholding and non-withheld tax evasion', *Journal of Public Economics* **35**(2), 183–204.

Yaniv, G. (1995), 'A note on the taxevading firm', *National Tax Journal* pp. 113–120.

Yaniv, G. (1996), 'Tax evasion and monopoly output decisions: Note', *Public Finance Review* **24**(4), 501–505.