Can Self-Regulation work?. A game-theoretic approach.*

May 18, 2001

Abstract

This paper analyses reputation-based incentives for self regulation from a principal-agent perspective, in a context of asymmetric information. Self Regulatory Organisations (SROs) have scant incentives to monitor quality and expose fraud because fraud exposure is often interpreted by Bayesian consumers as a bad signal of SRO quality. However, public regulation in parallel to Self Regulation can enhance SRO incentives to monitor quality and reduce fraud. Therefore, defying conventional wisdom, a mix of public and self regulation may be preferred because it would benefit from SROs informational advantage about quality, while public regulation would provide the incentives to monitor quality that may be absent otherwise.

1 Introduction

> The value, or WORTH of a man is, as of all other things, his Price..., and therefore is not absolute; but a thing dependant on the need and judgement of another...And, as in other things, so in men, not the seller, but the buyer determines the Price. For let a man (as most men do) rate themselves as the highest Value they can; yet their true Value is not more than it is esteemed by others. Thomas Hobbes, *Leviathan*, Chapter X.

Since Akerlof's (1970) seminal paper, the study of the implications of asymmetric information between consumers and suppliers about product quality has

*I am deeply indebted to John Vickers for his support, comments and insight. The author also thanks Peter Sorensen, Hyun Song Shin, Rupert Gatti, Meg Meyer and Tim Besley, who provided very valuable comments. All remaining errors are the author's.

constituted a very dynamic area of research. Almost 30 years after his contribution, there is a substantial amount of theoretical knowledge on many different aspects of this phenomenon. This vast research effort seems justified by the abundant evidence of consumer fraud world-wide and its severe consequences.[1] An important part of this research effort has focused on the fundamental question of whether the social costs of asymmetric information can be dealt with *within* and *by* the market system, or whether public regulation may be required. The literature has identified several ways in which the existence and implications of asymmetric information can be overcome by market mechanisms. In a classical work, Nelson (1970) introduced the now well-known distinction between *search* and *experience* goods. Despite the distinction, these two types of goods share the fact that quality is potentially observable, which, naturally, can reduce the extent and severity of asymmetric information. [2] Other flows of information described in the literature involve several forms of quality signalling mechanisms. By employing quality-related features such as pricing (Bagwell and Riordan (1991)), advertising (Nelson, (1974), Milgrom and Roberts (1986)), warranties (Cooper and Ross (1985)) and goods attributes (Spence (1977)), firms can often achieve separation of quality types, easing, therefore the problem of asymmetric information and enhancing trade efficiency.[3]

However, the role and effectiveness of these market mechanisms in overcoming the negative implications of asymmetric information can be subjected to several lines of criticism. From an empirical point of view, there is ample evidence suggesting that fraud and the provision of low quality are widespread practices in most industries world-wide, despite the existence of these market mechanisms.[4] Second, this should not be surprising, as this evidence is consistent with a great deal of the theoretical literature on market mechanisms. Indeed, in most of the contributions mentioned above, although asymmetric information is usually ameliorated, it still persists to a significant degree so that "full information" and efficient trade are rarely achieved in equilibrium.

However, there is a perhaps more limiting feature which cuts across most

[1] See Levi (1987), pp. 6-9 for some evidence. For example, in the UK the cost of fraud in 1985 was estimated at 3 billion pounds, somewhat short of 1 per cent of the GDP. In the last 20 years there have been around 130.000 cases of fraud in England and Wales every year (Home Office, 1997). There is also ample evidence that fraudulent behaviour and white collar crime in general is on the rise in many industrialised economies and in the developing world. However, this picture is probably too conservative because a great deal of fraud goes unreported, but also because the definitions of "social cost" fail to acknowledge "hidden" efficiency costs such as deals not realised, costly contracting and enforcement.

[2] In the case of search goods, consumers would have an incentive to bear some search costs and choose an optimal search level. This in turn provides incentives for the sellers to increase quality beyond the situation involving no search by consumers. In the case of experience goods, there are two fundamental ways in which the asymmetric information problem can be overcome. First, since quality is observable *ex-post* then sellers and consumers can agree on a quality-contingent contract where the seller is liable for negligence or low quality. Second, if there is the possibility of repeated purchases by consumers, then sellers may have an incentive to build a reputation for good quality among consumers by supplying higher quality. See Shapiro (1982, 1983).

[3] See Spulber (1989) for some examples and references.

[4] See footnote 1 and Levi (1987) for some evidence.

of the literature on market mechanisms. This is the fact that a fundamental assumption for most of these market mechanisms to exist and operate is that quality must be ultimately *verifiable* by consumers, either prior or after purchase. This is, by definition, the case in search goods. It is also the case in experience goods, and consequently also the case in the literature on seller liability and contingent contracts, and on seller reputation. Moreover, most of the quality-signalling mechanisms mentioned above also require some degree of quality verifiability.[5]

The requirement of quality verifiability is not a satisfactory assumption for a wide type of goods and services in many industries. There are many situations in which quality is not observable even after purchase, and therefore the informational flows described above may be not only insufficient but irrelevant to overcome the informational asymmetry between firms and consumers. For example, consumers of wooden furniture may not know how "sustainably" the forests are being managed, or whether advice from a financial consultant to invest in an asset is being given truthfully because of its risk-adjusted profitability or just to serve some hidden agenda of the consultant firm. It is also difficult to figure out how "organic" a product is and what this term exactly means and how to recognise such attribute, and it is hard to check how much effort lawyers are putting into one's judicial case. [6] In other industries, quality may be potentially verifiable *ex-ante*, but the search costs may be prohibitive, or perhaps the good may not be purchased often enough to justify investing in costly search, or to allow consumers to learn from experience. In conclusion, in such industries observation of quality is not a real possibility, and therefore other means of enhancing trade efficiency must be sought in such cases. In other words, the real challenge ahead is to find means of establishing effective consumer protection *beyond* what consumers can do for themselves.

1.1 Public Regulation

The natural alternative to the market mechanisms described above is public regulation. Regulation of quality can take several forms. Among the most common are $certification$ and $licensing$.[7] By reducing asymmetric information,

[5]For example, much of the literature on advertising as a quality signal require the traded good to be an experience good, the central idea being that if advertising is more valuable when resale is higher, then advertising may be a useful signal to convey quality information to consumers. See for example the seminal paper by Nelson (1974) and the survey in Spulber (1989).

[6]In the literature, these types of goods are often referred to as "credence goods". See for example Emons (1997).

[7]The former involves basically the public enforcement of certain quality standards by the certified suppliers, plus some device to inform consumers that such criteria were actually met. This form of regulation aims fundamentally at informing consumers about quality in an industry characterised by heterogeneous suppliers. The latter, instead, involves the granting of permission to operate in an industry, subject to the fulfilment of some minimum established criteria. Therefore, licensing aims at limiting entry into the industry exclusively to quality suppliers. In both cases, regulation can be based on *input* or *output* regulation.Examples of input regulation are the requirement of a minimum level of human capital investment by

both forms of public regulation enhance trade efficiency.

Public regulation has many limitations, however. One of the most significant ones is that public enforcement often requires the establishment of verifiable standards that may be difficult to institute in many industries. For example, most professional services are not standardised but instead are often tailored to the consumer's needs, so that minimum quality standards will be virtually impossible to establish, and quality may be difficult to assess. In some cases there might be unforeseeable contingencies that make quality standardisation impossible. In other cases quality standards, either on inputs or outputs, may be only imperfectly related to overall quality, which may generate inefficiency by distorting the producer's incentives towards overall quality provision.[8] Another, perhaps more fundamental problem of public regulation is that often the regulator is poorly informed about quality in comparison to the suppliers. This problem may be particularly acute in industries such as the professions and in financial markets. In these circumstances, the problem will not only be that the enforcement technology of the public regulator may be very costly and ineffective. In addition, optimal regulation is likely to imply a trade-off between informational rents and efficiency, both of which are socially costly.[9] Finally, the enforcement of public regulation is often delegated to institutions quite independent of the public agency itself, which often leads to agency problems in the enforcement of quality.[10] These limitations suggest that other means to achieve efficiency in industries with asymmetric information must be sought. A concept that has been thought to come to the rescue is that of self-regulation, which is explored next.

1.2 Self-Regulation

Self-regulation has often been proposed as an alternative to address the problems of asymmetric information, and it exists in many industries. Self Regulatory Organisations (SROs) have been claimed to have many advantages over public regulation since they do not exhibit many of the limitations of public regulation mentioned above. First, SROs are usually better informed about quality than public agencies, and therefore they are thought to be more efficient at finding fraud and wrongdoing than public regulators. Also, because of their informational advantage SROs would avoid the costly trade-off between informational rents and efficiency existing in public regulation. Second, since SROs enforce quality themselves, they lack the agency problems that emerge from the delegation of quality enforcement. Finally, public regulation enhances trade efficiency and therefore its benefits are appropriated mostly by consumers and suppliers.

professionals. An example of output regulation is the requirement that a certain product is free of some harmful substance. See Shapiro (1986).

[8]See for example Spulber (1989).

[9]See Laffont and Tirole (1993) for examples of trade-offs between informational rents and efficiency in the context of regulation.

[10]See for example Besanko and Spulber (1989).

Therefore, from a distributive perspective, self-regulation is preferable to public regulation because the costs of quality regulation are borne by the favoured parties, whereas in public regulation costs are financed from public funds.

There are many examples of SROs. For instance, the self-imposition of quality standards such as in "organic" foodstuffs, are usually enforced by the firms themselves. In many professions, practitioners are often required to be registered in a club which is intended to enforce and guarantee a good quality service to customers. One of the industries where self-regulation has been most common is in financial services. In the UK, for example, the Financial Services Authority (FSA) (formerly the Securities and Investments Board) currently recognises three SROs, namely, the Investment Management Regulatory Organisation (IMRO), the Personal Investment Authority (PIA) and the Securities and Futures Authority (SFA).[11] They have the power to authorise firms to conduct investments business in the UK.[12] A remarkable feature of financial self-regulation in the UK is that the SROs have legal immunity from lawsuit in respect to their regulatory activities.[13]

Self-regulation is also common in many professions. For example, in addition to the SROs above, the FSA also recognises nine Recognised Professional Bodies (RPBs).[14] These organisations are in charge of setting and enforcing professional disciplinary standards among their members. Also, the General Medical Council and the United Kingdom Central Council for nursing regulate professional practice in the health sector. Finally, engineers and other trades such as advertising, among many others, are often regulated by SROs.

However, self-regulation is a much broader practice than the examples mentioned above. In fact, perhaps the most common example of a SRO is just any firm. Indeed, presumably in the majority of firms, the firm's principal does not directly determine quality, but the agents who are supervised by the principal instead determine it. In that sense, firms are also SROs since the principal, as the party finally affected by the firm's reputation, has only partial information about the quality of the goods she sells, and has to consider how much to monitor the agents to preserve a reputation for good quality. Finally, self-regulation is by no means restricted to markets. For example, the principle of self-regulation can also be found in politics; political parties are certainly concerned by their collective reputation and would accordingly want to encourage diligent and hon-

[11] Two other SROs (Financial Intermediaries, Managers, and Brokers Regulatory Association (FIMBRA) and the Life Assurance and Unit Trust Regulatory Organisation (LAUTRO)) are in the process of winding down their regulatory responsibilities, following the creation of the PIA. However, new laws are being introduced to merge all the financial regulatory organisations into the FSA (Financial Services Authority (1997)).

[12] Under the Investment Services Directive they can also authorise firms to operate throughout the European Economic Area (Financial Services Authority (1997)).

[13] See Levi (1987, p. xvii.)

[14] They are the three Institute of Chartered Accountants (England and Wales, Scotland and Ireland), the Association of Chartered Certified Accountants, the Law Society and its two counterparts in Northern Ireland and Scotland, the Institute of Actuaries and the Insurance Brokers Registration Council.

est behaviour by its members. However, party "quality" (honesty, effort, and transparency) is by and large determined by the behaviour of individual party members, which other members acting as principals can only enforce partially.

There are two fundamental functions that SROs should carry out in order to function properly. First, SROs must *enforce* high quality provision among its members. Secondly, SROs must also *disclose* to consumers information of any evidence of fraud and wrongdoing by its members. These two functions play different roles; Quality vigilance has the purpose of ensuring quality standards, but disclosing information plays the role of informing consumers about the quality provided by different members of the SRO. There is no *a priori* reason why an SRO may either fulfil or fail to accomplish these objectives. These functions are the two yardsticks employed in this work to study and assess the performance of SROs.

However, an issue of considerable importance is whether SROs will have the adequate *incentives* to monitor quality effectively, as well as to inform customers about fraud and malpractice among its members. In fact, because SRO decisions about monitoring and information disclosure are *voluntary*, then, self-regulation will imply *regulatory capture* by definition as the outcome of self-regulation as regulatory scheme will be perfectly aligned with the private objectives of the SROs. This is particularly true in cases where the SROs have legal immunity against lawsuit for negligence in their regulatory activities.

At first, answering the question of whether self-regulation works may seem to require only an empirical approach. There is no doubt empirical analysis would be of invaluable help. However, there are two fundamental weaknesses of an empirical assessment of the performance of self-regulation. The first is the fact that fraud happens precisely *because* it is very difficult to detect by consumers. Indeed, "the point about fraud (and credence goods more generally) is that the offenders can manipulate the victims perceptions of 'what happened' so that it may never even occur to them that they have been victimized improperly or criminally".[15] In this context, "no news may be bad news" because fraud may be happening systematically but it is just not being detected and reported. Second, even if data on fraud were reliable, it is not clear whether evidence of fraud exposure would indicate that self-regulation is or is not working properly. In fact, *a priori*, substantial evidence of fraud exposure may be consistent with *high* and effective vigilance, but it could be also the result of *low* vigilance effort, from which widespread fraud is a consequence. Therefore, empirical research would still need theoretical foundations to interpret the evidence, let alone to explain it. In conclusion, in these circumstances theoretical insight and formal analysis of the SRO incentives towards vigilance and disclosure seem not only useful, but also essential.

The arguments in favour of self-regulation share the view that SROs should be concerned by their reputation; if reputation is valuable in an industry char-

[15]Levi (1987), p. 28. Text in parenthesis added.

acterised by asymmetric information about quality, then SROs will attempt to build a reputation of good quality. This would be done, firstly, by monitoring and enforcing quality, and secondly, by informing consumers about quality provision. A theoretical investigation of this hypothesis requires the analysis of how reputation can be built and sustained by the SRO. Reputation, in turn, would depend on the specific informational structure available to consumers from which expectations about quality are to be inferred and updated. SROs may have proper incentives under some informational structures but not in others. In the case of credence goods, the most important source of quality information available to consumers will be the *voluntary* and *direct* disclosure of quality information by SROs, understood either as individual firms or as a "club" of firms. Interestingly, the idea of self-regulation is usually proposed and found in operation precisely in credence-goods industries. In this work, the main source of information available to consumers is the direct and voluntary exposure by the SRO of fraud and wrongdoing done by its members.

1.3 Parallel Regulation

A pervasive idea in the regulation literature is that public regulation and self-regulation are seen and understood as *mutually exclusive* alternatives. This idea has not only been an implicit assumption in much of the theoretical literature, but also it has also been very influential in the design of regulatory policy in practice. In this view, the basic guiding principle is that public regulation is justified only if it can reduce asymmetric information more efficiently than "the market" or other private alternatives. For example, after surveying the literature on quality signalling and disclosure, Spulber (1989) asserts that;

> . . . it may be concluded that non observability of product quality or producer characteristics can never serve in itself as a justification for regulation or mandated disclosure. Rather, it must be demonstrated that incentives are absent for voluntary private disclosure.... To discern a potential role for regulation, it must first be established that the market does not supply consumers with the means to make informed choices and that the regulatory agencies can provide such information at less costs that private alternatives (p. 451).

However, there does not seem to be a reason why this must be the case, *a priori*. In fact, the view above rules out the possibility that *parallel regulation* (i.e. public and self-regulation in tandem) may provide a better alternative of dealing with asymmetric information. The implied suboptimality of joint public and self-regulation must not be an assumption. Instead, it must be an hypothesis to be investigated formally. Apart from the obvious costs of duplication, there are three possible implications of simultaneous public and private (self) regulation. First, public regulation would, *ceteris paribus*, lead to the detection of fraud that would otherwise have gone undiscovered and unpunished. Second, public regulation may change the behaviour of the SRO members, who ultimately determine quality. Third, parallel regulation may also alter the incentives towards vigilance that SROs face. The main conjecture and fundamental motivation for this line of enquiry is that parallel regulation can benefit from the informational advantage of the SROs, while it can provide the incentives for adequate quality enforcement and disclosure that the unregulated SRO may lack.

Perhaps as a result of avoidance of enforcement duplication and overlapping, there are relative few examples of parallel regulation in practice. In some cases a form of parallel regulation is exercised when the authority appoints inspectors to investigate possible malpractice.[16] In this case, however, parallel regulation is *ad hoc* and it is not exercised *ex-ante*, but only when some evidence of fraud already exists. Therefore, its effects are likely to be different to the conjectures above because the SRO would presumably wish to minimise everyone's *knowledge* of fraud, as opposed to the fraud *level*. The form of parallel regulation studied in this work resembles *ex-ante*, permanent public vigilance, in tandem with SROs.[17]

The views expressed above constitute the core of the concerns of the present work, and also the main ways in which it departs from the existing literature. This work begins by asking whether self-regulation can work, in terms of succeeding in enforcing high quality provision and disclosing quality information to consumers. The main distinguishing features of the analysis is that the incentives to enforce and disclose quality information are investigated starting from two fundamental assumptions. First, quality is not observable by consumers (i.e. credence good) and therefore the main source of information available to consumers is the disclosure of information by the SRO. Secondly, quality is neither an exogenous nor a choice variable, as in most contributions, but instead it is endogenously determined by principal-agent strategic interaction within the SRO. Third, this work explores the existence of possible complementarities between public regulation and self-regulation, and the potential benefits of having both schemes in parallel.

Before discussing the main ideas on SROs, it is important to begin by raising the fundamental issue of the motivations SRO should have, that is, what assumptions should be made about the objective function SROs might try to optimise. On the one hand, we might consider SROs to behave "altruistically" by exercising their enforcing power to achieve some degree of consumer protection. On the other hand, SROs might behave "selfishly" by attempting to maximise their "per capita income" or to maximise producer surplus. The point to make here is that it seems unreasonable, and certainly risky to base the system of

[16]This is the case, for example with the Secretary of State for Trade and Industry in the UK, who exercised these powers to investigate the notorious Guinness scandal in 1986. Levi (1987), Kay and Vickers (1988).

[17]For example, in airline safety and restaurant hygiene, public regulation co-exists with internal monitoring. See Kay and Vickers (1988) for an informal discussion of parallel regulation.

self-regulation on the assumption that SROs should care about consumer protection *per se*. Moreover, such assumption removes any meaningful motivations to carry out a theoretical enquiry on the incentives SROs face. In this work, it is assumed that SROs try to maximise a measure of income net of vigilance costs and do not exhibit any concerns for consumer protection *a priori*. The point is precisely to analyse whether a concern for quality and consumer protection can emerge endogenously from reputation-based incentives.

Reputation in economic theory is often understood and modelled in quite distinctive ways. Perhaps the most common taxonomy of it is the distinction between reputation driven by *incomplete information* versus reputation driven by *repeated interaction* in games of complete information. This work deals with reputation in the former sense; consumers wish to know the *type* of the SRO they are purchasing from, which may give the SRO an incentive to build a reputation of good quality. In contrast, Nunez (2000 B) deals with reputation in a repeated moral hazard framework.

The basic ideas on SRO in an incomplete information context are as follows. Suppose that quality is a function of the effort, skills and honesty of the SRO members. For simplicity, assume that quality is inversely related to the level of fraud or wrongdoing within the SRO, and that the direct (as opposed to reputational) costs of fraud are borne by consumers. Assume also that the SRO (the principal) can invest effort in investigating fraud by its members (the agents), and that fraud can be punished by the SRO by imposing some form of penalty. Assume further that consumers cannot observe the provision of quality (i.e. a credence good) and therefore the main source of quality information for consumers comes from the SRO itself. Finally, assume that the SRO is concerned by its reputation, because reputation determines the income of the whole profession. The central questions are; how does fraud exposure affect consumers' beliefs about the general "quality" of the SRO?, and secondly, what is the likely behaviour of the SRO in response to consumers' expectations?. In this context, fraud exposure can affect the SRO reputation in several ways. On the one hand, exposure can be a *positive* signal of the SRO vigilance level because, other things constant, the probability of exposure is increasing in SRO vigilance. On the other hand, exposure could be a *negative* signal about SRO vigilance. The reason for the latter is that if SRO members are better informed than consumers about the level of SRO vigilance, and if fraud is endogenously determined by SRO vigilance, then exposure may well be interpreted as the existence of little SRO vigilance. This, in turn, makes fraud more widespread and increases the probability of exposure. Third, exposure could again be a *positive* signal because it may mean that the average quality of the remaining members is higher once the "bad apple" has been identified. Finally, exposure can again be a *negative* signal about SRO quality since it could be informative about the propensity of the SRO members towards fraud and wrongdoing due to some form of group characteristic. From this informal discussion two issues emerge. First, it follows that consumers can have uncertainty about the SRO (principal's) type and/or about the SRO member's (or agent's) type. The first two effects above are related to uncertainty about the principal's type, while the latter two are related to uncertainty about the agents' type. This distinction may matter in terms of the inference process that consumers make, and therefore it may also have an impact on the incentives of the SRO towards vigilance and exposure. Second, the discussion suggests that in both types of consumer uncertainty presumably there will be a positive and a negative reputational effect of exposure. There are no obvious ways to tell *a priori* how exposure will affect the SRO reputation in each case, and what will be the optimal strategy the SRO would undertake.

This work addresses the situation where consumers care about the type of the SRO principal, dealing therefore with the first two effects mentioned above.[18] The basic structure of the model developed in the next sections is the following. A single agent chooses a level of fraud while the principal chooses a level of vigilance. This is modelled assuming Cournot and Stackelberg behaviour by the SRO, in terms of their internalisation of the effect of vigilance over the agent's fraud choice. I analyse whether and under what conditions the negative signal of exposure may dominate the positive one. If this is the case, then exposure would fail to signal good quality and achieve separation of types, and therefore we would have a self-regulation theory of cover-ups.

Next, suppose parallel regulation is introduced so that a public agency can investigate and expose fraud in parallel to, but independently of the SRO. I investigate the effect of parallel regulation on SRO vigilance, SRO disclosure decision and SRO member's fraud level. The effects of parallel regulation can be multiple. First, parallel regulation is likely to have a direct disciplinary effect on the SRO members fraud choice. Second, by threatening the SRO to find and disclose fraud, SROs may have an additional incentive to be vigilant and reduce fraud to avoid that possibility. Third, parallel regulation may increase the SROs incentives to expose fraud because exposure by the public regulator may damage the SRO reputation, and therefore the SRO might want to preempt the regulator by exposing first. However, if SRO exposure hurts its own reputation, then the SRO may decide to run the risk of not exposing, hoping that the regulator does not find any fraud.

1.4 Related Literature

Despite its relevance, the formal analysis of the economics of self-regulation has received relatively limited attention from theorists.[19] There is some applied literature on industry-specific case studies, but they often contain limited theoretical insight and in most cases no formal analysis. Some theoretical contributions have focused on the relationship between self-regulation and the incentives to

[18] Nunez (1999 A) develops another model addressing SRO incentives when consumers have uncertainty about attributes of the SRO members.

[19] See for example the brief survey in Gehrig and Jost (1995).

establish barriers to entry and limit competition within the industry.[20] These works, however, do not address issues of asymmetric information and the incentives of SROs towards quality enforcement and information disclosure, which constitute the main concern of this work.

In Shaked and Sutton (1981) firm's qualities are exogenously given, and a self-regulatory club determines the minimal quality necessary to enter the market. Their analysis focuses on the divergence between the interest of the club and that of a social planner. However, since quality is exogenous, they do not analyse the incentives for quality provision. Leland (1979) analyses the impact of licensing on professionals and consumers from an adverse selection perspective. Licensing is viewed as truncating the bottom of the exogenous quality distribution, which will have an impact on equilibrium price and quality provision in the market. He also asks the implications of having the profession choose the quality standard (self-regulation), as opposed to a social planner, concluding that standards are likely to be too high because higher quality raises market price. Again, since qualities are exogenously given, his work addresses adverse selection issues, as opposed to incentive (moral hazard) ones. In addition, quality standards are assumed to be always met by all firms in the industry, and there is no inquiry into the mechanisms whereby the standards are enforced or whether the profession would disclose any failure to meet the standards by any member.

Some authors have studied incentives for self-enforcement of quality standards. Gehrig and Jost (1995) show that the suppliers of experience goods will often have incentives to provide and enforce quality standards. This work differs from this work in many respects, however. First, the key assumption in their work is that SRO members sell *experience* goods and not *credence* goods since quality is observable after purchase. As argued earlier, this assumption is not satisfactory in many industries where quality cannot be assessed by consumers, and the main source of information for them has to come from the SROs themselves. In their work, the incentives for the enforcement of quality come from the possibility of consumer punishment in the future if bad quality is observed. If a member of the SRO (or "club") deviates and provides bad quality, then the expectation of future quality of all SRO members will decrease, and therefore the club (or SRO) will have incentives to enforce quality among its members. Second, unlike the case studied in this work, in their model SROs have no information disclosure decision to make, because quality is directly observable by consumers. Third, in their model firms choose quality directly, and therefore quality is not determined endogenously as the result of principal-agent interaction. Finally, their work sees self-regulation and public regulation as substitutes. Indeed, their work analyses the conditions in which each regulatory framework is preferred, as oppose to make an attempt to investigate potential complementarities between them.

[20]For example the influential work by Shaked and Sutton (1981).

Voluntary quality disclosure has been analysed by some authors. Milgrom (1981) shows that full disclosure will take place in a sequential equilibrium, and Milgrom and Roberts (1986) extend this discussion and conclude that firms will often have an incentive to disclose product quality to consumers.[21] However, as it is the case with most quality signalling literature, these contributions assume that quality is *exogenous*. In such setting, signalling does not alter intrinsic quality, and therefore the signals (direct quality disclosure, pricing, advertising or contract characteristics) are simply variables to be chosen by firms. This assumption is central to many of the results and claims in the literature, yet there are plausible reasons to expect that quality may be endogenous, and therefore that signals will also be endogenously determined as the result of strategic interaction within the SRO, as discussed above.

Shapiro (1982) has studied moral hazard issues in quality provision. His work considers the reputation of a monopolist that supplies a good of unobservable quality. However, in this model quality is directly chosen by the monopolist, and therefore it is not the result of strategic interaction between the SRO which is the main concern of this work. In addition, his work does not study either incentives for voluntary exposure of quality or parallel regulation.

This work is organised as follows. Section 2 presents the case of an unregulated SRO. Section 3 analyses equilibria under different circumstances. In section 4 parallel regulation is introduced and its effects are studied. Finally, section 5 presents the main conclusions of this work.

2 The Basic Model

There are two players; the SRO, behaving as the principal and the SRO member, as the agent. The agent chooses a level of fraud $x \in [0, \infty)$ and the SRO chooses a level of vigilance $y \in [0, \infty)$. The probability of fraud discovery is $p(x, y) \leq 1$ and it is assumed to be twice differentiable. Also, $p_x > 0$, $p_y > 0$, $p_x(0, y) = 0$, $p_y(x, 0) = \infty$, $p_{xx} > 0$, $p_{yy} < 0$ and $p_{xy} > 0$. This latter assumption implies that the marginal positive effect of fraud on p is increasing in the level of vigilance. The penalty faced by the agent if discovered is $T > 0$. Some of these assumptions deserve a comment. First, the assumption that evidence of fraud emerges with probability p may be an extreme one in some cases. Although in many cases fraud is either discovered or unnoticed in a binary way, in some cases there can be other non-binary signals the principal could use to assess the likelihood and extent of fraud, wrongdoing and negligence.[22] It is unclear, however, how and whether different modelling options would change the basic results of this work. Second, the penalty T can involve both penalties determined and imposed by the principal (with some possible limit) or other costs such as the transitory or

[21] See also the survey in Spulber (1989).

[22] For example, these include information on investment and trading patterns in financial services, or the number of patient checks per unit of time in health services.

permanent suspension of a licence to operate, or the costs of compensation to consumers, for example as in most professions. In the case of an SRO as a firm, these costs may include the costs of becoming unemployed, or not promoted, as well as the nonpayment of discretionary performance payments. Third, the level of fraud x can represent direct benefits due to fraud, such as extracting resources unlawfully and directly from consumers, or from inducing unnecessary treatment or repairs to consumers.[23] Alternatively, x can reflect the saving of costs if the provision of care and good quality is costly for the agent.

Finally, throughout this work it is assumed that the agent has no concerns for reputation. This assumption is likely to be reasonable for the cases where SROs are interpreted as firms, because an employee's identity is often not known to consumers. In other examples of SROs, however, this assumption may prove more limiting. In some professions, for example, practitioners usually have a reputation of their own, which affects their payoffs. However, the relevance and role of agent reputation may be limited by consumer's search costs and limited memory, as well as by limited consumer choice. For example in the NHS in the UK, patients play little role in choosing their GP, and in state-provided legal aid, defendants are usually assigned an attorney directly, and there is often very little they can do to change him/her.[24] In the conclusions of this work the possible limitation of the assumption of no agent reputation is aknowledged and discussed as potential future research.[25]

2.1 Agent's Behaviour

It is assumed that the SRO members utility is increasing in fraud. For simplicity, it is assumed that utility is linear in the amount of fraud. It is also assumed that the agent dislikes being punished after being found engaged in fraud. Then the agent maximises,

$$x - p(x,y)T \tag{1}$$

and there is an interior solution if the first-order condition (FOC) is satisfied; $0 = 1 - p_x T$ and $p_{xx} > 0$ is met. Therefore, the agent's reaction function will have negative slope given by $\frac{dx^*}{dy} = \frac{-p_{xy}}{p_{xx}} < 0$, reflecting that optimal fraud is decreasing in SRO vigilance. Let $x_{max} = x^*(y = 0)$. It is assumed that the agent knows the type of the SRO principal.

[23] See for example Emons (1997) for the incentives to provide fraudulent advice.

[24] Moreover, the need for Self-regulation is higher precisely in industries where agent reputation is likely to be limited or not very effective.

[25] In particular, I conjecture that agent's reputation may change the principals' vigilance and exposure incentives.

2.2 Consumers' Beliefs

As mentioned in the introduction, consumers are assumed to have uncertainty about the type of the SRO, in terms of its honesty and propensity towards quality monitoring within the SRO. Let the SRO have two "incarnations" or types denoted by i, $i = 1, 2$, and let c_i denote the unobservable unit costs of vigilance of the two SRO types, with $c_1 < c_2$. This assumption captures consumers' uncertainty about the "propensity" of the SRO towards quality vigilance. The credence-goods assumption implies that consumers can only observe fraud exposure, or absence of it (not x or y). Let λ be consumers' prior belief that the SRO is of the "vigilant" type. Let b_e and b_n be the *updated* beliefs with and without SRO exposure, respectively. In order for consumers to observe fraud exposure, it is required not only that fraud is *discovered* by the SRO, but also that fraud is *exposed* to the public by the SRO. Therefore, the decision of exposure is different from the vigilance decision. In fact, the SRO may be willing to be vigilant but it may not be willing to expose fraud, if discovered and vice versa. The performance of self-regulation must be assessed in these two respects.[26] Let $e_i = 0, 1$ be the SRO non-exposure and exposure decisions, respectively. Then, under Bayesian updating by consumers and with $e_i p_i$ denoting the equilibrium probabilities of exposure by each SRO type, then

$$\begin{aligned} b_e &= \frac{e_1 p_1 \lambda}{e_1 p_1 \lambda + e_2 p_2 (1-\lambda)} \\ b_n &= \frac{(1-e_1 p_1)\lambda}{(1-e_1 p_1)\lambda + (1-e_2 p_2)(1-\lambda)} \end{aligned} \tag{2}$$

Note that when $e_i = 0$, then b_e becomes indeterminate. As it is customary, it is assumed that in such cases if consumers are confronted with (out-of-equilibrium) exposure, they would compute b_e as if $e_i = 1$.

2.3 SRO Behaviour and Types

The SRO has two choices to make. First, it must decide its vigilance effort, and secondly it must decide whether to expose fraud, if discovered. Then the SRO of type i chooses e_i and y_i to maximise the expected value of reputation;

$$[e_i p(x, y_i) b_e + (1 - e_i p(x, y_i)) b_n] W - c_i y_i \tag{3}$$

where term $W = w_1 - w_2$ denotes the net difference of the reputation value of each SRO type. This maximisation is equivalent to maximising the difference between the reputation gain and the vigilance costs;

$$e_i p(\cdot) R - c_i y_i \tag{4}$$

[26] An alternative modeling option would be to assume that once fraud is discovered by the SRO, the evidence of fraud leaks out to the public with some probability even against the SRO's will. This possibility is explored in Nunez (2000 B). However, the implicit assumption of "no leaks" captures the idea that SROs are usually the only ones properly informed about product quality.

where $R = (b_e - b_n)(w_1 - w_2) \in [-W, W]$ denotes the "reputation gain" of the SRO, which corresponds to the value of the change in the probability of being perceived by consumers as being of the vigilant type. The term $W = w_1 - w_2$ is motivated in different ways. First, this is customary in many signalling models in which different types are simply assumed to imply different payoffs for the party observing the signals. For example, consumers may value more a "vigilant" SRO because they expect less fraud, more consumer protection and more likelihood of compensation for product failure. Second, since assumption $W > 0$ can only promote finding positive vigilance and voluntary exposure in equilibrium, proving non-existence of such an equilibrium even under these favourable conditions would constitute a significant "impossibility" result. Third, suppose consumers value (unobservable) "equilibrium fraud" such that $w_i(x^*(y_i^*))$. If an equilibrium satisfying $x^*(y_1^*) < x^*(y_2^*)$ exists for *any* $W > 0$, then there would be a qualitatively similar equilibrium if terms w_i were defined as decreasing functions of $x^*(y_i^*)$.

At this point it is necessary to specify whether the SRO types internalise the reaction of the agent in their maximisation problem. If they do, then they behave as Stackelberg players, and as Cournot players otherwise. Therefore, under each behavioural assumption, there is an interior optimum if the FOCs are satisfied. If $e_i = 1$ then the FOCs are;

Stackelberg SRO:

$$\frac{dp(x^*(y_i), y_i)}{dy_i} R - c_i = 0 \tag{5}$$

Cournot SRO:

$$\frac{\partial p(x, y_i)}{\partial y_i} R - c_i = 0 \tag{6}$$

and the second-order conditions $\frac{d^2p}{dy^2} < 0$ and $p_{yy} < 0$ are met in each case, respectively.

3 Equilibria

An equilibrium in this model consists of exposure choices e_i^*, vigilance levels y_i^* for each type of SRO, fraud levels $x^*(y_i^*)$ in response to each type of SRO, and beliefs b_e^*, b_n^* by consumers that satisfy Bayes' law. [27] In this section the existence of equilibria of the model presented above is analysed by proving the existence of a fixed point $R^* \in [-W, W]$.

[27] In what follows, the superscript * will be used to denote equilibrium levels of the corresponding variables.

From the SROs maximisation problem it is straightforward that a necessary condition for $e_i^* = 1$ is $R > 0$; voluntary fraud exposure would occur only if it yields a reputation gain. Otherwise, $e_i^* = 0$. Naturally, in the latter case because exposure does not happen, optimal vigilance is zero and consequently fraud is maximum. Therefore, this fact implies that only the case $R > 0$ remains to be examined in what follows.[28] The Stackelberg case is analysed first. From the FOCs the optimal level of vigilance by each SRO can be derived, which can be expressed as;

$$y_i^* = \sigma(c_i/R) \tag{7}$$

where $\sigma^{-1}(\cdot) = (\frac{dp}{dy})$. Then, for a given R_t the optimal values y_i^* and $x^*(y_i^*)$, and the equilibrium probabilities p_i^* and beliefs b_e^*, b_n^* can be obtained, from which R_{t+1} can be computed. Therefore, equilibrium can be proven to exist by finding a fixed point $R^* = R_{t+1} = \omega_s(R_t) = R_t \in [-W, W]$, where ω_s is a function that relates the levels of R with the SRO types and the agent behaving optimally in response to different "initial" levels of R.

In the Cournot case, optimal vigilance can be derived assuming that the agent is also optimising for all levels of y_i. This is given by;

$$y_i^* = \gamma(c_i/R) \tag{8}$$

where $\gamma^{-1}(\cdot) = \partial p/\partial y$. Again, an equivalent of the function ω_s, name it ω_c, can be obtained, which relates levels of R. As in the Stackelberg case, existence of equilibrium in this case can be established by finding a fixed point $R^* \in [-W, W]$ for the function ω_c. Let $x_{max} = x^*(y = 0)$. The next result characterises one equilibrium (a fixed point) for both the Cournot and Stackelberg SRO games.

Result 1. *Equilibrium exists.* $R^* = y_i^* = 0$, $e_i^* = 0$ *and* $x^*(y_i^* = 0) = x_{max}$ *is an equilibrium in both the Stackelberg and Cournot cases.*

Zero vigilance and maximum fraud is an equilibrium because if consumers have prior beliefs such that exposure is as likely to be carried out by either SRO type, that is $R_t = 0$, then such beliefs will be fulfilled in equilibrium. This will happen because both SRO types will indeed choose not to be vigilant and expose fraud at all, which leads to $b_e^* = b_n^*$, and consequently to $R^* = 0$. Therefore, this situation constitutes a fixed point. Also, note that $e_i^* = 0$ because exposure does not lead to a reputation gain. The next section analyses the existence of other equilibria in the Cournot and Stackelberg games. For reasons that will become apparent, this is done by analysing the cases where $\frac{dp}{dy}$ is non-positive and positive.

[28] As it will become clear later, the existence of the exposure choice helps rule out some implausible equilibria, namely $R^* < 0$ in the Stackelberg case.

3.1 Equilibrium with $\frac{dp}{dy} \leq 0$

Given the assumptions of this model, the probability of fraud discovery can be increasing or decreasing in SRO vigilance, once the effect of the agent's response to SRO vigilance is taken into account. This section examines the case where $\frac{dp}{dy} \leq 0$.

3.1.1 Stackelberg Case

Since this is only a particular case of the previous section, the equilibrium stated in Result 1 is still an equilibrium in this case. However, we are interested in the likelihood of separating equilibria with positive vigilance and fraud deterrence, which is analysed next. For $R_t > 0$, by simple inspection of the FOCs, neither SRO type will choose positive vigilance if $\frac{dp}{dy} \leq 0$. Then, no beliefs updating will exist, which implies that $R_{t+1} = 0$. Therefore, the only equilibrium is the one presented in Result 1, which involves zero vigilance and maximum fraud.[29]

3.1.2 Cournot Case

As in the previous case, the equilibrium in Result 1 holds. For the cases where $R_t > 0$, then $y_i^* > 0$, that is, it is optimal for the SRO types to set positive vigilance. However, this in turn would imply $R_{t+1} \leq 0$ if $\frac{dp}{dy} \leq 0$, as the next result shows, which follows from the expressions for b_n and b_e.

Lemma 1. *$R > 0$ if and only if $p_1 > p_2$.*

This result follows from consumers' Bayesian beliefs, and it implies that exposure would improve the SRO reputation if and only if exposure is a more likely event when the SRO is of the low vigilance cost type, rather than the high cost type.

Lemma 1 implies that for all $R_t > 0$ then it is necessarily the case that $R_{t+1} < 0$, because $y_1 > y_2$ and $\frac{dp}{dy} < 0$ jointly imply that $p_1 < p_2$. These facts lead to the following result.

Result 2. *If $\frac{dp}{dy} \leq 0$ then $R^* = y_i^* = 0$, $e_i^* = 0$ and $x^*(y_i^* = 0) = x_{max}$ constitute the unique equilibrium in both the Stackelberg and Cournot cases.*

This result proves the impossibility of an equilibrium with any positive vigilance, any voluntary fraud exposure and any level of fraud deterrence, however small, for any SRO type if an additional unit of vigilance does not increase the probability of fraud discovery (and punishment). Moreover, this result exists regardless of whether the SRO behaves as a Cournot or Stackelberg SRO. The

[29] The role of the exposure decision $e_i^* = 0$ is central to this result. In fact, if the SRO could not cover up fraud, then an equilibrium with $R < 0$ is plausible, in which the Stackelberg SRO may have incentives to increase vigilance to reduce the chances of fraud discovery and avoid a reputation loss. This idea is, however, discussed later in the context of parallel regulation, and in Nunez (2000 A,B).

intuition behind this result is the following. In the Stackelberg case, the SRO chooses zero vigilance because vigilance *decreases* the probability of fraud exposure, which must happen necessarily in order to achieve a reputation gain. In the Cournot case instead, the SRO chooses zero vigilance because they know that consumers believe that fraud exposure is a more likely event the lower is SRO vigilance, and therefore exposure would be interpreted by consumers as a signal of low vigilance rather than high vigilance.

It is important to verify whether an equilibrium with positive vigilance exists if $\frac{dp}{dy} > 0$, which is analysed next.

3.2 Equilibrium with $\frac{dp}{dy} > 0$

Result 3. *If $\frac{dp}{dy} > 0$ there must be an equilibrium in the Cournot case that satisfies $R^* > 0$, $e_1^* = e_2^* = 1$, $y_1^* > y_2^* > 0$ and $x^*(y_1^*) < x^*(y_2^*) < x_{max}$. In the Stackelberg case this equilibrium is possible but cannot be guaranteed.*

Proofs for the Stackelberg and Cournot cases are provided in Appendix 1. Note, however, that this equilibrium is not unique because Result 1 shows that the situation $R^* = y_i^* = 0$, $e_i^* = 0$ and $x^*(y_i) = x_{max}$ is also an equilibrium in the Cournot and Stackelberg cases. This situation is portrayed in Figure 1. E_1 is the equilibrium stated in Result 1 where there is no exposure, no vigilance and fraud is maximum. E_2 is the equilibrium involving positive vigilance and some fraud deterrence, which can only exist if $\frac{dp}{dy} > 0$. Finally, Figure 1 also portrays a situation where a positive vigilance stable equilibrium exists in the Stackelberg case, and another where such equilibrium fails to exist.

3.3 Likelihood of a Positive Vigilance Equilibrium

Previously it has been stated that in the model above a necessary condition for an equilibrium with positive vigilance and some fraud deterrence is that $\frac{dp}{dy} > 0$ is satisfied. The likelihood of this case is explored next.

The reaction function of the agent slopes downwards and its slope is given by $\frac{dx^*}{dy} = -\frac{p_{xy}}{p_{xx}} < 0$. As already shown, an equilibrium with $y_i > 0$, can exist only if $\frac{dp}{dy} > 0$. However, it is unclear that this will be the case. Even though the asymmetry in vigilance costs of the SROs implies that $y_1^* > y_2^*$, it will also be the case that $x^*(y_1^*) \leq x^*(y_2^*)$, and therefore the net effect on $p(\cdot)$ remains ambiguous. Since any equilibria requires that the agent is maximising, an interior equilibrium with positive vigilance may exist only if $p(\cdot)$ is positively related to vigilance as one moves along the agent's reaction function. To analyse this, differentiate as follows;

$$\frac{dp(x^*(y), y)}{dy} = p_x \frac{dx^*(y)}{dy} + p_y = \frac{-p_x p_{xy}}{p_{xx}} + p_y \tag{9}$$

This has the same sign as $\frac{\partial[p_x/p_y]}{\partial x}$. Therefore, the possibility of an equilibrium with fraud deterrence can be reduced to the sign of expression (9), which is stated in the next lemma.

Lemma 2. *If* $\frac{dp}{dy} > 0$, *then it cannot be the case that* $\frac{\partial[p_x/p_y]}{\partial x} < 0 \; \forall \; x, y$.

However, it is entirely possible for $\partial[p_x/p_y]/\partial x$ to be negative, in which case there would be an equilibrium with no vigilance and maximum fraud.[30] Therefore it is interesting to analyse the likelihood of this condition being satisfied by alternative functional forms of $p(\cdot)$ that satisfy the initial assumptions. The first appealing possibilities to explore are separable forms. Additive separable forms are, however, ruled out because the assumption that $p_{xy} > 0$ is violated. Therefore, only multiplicative separable forms can be considered, that is, $p(x, y) = f(x)g(y)$ that satisfy the initial assumptions. It can be shown that,

$$sgn[\frac{dp(x(y), y)}{dy}] = sgn[\frac{d(f_x(x)/f(x))}{dx}] \tag{10}$$

This implies that a *necessary* condition for an equilibrium with positive vigilance and fraud deterrence when $p(x, y)$ is multiplicatively separable is that $f(x)$ must be strictly log convex, that is, it must increase in x more rapidly than exponentially since its rate of growth must be *increasing* in x. [31]

[30] This condition would imply that in the (x, y) plane, the slope of the iso-probabilities for a fixed value of y are increasing in x (they become less negative or flatter).

[31] The intuition of this condition can be illustrated using a heuristic interpretation. Suppose $p(\cdot)$ increases by some amount due to increased vigilance. Then fraud would "initially" fall. This would then decrease $p(\cdot)$ somewhat, which would then increase fraud and so forth. The more convex $f(x)$ is, then the larger the positive feedback effects of fraud on $p(\cdot)$ will be relative

This condition implies that an ample range of possible functional forms of $f(x)$ will fail to sustain positive vigilance and fraud deterrence in equilibrium.[32] Yet, this condition is *not* a sufficient condition for positive vigilance. It is natural to assume that finding and exposing fraud requires that *some* degree of fraud must indeed exist. In other words, it is assumed that fraud is potentially verifiable by a third party, and therefore that no "type I" errors can occur in equilibrium. Therefore, either the SRO finds fraud only if some fraud actually exists, or if an innocent SRO member is wrongly and publicly accused of fraud, then such accusation can be proven to be false and therefore consumers behave as if exposure did not occur. This assumption is presented next.

Assumption 1: $p(0, y) = 0$.

This assumption implies that $f(0) = 0$; if no fraud exists, then the chances of (rightful) fraud discovery and exposure are nil. This assumption imposes a new constraint and further reduces the chances of an equilibrium with positive vigilance. For example, functions of the general form $f(x) = ab^{cx}+d,\ a, b, c > 0$, $d > 0$, and $f(x) = ae^{h(x)},\ \ h(0) = 0,\ \ h_x > 0$ do indeed satisfy $\frac{dp}{dy} > 0$, but fail to satisfy $p(0, y) = 0$. The intuition behind this condition is that since the probability of fraud discovery must be positive at x_{max} but zero at $x = 0$, then $p(\cdot)$ must necessarily increase in x between these two values along the agent's reaction function (although not necessarily monotonically). In other words, this loosely means that $p(\cdot)$ must decrease in vigilance level y in overall terms as y increases along the agent's reaction function. Therefore, the possibility of an equilibrium with positive vigilance requires $\frac{dp}{dy}$ to be positive, which can only happen in a local, restricted neighbourhood along the agent's reaction function. Therefore, this assumption further reduces the likelihood of an equilibrium with positive vigilance and fraud deterrence, in addition to the condition of log convexity of $f(x)$.

The general message of this section is that, for a very ample range of circumstances, self-regulation will not generate *any* incentives to expose fraud or to be vigilant, and fraud would be maximum as a result. Finally, even if $\frac{dp}{dy} > 0$ *and* assumption $p(0, y) = 0$ are satisfied, a situation with zero vigilance and full fraud is still a possible equilibrium as Result 1 has shown.

4 Introducing Parallel Regulation

It is commonly sustained that self-regulation and public enforcement of quality are mutually exclusive alternatives to deal with the negative effects of asymmet-

---

to the negative feedback effects, and therefore their combined effect would counterbalance less the initial increase in $p(\cdot)$. If $f(x)$ is log-convex, then the combined negative feedback effect of fraud would not outweigh the initial increase in $p(\cdot)$, making $p(\cdot)$ increase on the whole.

[32] For example, all functions of the form $f(x) = ax^b$, $\forall a, b > 0$, as well as all exponential functions of the general form $ab^x, \forall a, b > 0$.

ric information. This idea, however, rests more on a generalised belief than in actual theoretical foundation.[33] In fact, this need not be the case and both regulatory schemes can co-exist in parallel and complement each other. This section examines the effects of public quality enforcement in parallel to self-regulation. In particular, it analyses whether parallel regulation enhances or decreases the SRO incentives to be vigilant and exposure fraud. This is important because parallel regulation may increase the effectiveness of self-regulation, while seizing the claimed informational advantage of the SRO over public regulation.

The presence of public regulation is captured by introducing a function $q(z, x) \geq 0$ that represents the probability of fraud being discovered by the public regulator, with z being public vigilance, and q_z, q_x and $q_{xz} > 0$. Also $q(0, x) = 0$; public exposure can occur only if there is positive public vigilance ($z > 0$). The total probability of fraud discovery now becomes $p + q - pq$, where pq is the probability that both parties discover fraud jointly.

4.1 Consumers' Beliefs

Introducing parallel regulation raises the issue of how many exposure events should consumers be allowed to use to make their inferences about the SRO type. The possible events are four, namely, exclusive exposure by either the SRO or the regulator, simultaneous exposure by both parties, and no exposure by any party. Allowing the possibility of joint exposure as an independent event is a rather unappealing formulation, and it also makes matters more complex. Therefore, the probability of joint or simultaneous exposure must be allocated to either fraud discovery by the SRO, or by the regulator. It is plausible that the SRO has the informational edge over the public regulator. Accordingly, it is assumed that the public regulator can expose fraud *only* if the SRO has not found and exposed fraud yet. This is natural if the SRO happens to be more quickly and better informed about fraud than the regulator. A justification for this assumption is that the SRO's informational advantage is one of the arguments in favour of self-regulation. Therefore, it is interesting to examine the behaviour of SROs precisely under this assumption. Then, the probabilities of the three events for each type are $e_i p_i$, $(1 - e_i p_i)q_i$ and $(1 - q_i)(1 - e_i p_i)$ for SRO and regulator exposure, and no exposure, respectively. Naturally, these events add up to unity. Therefore, consumers' updated beliefs are;

$$\begin{array}{rcl} b_s & = & \phi(\frac{e_1 p_1}{e_2 p_2}) \\ b_r & = & \phi(\frac{q_1(1-e_1 p_1)}{q_2(1-e_2 p_2)}) \\ b_n & = & \phi(\frac{(1-q_1)(1-e_1 p_1)}{(1-q_2)(1-e_2 p_2)}) \end{array} \tag{11}$$

where $\phi(\rho) = \frac{\lambda \rho}{\lambda \rho + (1-\lambda)}$ and where b_s, b_r and b_n stand for the updated beliefs under SRO and public exposure, and no exposure, respectively. It is simple to

[33] See for example Spulber (1989).

show that if $x(y_1) = x(y_2)$, then $b_r = b_n$ as $q_1 = q_2$. However, if $x(y_1) < x(y_2)$, then $q_1/q_2 < 0$ and therefore $b_r < b_n$, which implies that the SRO experiences a reputation loss if fraud is exposed by the regulator, because exposure is more likely to happen for higher fraud levels.

4.2 Agent's Behaviour

Now the agent chooses x to maximise;

$$x - [p(x, y) + (1 - p(x, y))q(z, x)]T \tag{12}$$

Note that the agent must face penalty T independently of who discovers fraud. Therefore, the FOC and SOC become;

$$\begin{gathered} 1 - [p_x(1 - q) + (1 - p)q_x]T = 0 \\ -[p_{xx}(1 - q) - 2p_xq_x + (1 - p)q_{xx}] < 0 \end{gathered} \tag{13}$$

By differentiating the FOC totally, the slope of the agent's reaction function with a public regulator now equals

$$\frac{dx^*}{dy} = -\frac{(1 - q)p_{xy} - p_yq_x}{p_{xx}(1 - q) - 2p_xq_x + (1 - p)q_{xx}} \tag{14}$$

The denominator in (14) is positive because it is the negative of the SOC. Therefore, the slope of the reaction function has the same sign as $-(1-q)p_{xy} + p_yq_x$. This implies some facts that defy intuition. In fact, it is possible for this term to be *positive*, which means that the agent would choose to *increase* fraud if SRO vigilance increases. Moreover, such possibility would also imply that if $y_1 > y_2$, then $x^*(y_1) > x^*(y_2)$ and therefore that $b_r > b_n$, which means that if fraud is discovered by the regulator, then consumers should infer that it is more likely that it was the *vigilant* SRO type. Again, this fact appears counter-intuitive. Therefore, to avoid this the following assumption is made;

Assumption 2: $\frac{dx^*}{dy} < 0$ and therefore $(1 - q)p_{xy} > p_yq_x$.[34]

Finally, a further assumption is needed. Under parallel regulation the total probability of fraud discovery amounts to $p + q(1 - p) > p$. This fact suggests that parallel regulation will hurt the agent who will possibly react to it by reducing fraud. Since zero vigilance was shown to be a very likely case without parallel regulation, we are interested in optimal fraud when $y = 0$ to obtain results comparable to the case without parallel regulation. This would allow the examination of the role of parallel regulation on fraud, in addition to its effects over SRO vigilance. Comparing the agent's FOCs with and without parallel regulation, it follows that a necessary condition for $x^*_{rmax} = x^*_r(y = 0) < x_{max}$

[34] This assumption is analogous to the assumption that $p_{xy} > 0$ in the case without public regulation; both imply that optimal fraud is decreasing in SRO vigilance.

is that $q_x(1-p) > qp_x$.[35] where r denotes the case with parallel regulation. This condition implies that the agent's reaction function lies to the left of the one without a parallel regulator. Consequently, for a given y fraud must be lower with parallel regulation. This rather likely condition is justified on two grounds. First, since for a given amount of fraud the probability of discovery is necessarily higher with parallel regulation, and given the assumption that $p(0,y) = q(z,0) = 0$, (i.e. fraud can only be discovered if it exists), then it must be the case that total probability increases faster in x under parallel regulation in at least some range. Second, q and T are variables that can, in theory, be chosen by policy-makers and therefore both z and T could be adjusted to satisfy the condition above.[36] Therefore, the following assumption is made, which leads to Lemma 3.

Assumption 3: $q_x(1-p) > qp_x$.

Lemma 3. *For $y = 0$ if $q_x(1-p) > qp_x$, then $x^*_{rmax} < x_{max}$.*

This implies that even if SRO vigilance is zero, optimal fraud will be smaller under parallel regulation because the agent fears public regulation and will reduce fraud to avoid being punished. It is important to verify whether public regulation changes the value of $\frac{dp}{dy}$. Recall that $\frac{dp(\cdot)}{dy} = p_x(\frac{dx^*}{dy}) + p_y$. Since $\frac{dx^*}{dy}$ can turn larger or smaller with public regulation than without it depending on the value of the relevant parameters, then $\frac{dp}{dy}$ can change sign in either direction. In the next section equilibria are characterised for the case where $\frac{dp}{dy} \leq 0$. The justification for this is that in the previous section it has been shown that full fraud and no vigilance prevailed when $\frac{dp}{dy} \leq 0$, and that some vigilance and fraud deterrence would exist otherwise. In order to assess the role of parallel regulation, this section assumes the most hostile environment for parallel regulation, which is the case where $\frac{dp}{dy} \leq 0$.

4.3 Equilibrium with $\frac{dp}{dy} \leq 0$

4.3.1 Cournot Case

The Cournot SRO types now choose e_i and y_i to maximise;

$$[e_i p(x, y_i) b_s + (1 - e_i p(x, y_i))(q_i b_r + (1 - q_i) b_n)]W - c_i y_i \tag{15}$$

where $(q_i b_r + (1 - q_i) b_n)$ is the expectation of SRO type i of being perceived by consumers as being of the vigilant type, given no SRO exposure. The maximi-

[35] To see this, fix y. In both cases the FOCs must be satisfied and therefore $p_x(1-q) + q_x(1-p) = p_x = 1/T$, Then, because $p_{xx} > 0$, then $x^*_r < x^*$ requires that $q_x(1-p) > qp_x$

[36] For example by setting a sufficiently high penalty, or increasing z to increase q_x.

sation problem is equivalent to maximising[37]

$$\begin{aligned}&[e_i p(x, y_i)(b_s - (q_i b_r + (1 - q_i)b_n) + q_i(b_r - b_n)]W - c_i y_i \\ &\quad = e_i p(x, y_i)R^i + q_i(b_r - b_n)W - c_i y_i\end{aligned} \tag{16}$$

where $R^i = [b_s - (q_i b_r + (1 - q_i)b_n)]W$. Note that now the value of reputation change would be different for each SRO type. The choice of e_i is done just by comparing the expression (16) for each case, which implies that $e_i^* = 1$ if $R^i > 0$, and $e_i^* = 0$ otherwise. This is entirely natural; the SRO will decide to expose fraud only if this leads to a reputation gain (in expected value). If $e_i = 1$, then the FOC becomes simply,

$$p_y R^i - c_i = 0 \tag{17}$$

and the SOC is $p_{yy} < 0$.

There are two interesting issues. The first is whether the equilibrium $R = y_i^* = 0$, $e_i^* = 0$ and $x^*(y_i) = x_{max}$ can be ruled-out when parallel regulation is introduced. Second, even if this equilibrium does not vanish, it is interesting to analyse whether parallel regulation can create an equilibrium with positive vigilance and fraud deterrence.

Consider $R_t^i = 0$. Then again $y_i^* = 0$ and therefore $x^*(y_1^* = 0) = x^*(y_2^* = 0)$, which implies that $q_1 = q_2$ and $p_1 = p_2$. As a result, then $b_s = b_n = b_r$ and therefore $R_{t+1}^i(R_t^i = 0) = 0$, which constitutes a fixed point. This means that the equilibrium with no exposure, no vigilance and full fraud is still an equilibrium when parallel regulation is introduced. However, equilibrium fraud will now be lower than x_{max} because of the direct effect of public vigilance on the agent, which follows from Lemma 3.

The problem now is to establish the sign of R_{t+1}^i for any $R_t^i > 0$. To investigate this, let $b_n^0 = \phi((1-p_1)/(1-p_2))$. This is equal to the ex-post belief, given no SRO exposure, that the SRO is of the vigilant type when either no parallel regulation exists or when $q_1 = q_2$. Note that if $p_1 < p_2$ then $b_n^0 > b_s$. Therefore, if it can be proven that $b_n^0 \leq q_i b_r + (1 - q_i)b_n$, then it would be proven that $R_{t+1}^i(R_t^i > 0) < 0$. This would finally imply that there would not be a fixed point with fraud exposure, positive vigilance and fraud deterrence.

Lemma 4. *If $p_1 < p_2$ and $q_1 < q_2$ then $q_2 b_r + (1-q_2)b_n < b_n^0 < q_1 b_r + (1-q_1)b_n$.*

The proof is presented in Appendix 1 and it is illustrated by numerical examples. The situation is portrayed in Figure 2.

[37] In this transformation, term $b_n W$ is lost, which is irrelevant for the maximisation, however.

Figure 2: Function $\phi(z)$

Lemma 4 implies that $R^1_{t+1} < 0$ for all R^1_t. In other words, the vigilant SRO type should always expect a reputation loss if its vigilance is higher that the lax SRO type. Under these circumstances, $e^*_1 = 0$ and $y^*_1 = 0$ for all $R_t > 0$. The intuition behind this result is the following. It is a fact that SRO exposure hurts the SRO because $b_s < \lambda$. However, no exposure may also hurt the SRO because it opens the way for regulator exposure, which may imply a reputation loss because $b_r < \lambda$ is a possible outcome.[38] Therefore the dilemma for the SRO is whether to expose and get a reputation loss with certainty, or not expose and get a lottery over two outcomes, namely regulator exposure and no regulator exposure. It is the case that the vigilant SRO type has lower fraud, and consequently it has a lower probability of fraud being discovered among its members by the public regulator. This implies that the chances of regulator exposure (given no SRO exposure to date) are low enough to make the vigilant SRO run the risk of having its member exposed by the regulator (with probability q_1). In other words, exposure will be a dominated strategy for the vigilant SRO if $\frac{dp}{dy} < 0$, and therefore $e^*_1 = 0$. To find an equilibrium, the values for R^i_{t+1} must be computed for $R^i_t > 0$. From consumers' inferences it follows that since $e^*_1 = 0$ and $y^*_1 = 0$, then $b_s = 0$, but b_r,$b_n > 0$. This implies that the only sustainable situation involves zero reputation change for both SRO types. Therefore, it cannot be the case that there is an equilibrium with positive reputation, positive vigilance in the Cournot case with $\frac{dp}{dy} < 0$. However, from Lemma 3, equilibrium fraud will be lower with parallel regulation due to its direct effect on the agent's behaviour.

Summing up, in the Cournot game parallel regulation cannot destroy the equilibrium with nil vigilance, and it fails to create an equilibrium with positive

[38]This will depend on the values for p_i and q_i.

vigilance and fraud deterrence. However, it can reduce the level of fraud due to its direct effect on the agent's behaviour.

4.3.2 Stackelberg Case

The Stackelberg SRO internalises the optimal reaction of the agent to different vigilance levels. Since the agent's fraud choice will also affect the probability of regulator exposure, the effect on $q(\cdot)$ will also be internalised by the SRO, unlike the Cournot case. Then, the Stackelberg SRO types choose e_i and y_i to maximise,

$$[e_i p(x, y_i)(b_s - (q_i b_r + (1 - q_i)b_n) + q_i(b_r - b_n)]W - c_i y_i \tag{18}$$

As in the Cournot case, the choice of e_i will depend only on R^i; $e_i^* = 1$ if $R^i > 0$ and $e_i^* = 0$ otherwise. If $e_i^* = 1$ the FOC is;

$$\frac{dp}{dy}[b_s - (q_i b_r + b_n(1 - q_i))] + \frac{\partial q_i(x^*(y))}{\partial y}(1 - p)(b_r - b_n) = c_i/W \tag{19}$$

If $e_i = 0$, then the FOC becomes only

$$\frac{\partial q_i(x^*(y))}{\partial y}(b_r - b_n) = c_i/W \tag{20}$$

As already explained above, Lemma 4 implies that $e_1^* = e_2^* = 0$ and therefore the second FOC is the relevant one. It can be shown that the SRO may have enhanced incentives to be vigilant. The left hand side of (20) is positive if $b_r - b_n < 0$, and it represents the value to the SRO of marginal units of SRO vigilance that help avoid the reputation loss that would exist if fraud is discovered by the regulator.

Redefine $R = (b_r - b_n)W$. Note that R is now the same for both types, unlike the Cournot case. Also note that because $e_i^* = 0$, R only depends on q_i and not on p_i. An equilibrium can be proven to exist by finding a fixed point for $R = (b_r - b_n)W$. For $R_t \geq 0$, then clearly $y_i^* = 0$ and $q_1 = q_2$. This implies that $R_{t+1} = 0$ for all $R_t \geq 0$. Consequently, $R^* = 0$ constitutes a fixed point. Therefore zero vigilance and full fraud still remains an equilibrium in the Stackelberg case. However, for $R_t < 0$, then there may be an equilibrium with positive vigilance and fraud deterrence. In fact, when R_t is marginally negative, then $y_1^* = y_2^* = 0$ because the vigilance costs are positive. As $R_t <$ decreases, then eventually the vigilant type will find worthwhile to begin choosing positive vigilance. At this point, however, the lax type will still stick to no vigilance. This implies that R_{t+1} becomes negative. As R_t becomes more negative, then it must be the case that the vigilant type is more vigilant than the lax type, and therefore $x^*(y_1) < x^*(y_2)$. This implies that R_{t+1} will remain negative as R_t decreases. A fixed point is possible, but it cannot be guaranteed if each type's vigilance cost are both positive. It is perfectly possible, however, that the function R_{t+1} crosses the 45-degree line, and given its continuity, then a fixed point would exist. This only requires R_{t+1} to be sufficiently decreasing as R_t falls to reach and intercept the 45-degree line. This would happen, for example, if $q(x)$ is sufficiently responsive to reductions in fraud. A fixed point would also exist provided that the cost asymmetry is sufficiently large.[39] The main conclusions of this section are summarised in the next result.

Result 4. *In the Stackelberg case with parallel regulation and* $\frac{dp}{dy} \leq 0$ *the equilibrium with no exposure, no vigilance and maximum fraud still exists. However, an equilibrium involving positive vigilance and fraud deterrence is possible.*

This situation is portrayed in Figure 3.

Figure 3: Equilibria in the Stackelberg case with parallel regulation

Therefore, the conclusion of this section is that even if the SRO is not willing to expose fraud, there is still the possibility of an equilibrium with positive vigilance and fraud deterrence.

5 Discussion and Conclusions

This work analyses the incentives of an SRO to devote resources to find and expose fraud and wrongdoing by its members, when there is uncertainty about its predisposition towards vigilance. A distinctive feature of this work is that product quality (inversely related to fraud and wrongdoing) is endogenous, and therefore the signals from which customers update their beliefs about the SRO

[39] Note that this equilibrium was possible because it was assumed that $e_i = 1$, which simplified the maximisation problem. This situation is true in equilibrium because $p_1 < p_2$ and $q_1 < q_2$, which are the necessary conditions for $e_1^* = 0$.

type are determined endogenously as the result of strategic interaction within the SRO. This work shows that under very ample conditions unregulated SROs will have no incentives either to search for fraud or to expose it, if it is discovered. This is always a possible equilibrium, and in a wide class of cases it is *unique*. This result held for both Cournot and Stackelberg SRO types.

This result suggests that reputation based arguments in favour of self-regulation may be weak in industries where voluntary disclosure is one of the main sources of information to consumers. A corollary of this is that there may be large scope for public regulatory policy in such industries. This was addressed in the latter sections of this work. Parallel regulation did generate new equilibria in which fraud is lower and SRO vigilance is positive. The forces at work are two. First, there is the effect of reduced fraud due to the direct deterrent effect of public vigilance on the agent. Second, the SRO now has incentives to increase its vigilance to avoid exposure by the public regulator. These results suggest that self-regulation and public regulation should not be seen as exclusive alternatives. Moreover, a cost-benefit evaluation of public regulation that does not take into account its effects on SRO reputation and behaviour will underestimate the net social benefits of public regulation.

The results about the role of public regulation were obtained under the assumption that the SRO could choose whether to expose or cover up the evidence of fraud. It was shown that under a broad range of circumstances the SRO would choose not to expose fraud. As explained above, even under such circumstances the Stackelberg SRO had incentives to be vigilant. However, this result has the important (and rather unfortunate) implication that the SRO will not expose fraud in equilibrium. In turn, this raises several issues. First, it suggest that consumers will only be able to get information from public exposure, and therefore the problem of asymmetric information may remain largely unsolved. Second, parallel regulation will not make any difference if the SRO behaves à la Cournot, apart from its direct effect on the agent's behaviour, since SRO vigilance would still be nil. Third, if voluntary exposure did happen in equilibrium, then it is plausible to expect that optimal vigilance would be higher given that fraud discovery is a necessary condition for fraud exposure. Fourth, if the agent also has concerns for reputation, then SRO exposure may further reduce fraud if exposure now hurts the agent, apart from the direct penalties faced by her. Therefore, an issue of most importance for future research is to investigate whether the SRO may have incentives to expose fraud voluntarily under different circumstances than the ones considered here. Proving so would imply that consumers will have more access to information (which alleviates the problem of asymmetric information), parallel regulation would improve matters even under Cournot behavioural assumptions, and also the SRO may have enhanced incentives to expand vigilance than it has been suggested in this work.

There are at least three environments in which voluntary exposure may emerge. First, it was shown earlier that parallel regulation may make the agent's optimal fraud less sensitive to SRO vigilance. This raises the possibility that parallel regulation may turn the sign of $\frac{dp}{dy}$ from negative in the unregulated case, to positive in the case with parallel regulation. By doing so, now the SRO may have incentives to expose fraud voluntarily because exposure now may be informative of a relative advantage in vigilance costs. Second, if the SRO is risk-averse, then exposure may have an additional advantage over non-exposure because SRO exposure has a unique outcome whereas the latter involves a lottery of two possible outcomes, namely public exposure and non-exposure. Risk aversion would imply that exposure would have a premium over cover-up, which may reverse the choice of non-exposure obtained in this work. Third, if it is assumed that the agent also has a reputation to protect, then the SRO may attempt to threaten the agent to expose fraud with the purpose of disciplining her.[40]

Therefore, these arguments suggest that perhaps the results obtained here may provide only a bottom-line for the potential benefits of parallel regulation and the complementarities it may have with self-regulation. If some of these ideas prove to be correct, then the underestimation of the social benefits of public regulation will be larger.

Finally, there are multiple other avenues for future research. For example, SRO reputation-based incentives can be studied assuming that consumers have uncertainty about the nature of the SRO *agents* instead. Second, there can be side transfers and "corruption" between the SRO principal and agents, which can affect SRO vigilance and exposure incentives. This can become even more interesting in the case where the agent also has a reputation to protect discussed above. Third Fourth, there can be reputation-based competition among many SROs in an industry, or competition for promotion among many principals within an SRO. These extensions can lead to rather different results than those reported here.[41]

Appendix 1

Proof Result 3. Cournot case. Function ω_c is continuous in interval $(0, W]$. In the neighbourhood of $R_t = 0$, $\underset{R_t \to 0}{Lim}\ R_{t+1} > 0$. Indeed, this has the same sign as $\underset{R_t \to 0}{Lim}\ [b_s - b_n]$. As R_t gets smaller, still $y_1 > y_2$ because $c_1 < c_2$, and because $p_y R \to \infty$ as $y_i \to 0$. Therefore, $p_1 > p_2$ for any $R_t > 0$. Consequently, $\underset{R_t \to 0}{Lim}\ p_1/p_2 > 1$ and $\underset{R_t \to 0}{Lim}\ (1 - p_1)/(1 - p_2) = 1$, which implies that $\underset{R_t \to 0}{Lim}\ R_{t+1} > 0$. On the other extreme, when $R_t = W$ then $y_1 > y_2 > 0$, and therefore $b_e < 1$ and $b_n > 0$. This implies that $R_{t+1} = \omega_c(R_t = W) < W$. Given continuity of ω_c in (0,W], a fixed point $R^* > 0$ must exist, in which $y_i^* > 0$, $e_i^* = 1$ and $x^*(y_i^*) < x_{max}$.

[40] There is, however, the issue of whether such a threat would be credible in a dynamic game; since fraud discovery precedes exposure, SRO exposure may not satisfy sub-game perfection.

[41] Nunez (2000 A, B) presents a discussion on some of these issues.

Stackelberg Case. For $R_t > 0$ but close to 0, $e_i^* = 1$ although $y_i^* = 0$ because $c_i > 0$. This yields $R_{t+1} = 0$ as $R_t \to 0$. As $R_t > 0$ increases, there is a point from which $y_1^* > y_2^* \geq 0$, and therefore $R_{t+1} > 0$. However, although a fixed point $R^* > 0$ is possible, it cannot be guaranteed, which would ultimately depend on c_i and $\frac{dp}{dy}$.

Proof Lemma 4

1. Prove that $q_2 b_r + (1-q_2) b_n < b_n^0$**;** The situation is portrayed in Figure 2 in the text. Recall that $b_n^0 = \phi(k)$, where $k = \frac{1-p_1}{1-p2}$. Then, $q_2 \frac{q_1}{q_2} k + (1-q_2) \frac{1-q_1}{1-q_2} k = k$ is the value in the x-axis that corresponds to $q_2 b_r + (1-q_2) b_n$. Since $\phi(\cdot)$ is concave, then it must be the case that $b_n^0 = \phi(z = k) > q_2 b_r + (1-q_2) b_n$.

2. Prove that $q_1 b_r + (1-q_1) b_n > b_n^0$**;** After dividing by λ and k this expression reduces to;

$$\frac{q_1^2}{q_1 f + q_2 g} + \frac{(1-q_1)^2}{(1-q_1) f + (1-q_2) g} > 1/(f+g) \tag{21}$$

where $f = \lambda k > 0$ and $g = (1-\lambda) > 0$. After expanding and cancelling terms out, (21) simplifies to;

$$-\frac{g^2 (q_1 - q_2)^2}{(q_1 f + q_2 g)[f(q_1 - 1) + g(q_2 - 1)](f+g)} > 0 \tag{22}$$

Therefore, provided that $q_1,\ q_2 \in [0,1)$, the statement in (22) is true and therefore $q_1 b_r + (1-q_1) b_n > b_n^0$. Note that $q_1 b_r + (1-q_1) b_n = b_n^0$ only if $q_1 = q_2$, and it will be positive otherwise. The next table indicates values for $q_1 b_r + (1-q_1) b_n$ for arbitrary parameters values.

$q_1 \downarrow$	$q_2 \rightarrow$	0.1	0.3	0.5	0.7	0.9	0.99
0.01		0.740	0.783	0.833	0.889	0.954	0.990
0.2		0.731	0.730	0.751	0.788	0.839	0.870
0.3		0.739	0.727	0.737	0.763	0.806	0.833
0.4		0.750	0.730	0.729	0.747	0.782	0.806
0.5		0.764	0.736	0.727	0.736	0.764	0.786
0.6		0.782	0.747	0.729	0.730	0.750	0.769
0.7		0.806	0.763	0.737	0.727	0.739	0.756
0.8		0.839	0.788	0.751	0.730	0.731	0.745
0.9		0.887	0.828	0.780	0.744	0.727	0.735

Table 1: Values for $q_1 b_r + (1-q_1)b_n$ for parameters $\lambda = 0.5$, $p_1 = 0.2$, $p_2 = 0.7$ and $b_n^0 = 0.727$

References

Akerlof, G. (1970) “The Market for Lemons: Qualitative Uncertainty and the Market Mechanism”, *Quarterly Journal of Economics*, 84, August, pp. 488-500.

Bagwell, K. and M. Riordan (1991) “High and Declining Prices Signal Quality”, *American Economic Review* 81, pp. 224-239.

Besanko D. and D. Spulber (1989) ”Delegated Law Enforcement and Noncooperative Behaviour”, *Journal of Law, Economics and Organization* Vol. 5, No. 1, Spring, pp. 25-52.

Cooper, R. and T. Ross (1985) “Monopoly Provision of Product Quality with Uninformed Buyers”, *International Journal of Industrial Organization*, 3(4), pp. 439-449.

Emons,W. (1997) “Credence Goods and Fraudulent Experts”, *Journal of Economics*, Vol. 28, No. 1, Spring, pp. 107-119.

Financial Services Authority (1998) “FSA today”, FSA web site.

Fudenberg, D. and J. Tirole (1991) *Game Theory*, Cambridge, The MIT Press.

Gehrig, T. and P-J. Jost (1995) “Quacks, Lemons, and Self Regulation: A Welfare Analysis”, *Journal of Regulatory Economics*, 7, pp. 309-325.

Green, E. and R. Porter (1984) “Noncooperative Collusion Under Imperfect Price Information”. *Econometrica*, 52: 87-100.

Hobbes, T., *Leviathan*, London: The Penguin Classics, 1985.

Home Office (1997) *Criminal Statistics of England and Wales, 1996.* London, HMSO.

Kay, J. and J. Vickers (1988) ”Regulatory Reform in Britain”, *Economic Policy* 7, pp. 285- 351.

Laffont, J.J. and J. Tirole (1993) *A Theory of Incentives in Procurement and Regulation*, MIT Press, Cambridge, Mass.

Leland, H.E. (1979) "Quacks, Lemons, and Licensing: A Theory of Minimum Quality Standards", *Journal of Political Economy*, 87, pp. 1328-1346.

Levi, M. (1987) *Regulating Fraud. White Collar Crime and the Criminal Process*, Tavistock Publications, New York.

Milgrom P. (1981) "Good News and Bad News: Representation Theorems and Applications", *Bell Journal of Economics* 12, Autumn, pp. 380-391.

Milgrom P. and J. Roberts (1986) "Price and Advertising Signals of Product Quality", *Journal of Political Economy*, 94, August, pp. 796-821.

Nelson, P. (1970) "Information and Consumer Behaviour", *Journal of Political Economy*, 78, March/April, pp. 311-329.

Nelson, P. (1974) "Advertising as Information", *Journal of Political Economy*, 82, July/August, pp. 729-754.

Nunez, J. (2000 A) "A Model of Self Regulation with uncertainty about Agents. Department of Economics, Universidad de Chile.

Nunez, J. (2000 B) " A Repeated game of Self Regulation". Department of Economics, Universidad de Chile.

Radner, R. (1985) "Repeated Principal Agent Games with Discounting". *Econometrica* 53: 1173-1198.

Shaked, A. and J. Sutton (1981) "The Self-Regulating Profession", *Review of Economic Studies*, 48, pp. 217-234.

Shapiro, C. (1982) "Consumer Information, Product Quality and Seller Reputation", *Bell Journal of Economics*, 13, Spring, pp. 20-35.

Shapiro, C. (1983) "Premiums for High Quality Products as Returns to Reputations", *Quarterly Journal of Economics*, 98, November, pp. 659-679.

Shapiro, C. (1986) "Investment, Moral Hazard and Occupational Licensing", *Review of Economic Studies*, 53, pp. 843-862.

Spence, A. M. (1977) "Consumer Misperceptions, Product Failure and Producer Liability", *Review of Economic Studies*, 44, October, pp. 561-572.

Spulber, D. (1989) *Regulation and Markets*, MIT Press, Cambridge, Mass.

Tirole, J. (1993) "A Theory of Collective Reputations, with Applications to the Persistence of Corruption and to Firm Quality", MIT working paper.