# A REPEATED GAME OF SELF REGULATION*

JAVIER I. NÚÑEZ ERRÁZURIZ

**Abstract**

*This paper analyses incentives for self-regulation of quality from a principal-agent perspective, in a context of repeated interaction between a Self-Regulatory Organization (SRO) and consumers who can not observe SRO vigilance choice or fraud perfectly. This work unveils five obstacles for positive SRO vigilance to occur in equilibrium. However, this article also shows that 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.*

**Resumen**

*Este artículo investiga teóricamente los incentivos presentes en la autorregulación desde una perspectiva agente-principal y en un contexto de interacción repetida entre una Organización Autorregulada (OA) y consumidores que no pueden observar perfectamente la vigilancia ni el nivel de fraude en la OA. Este trabajo revela cinco obstáculos para la existencia de un equilibrio con vigilancia voluntaria por parte de la OA. Además, se muestra que la regulación pública en paralelo con la autorregulación puede incrementar los incentivos de la OA para monitorear la calidad y reducir el nivel de fraude. De este modo, y contrariamente al conocimiento convencional, una mezcla de regulación pública y autorregulación puede ser deseable, porque se aprovecharía la ventaja informacional de las OAs, mientras que la regulación pública proporcionaría los incentivos para monitorear la calidad que las OAs podrían no poseer de otro modo.*

JEL classification: *L51; K20.*

Keywords: *Self-Regulation, Quality Regulation, fraud, corruption.*

---

* I am indebted to John Vickers, Peter Sorensen, Hyun Song Shin, Rupert Gatti, Meg Meyer and Tim Besley for very valuable comments. All remaining errors are the author's.

□ Departamento de Economía, Universidad de Chile.
E-mail: jnunez@econ.uchile.cl

## 1. Introduction

Self regulation (SR) abounds in a wide range of social institutions such as, for example, financial and professional services, judicial and other public institutions, political parties and even in the monitoring of quality within most firms. Despite its empirical relevance and potential applications, SR has been neglected in the regulation literature. SR has the interesting advantage of avoiding the common problem of asymmetric information between suppliers and an external regulator. However, SR also implies *regulatory capture* by definition, and therefore the mere existence of SR in some industries does not guarantee efficiency. Moreover, theoretical investigation of SR is particularly relevant considering that it is very difficult to empirically assess the performance of an SRO, because it is unclear whether lack of evidence of fraud exposure, for example, actually means that no fraud exists. In this context, the fundamental theoretical question is whether Self Regulatory Organizations (SROs) have the correct *incentives* to monitor quality and expose fraud voluntarily to the public. SR is characterized by three essential facts that must be taken into account in a theoretical attempt to understand it: i) SROs face a principal-agent problem whereby quality is determined ultimately by the SRO agent, ii) SRO incentives to monitor quality are reputation-based, and iii) SR usually occurs in *credence goods* industries.[1] These features of SR have not been properly addressed in the literature. Indeed, in most related works quality is assumed to be observable by consumers. Moreover, quality is assumed typically to be either exogenous, or a direct choice variable of the principal.[2]

A theoretical investigation of the incentives that SROs face requires the analysis of how reputation can be built and sustained by SROs. Reputation, in turn, would depend on the specific informational structure available to consumers from which expectations about quality are to be inferred and updated. 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.

Reputation in economic theory is often understood and modeled 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. Núñez (2000, A, B) study self-regulation from an incomplete information perspective, in which two principal-agent models of self-regulation are developed with endogenous quality. In Núñez (2000 A) consumers have uncertainty about the predisposition of the SRO principal towards vigilance and quality enforcement. In Núñez (2000 B) consumers have uncertainty about the predisposition of the SRO members

[1] Where consumers cannot properly observe product quality either prior or after purchase.

[2] See for example Milgrom (1981), Shapiro (1982), Milgrom and Roberts (1986), Gehrig and Jost (1995) and Emons (1997).

towards fraud and wrongdoing. In both models most equilibria involve no vigilance by the SRO and maximum fraud by the SRO members. In the case of asymmetric information about the SRO principal, in a wide range of cases exposure fails to be an separating equilibrium signal. This is because consumers would not interpret exposure as a positive signal about the commitment of the SRO principal towards vigilance, but rather as a negative one. Thus, the SRO would not wish to engage in any vigilance effort and therefore fraud would be maximum. In the model with uncertainty about the SRO members, for a broad range of circumstances fraud exposure would reduce the expected quality of the remaining (non-exposed) member, eliminating, therefore, any incentives to be vigilant.[3]

The few equilibria with positive vigilance and some fraud deterrence found in these models require rather restricted parameter values and therefore these equilibria appear unlikely. In conclusion, the general picture that emerges from these works suggests that effective self-regulation may be an unlikely outcome, and that at best, it can be effective only under some narrow, specific conditions.

However, there are two common elements in these models that may be driving these results. The first is that these models do not involve intertemporal issues, in the sense that the SRO does not internalize the effects of its behavior over several periods. As it has been shown in many works on reputation building and co-operation, dynamic models often lead to concerns for reputation that are absent or ill-defined in static models, which often lead to the emergence of co-operative behavior.

Secondly, these models focus only on the reputation-based incentives of the SRO in a hidden information framework, where fraud exposure (or absence of it) are the signals employed by consumers to update their beliefs about the SRO type. However, reputation-based incentives for co-operation can also emerge as the result of repeated interaction in a moral hazard framework, instead of the result of asymmetric information. Therefore, it is interesting to investigate whether the results obtained in these works are reinforced in the context of repeated interaction, or whether new incentives for successful self-regulation can be identified.

This work is an attempt to address these two issues. A formal model is developed to introduce intertemporal "reputational" effects on the SRO vigilance decision by considering an infinitely repeated relationship between an SRO and consumers. Therefore, the issue is no longer that of a signaling game as in Núñez (A, B).

There are at least three respects in which the present work departs from the abundant literature on repeated games and co-operation.[4] First, in most of the literature it is assumed that actions are observable. This is not an appropriate assumption for the problem studied here because what characterizes SROs is precisely the difficulty of assessing their vigilance effort and the quality provided. Therefore, in the present work consumers are only able to make decisions based on the existence or absence of fraud exposure. Second, the prob-

---

3 A more informal but extense discussion on these issues as well as other institutional issues on self regulation can be found in Núñez (2000 C).

4 See Fudemberg and Tirole (1993) for an exposition of the literature on repeated games.

ability of exposure is not determined only by the SRO, but instead it is endogenously determined by the interaction between the SRO and its members. The implication of this is that the relationship between the probability of exposure and the SRO vigilance effort may not be positive (or positive enough) to obtain results similar to the relatively few works on repeated games with unobservable actions.[5] Third, this model also studies the role of a third party, namely parallel regulation, in shaping the incentives for co-operation. Núñez (2000 A, B) have shown that parallel regulation is likely to enhance the SRO's incentives to be vigilant. It is interesting to analyze whether this would still be the case under the repeated interaction context studied here.

Section 2 presents a simple repeated game of SRO co-operation, with discrete vigilance choice and with permanent punishment by consumers if fraud is exposed. The necessary condition for co-operation is then compared to the full observation benchmark. Section 3 presents a more elaborate model with continuous vigilance effort and transitory punishment of variable length. The assumption of continuous vigilance will allow to analyze the effect of several parameters on the equilibrium levels of vigilance and fraud. Section 4 analyses the role of parallel regulation. Finally, section 5 discusses the main conclusions and limitations of this work.

## 2. A Simple Model

In this section a simple model is presented as a way of example to illustrate the main issues and forces at work. In the next sections a more formal and general model is developed.

**SRO Member's Behavior.** In this model, the principal is the SRO and the only agent is the SRO member. Unlike Núñez (2000 A, B) both players are of known type, so that there are no hidden information issues. There are infinite periods. In every period, the SRO member maximizes;

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

where $x$ is the level of fraud, $p(x,y)$ is the probability of fraud discovery by the SRO, which depends on the fraud level and on the level of SRO vigilance $y$ and $T > 0$ is the penalty imposed on the member if found doing fraud. It is assumed that $p_x$, $p_y$ and $p_{xy}$ are positive. The first and second-order conditions are $p_x T = 1$ and $p_{xx} T < 0$, respectively. Accordingly, it is assumed that $p_{xx} < 0$. Some of these assumptions deserve detailed justification. 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.[6] It is

5 Such as Green and Porter (1984) and Radner (1985).

6 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.

unclear, however, how and whether different modeling 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 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 to the agent due to fraud, such as extracting resources unlawfully and directly from consumers, or from inducing unnecessary treatment or repairs to consumers.[7] Alternatively, *x* can reflect the saving of costs (such as effort) 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.[8]

**Consumers' Behaviour.** There are two possible information structures for the consumers. In the observable actions case, consumers are able to observe either the level of fraud, or the level of SRO vigilance. In the case with non-observable actions, consumers can only observe whether exposure occurred or not. Both informational structures will be considered in turn, and the different outcomes are compared to illustrate the role of the information available to consumers on the effectiveness of SROs. In the full observation case, it will be assumed that if consumers observe either a low level of vigilance or a high level of fraud, then consumers will punish the SRO by reducing their purchases from the SRO and turning to other substitutes thereafter. In the case where neither vigilance nor fraud levels are observable, consumers have to rely only on the evidence of fraud discovery. It will be assumed here that if fraud becomes known to the public, then consumers will punish the SRO by cutting their purchases from then onwards. Let $_L$ denote the per period payoff to the SRO should this happen. On the other hand, if a high level of vigilance or low fraud is observed in the full-observation benchmark, or if fraud is not observed in the non-observation benchmark, then consumers will continue to purchase the good. Let $_H$ be the corresponding per period payoff to the SRO in these cases.

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

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

### 2.1. Full Observation Benchmark

For simplicity, and as it is customary in repeated games, it is assumed that the SRO has a discrete choice to make between a high and low level of vigilance. Let $y_h$ and $y_l$ denote the high and low levels of vigilance, respectively. Vigilance is costly for the SRO. Let the cost difference between high and low vigilance be equal to $c$. Let $\delta$ be the discount factor. Then, the SRO chooses between $y_h$ and $y_l$ to maximise the present value of net income flows. Choosing high vigilance yields;

(2) $$\frac{\delta}{1-\delta}\left(\pi_H - c\right)$$

whereas choosing low vigilance yields;

(3) $$\delta\pi_H + \frac{\delta^2}{1-\delta}\pi_L$$

Therefore, the SRO will be willing to choose to high vigilance only if

(4) $$\delta > \frac{c}{\pi_H - \pi_L}$$

If it is assumed that $\pi_H - c > \pi_L$, then this condition is always satisfied if $\delta$ is sufficiently close to unity. Condition (4) implies that if the SRO cares enough about the future and if vigilance is not too costly relative to the cost of having a lower demand, then the SRO will be willing to be vigilant. However, this story assumes that either vigilance or fraud is observable by consumers at all times. As discussed in the introduction, this assumption is not satisfied in credence good industries. If neither vigilance nor fraud are observable, then it is possible that the SRO might be willing to hide fraud from the public, once it is privately discovered within the SRO. However, there is a chance that consumers can become aware of the existence of fraud by other means. For example, evidence of fraud can leak out from the SRO to the public or to the media, which could be sometimes out of the control of the SRO. In addition, there can be parallel regulation by a public agency or any other third party, which can also lead to involuntary exposure. The next model addresses these issues, which may change the SRO's incentives towards vigilance and fraud exposure.

### 2.2. Non Observation Benchmark

It is assumed that consumers cannot observe vigilance or fraud levels. Instead, in every period consumers will only be able to observe either fraud exposure, or non-exposure. The event of non-exposure can happen with both high and low vigilance. Let $\gamma$ be the probability that the discovery of fraud leaks out to the general public. The SRO has to choose between exposure and non-exposure, in addition to the vigilance choice. Let $e = 0,1$ denote the choices of exposure and non-exposure, respectively. Therefore, the probability of exposure given some level of vigilance becomes;

(5) $$\text{Prob (exposure)} = \begin{cases} p(\cdot) & if\ e^* = 1 \\ \gamma\, p(\cdot) & if\ e^* = 0 \end{cases}$$

Given that now the event of non-exposure does not necessarily imply that $y = y_h$, then consumers will be willing to purchase and pay only according to the expected value of fraud, which equals $E(x) = Prob(y = y_h)\, x^*(y_h) + Prob(y = y_l)\, x^*(y_l)$. The probabilities of vigilance being high or low will be consistent with optimal behavior by the SRO. In particular, if $y_h$ is optimal for the SRO, then consumers will infer that $Prob(y = y_h) = 1$ and therefore $E(x) = x^*(y_h)$. Otherwise, consumers should infer that $E(x) = x^*(y_l)$. Accordingly, redefine $\pi_H = \pi(x^*(y_h))$ and $\pi_L = \pi(x^*(y_l))$. The two choices to be made by the SRO are analyzed next. It is simple to show that if fraud has been discovered, then the SRO would always avoid its exposure, which is stated in the next result.

**Result 1.** $e^* = 0$. *Therefore, fraud cover-up is optimal*.

The exposure choice is relevant only for the cases in which fraud has been discovered. The present value of exposing fraud is;

(6) $$\begin{aligned} &\delta(\pi_H - c) + \frac{\delta^2}{1-\delta}\pi_L \quad if\ y = y_h \\ &\delta\pi_H + \frac{\delta^2}{1-\delta}\pi_L \quad if\ y = y_l \end{aligned}$$

The present value of not exposing fraud is more complex, because, unlike exposure, non-exposure may lead to different histories in future periods. However, it can be shown that exposure is a dominated action because non-exposure will outperform exposure in all circumstances. Indeed, suppose that, given that exposure has not occurred until then, from the third period onwards the SRO gets the *lowest* possible payoff, namely $(1-\gamma)\frac{\delta^3}{1-\delta}\pi_L$. Then the minimum payoff to the SRO if non-exposure is chosen will be; $\delta(\pi_H - c) + \gamma\frac{\delta^2}{1-\delta}\pi_L + (1-\gamma)\delta^2(\pi_H - c) + (1-\gamma)\frac{\delta^3}{1-\delta}\pi_L$ if $y = y_h$, and $\delta\pi_H + \gamma\frac{\delta^2}{1-\delta}\pi_L + (1-\gamma)\delta^2\pi_H + (1-\gamma)\frac{\delta^3}{1-\delta}\pi_L$ if $y = y_l$. Therefore, by comparing the payoffs from exposure and the worst possible outcome under non-exposure, it can be seen that non-exposure always outperforms exposure, independently of future circumstances which groves result 1. The next step is to analyze the SRO's vigilance incentives. If fraud leaks out to the public, then from the subsequent period consumers would cut their purchase of the good and therefore the SRO would get $\pi_L$ thereafter. If the SRO chooses not to expose fraud (if found), then the probability that fraud is actually discovered *and* leaks out to the public is $\gamma p(\cdot)$. Therefore, the probability that evidence of fraud becomes available to the public for the first time at year *t* is $(1-\gamma p(\cdot))^{t-1}\gamma p(\cdot)$. Naturally, $\sum_{i=1}^{\infty}(1-yp(\cdot))^{i-1}\gamma p(\cdot) = 1$. If fraud becomes known to the public for the first time at time *t*, then the present value to the SRO from co-operation corresponds to $(\pi_H - c)\sum_{j=1}^{t}\delta^j + \frac{\delta^{t+1}}{1-\delta}\pi_L$. The SRO gets $\pi_H - c$ in periods 1 to $t$, and $\pi_L$ from $t$ thereafter. If the SRO chooses high vigilance as long as no exposure occurs, then the present value corresponds to;

$$\sum_{i=1}^{\infty}(1-\gamma p_h)^{i-1}\gamma p_h\left[(\pi_H - c)\sum_{j=1}^{i}\delta^j + \frac{\delta^{i+1}}{1-\delta}\pi_L\right] \tag{7}$$

where for simplicity of notation, $p_h = p(x^*(y_h), y_h)$. Similarly, the present value of choosing low vigilance in every period is;

$$\sum_{i=1}^{\infty}(1-\gamma p_l)^{i-1}\gamma p_l\left[\pi_H\sum_{j=1}^{i}\delta^j + \frac{\delta^{i+1}}{1-\delta}\pi_L\right] \tag{8}$$

where $p_l = p(x^*(y_l), y_l)$. In order to assess the SRO's optimal strategy, expressions (7) and (8) must be compared to each other. Regrouping terms, factorizing by $\pi_H$, cancelling terms and collapsing the discount series, it follows that the SRO will be willing to commit to high vigilance if;

$$c < (\pi_H - \pi_L)\frac{\gamma\,\delta(p_l - p_h)}{1-\delta(1-\gamma\, p_l)} \tag{9}$$

It is interesting to compare condition (9) to the situation where either vigilance or fraud levels are fully observable. Note that if $p_l = 1$, $p_h = 0$ and $\gamma = 1$ then condition (9) is identical to the condition for high vigilance in the full observation benchmark, namely $c < (\pi_H - \pi_L)\delta$. This fact confirms intuition; if exposure is fully informative about low vigilance, and if evidence of fraud always leaks out to the public, then observing exposure is equivalent to observing either fraud or vigilance levels. The full observation case is, therefore, a special case of this model. Note that condition (9) is more restrictive than the case where SRO vigilance or fraud are observable by consumers. In other words, high vigilance and fraud deterrence are more difficult to achieve when consumers can only rely on a noisy, imperfect signal of product quality. This leads to the following result.

**Result 2**. *A necessary condition for an equilibrium with high vigilance and some fraud deterrence is that* $c < (\pi_H - \pi_L)\frac{\gamma\delta(p_l - p_h)}{1-\delta(1-\gamma p_l)}$. *However, this is not a sufficient condition because the pair* ($y^* = y_l$, $x^* = x(y_l)$) *is also an equilibrium of the repeated game.*

This condition is necessary, but not sufficient for high vigilance and low fraud because the situation where the SRO optimally chooses low vigilance in response to consumers' prior expectations of low vigilance constitutes another equilibrium of the game. This fact raises the problem of multiple equilibria and equilibrium selection that are common in the literature of repeated games.

Condition (9) can be broken down into several necessary conditions, namely, i) $\pi_H - c > \pi_L$, ii) $\delta > 0$, iii) $\gamma > 0$, and iv) $p_l > p_h$. The first condition implies that the difference between stage payoffs from high and low vigilance must exceed the vigilance cost differential. The second condition states that the SRO must have some concern for future payoffs, otherwise the SRO will not have an incentive to invest in its reputation. These two conditions resemble the customary requirements for co-operative behaviour in repeated games. However, the last two conditions are less straightforward, and also it is unclear whether they are likely to be satisfied. The third condition suggests that if the probability that fraud leaks out to the public is sufficiently low, then the SRO will not have adequate incentives to engage in high vigilance. Note that $\lim_{\gamma \to 0} (\pi_H - \pi_L) \frac{\gamma \delta (p_l - p_h)}{1 - \delta (1 - \gamma p_l)} = 0$. As $\gamma \to 0$ then fraud leaks out for the first time increasingly later in time. Thus, the present value of choosing low vigilance will approach $\frac{\delta}{1-\delta} \pi_H$, which exceeds the present value of choosing high vigilance. Finally, the fourth condition states that high vigilance requires the probability of fraud discovery to be decreasing in vigilance. This condition, however, is not necessarily satisfied. Recall that the probability of fraud discovery depends not only on vigilance, but also on the fraud level. In Núñez (2000 A) it was shown that for a wide range of conditions the probability of fraud discovery is indeed decreasing in vigilance. However, this is not necessary the case, and in any case the rate at which $p(\cdot)$ decreases in $y$ might not be high enough to satisfy condition (9).

Summing up, the basic implications of this model are the following; i) it must be possible that fraud can become known to the public despite the attempts of the SRO to cover it up, and ii) fraud discovery must be more likely the lower is the vigilance level. However, these are not sufficient conditions because low vigilance and high fraud also constitute a plausible equilibrium. With respect to the first condition, the SRO might want to commit to fraud exposure, but that commitment was proven to be not credible, because fraud exposure can only trigger a punishment phase. However, the SRO would like to find some means of making fraud known to the public, should it be discovered. This is because the present value of building a reputation for high vigilance and low fraud yields a higher payoff that staying in a low vigilance, high fraud equilibrium. Some actions the SRO may encourage are, for example, an open-book policy, provide easy access to a public agent, to consumers, or to a third party. If these means become credible to the public, then the SRO could optimally build up a reputation for high vigilance and low fraud. However, the likelihood and credibility of such forms of commitment to exposure are uncertain and are a rather empirical matter.

There are, however, at least four main limitations of this simple model. The first is the assumption of punishment of infinite duration. As it has been pointed out in the literature, this seems a very grim strategy. A preferable approach would be to have a variable length of punishment, which can have two possible interpretations. First, consumers simply may not be willing to engage in punishment for ever, and even they may want to choose an optimal length of punishment in order to maximize their net present value. Second, a variable length of punishment can reflect the possibility that consumers may have a limited "memory".[9] Therefore, after some periods of punishment, consumers can forget about the SRO past mistakes and a high vigilance equilibrium becomes possible again. In this context, it would be interesting to study the role of the nature and duration of consumers' memory on the SRO incentives to be vigilant.

---

9 This may also be the case in an overlapping generations model in which new generations cannot observe events before their appearance, supposing they cannot (or would not) learn from their ancestors' testimonies.

The second limitation of this model is that consumers may not be willing to engage in a punishment phase after observing only *one* outcome of exposure. This may be too harsh on the SRO because exposure is a possible event even if the SRO is being vigilant. In this context, consumers may be willing to observe the outcome of several periods before deciding whether to start a punishment phase. This would have the advantage that the SRO's behaviour can be estimated more accurately than the former case.

Third, the basic model presented above can be modified to allow the SRO to have a continuous choice of vigilance instead of a discrete one. Apart from being a more satisfying assumption, this would allow studying the effects of several parameter values on the equilibrium levels of fraud and vigilance.

Finally, this model did not address many issues related to the role of parallel regulation in creating incentives for successful self-regulation. These four issues are tackled in the next section in a more general repeated game of self-regulation.

## 3. A More General Model

### 3.1. Variable Length of Punishment

Every period the SRO chooses a level of vigilance $y \in (0, \infty]$, and the SRO member chooses a level of fraud $x \in (0, x_{max})$ after observing the SRO's move. $x_{max}$ is the full level of fraud and is defined by the member's optimal response to zero vigilance, that is, $x_{max} = x^* (y = 0)$. Therefore, the SRO behaves as a Stackelberg player because it internalizes the effect of vigilance on the SRO member's behaviour. Consumers are unable to observe the SRO's vigilance level or fraud directly, and they can only verify whether exposure happened in the last period. If exposure occurs, then a punishment phase of length $P \leq \infty$ starts. Once the punishment is over at period $P + 1$, a new cycle of high vigilance, low fraud and high consumer demand may start again.

Let $v(\delta)$ be the discounted present value of payoffs to the SRO, which equals

$$(10)\quad v(\delta,y) = (1-\delta)\,(\pi^*-c(y))+\alpha_e[\delta(1-\delta^P)\pi_L+\delta^{P+1}v(\delta,y)]+(1-\alpha_e)\delta v(\delta,y)$$

where $\alpha_e$ is the probability of exposure, $\pi^*$ is the equilibrium transfer from consumers to the SRO, and $c(y)$ represents the cost of vigilance, which satisfies $c_y$, $c_{yy} > 0$ and $\lim_{y\to\infty} c_y = \infty$.

Unlike the previous model with binary vigilance choice, in this model $\pi^*$ is not a decreasing function of expected fraud. Instead, $\pi^*$ is now assumed exogenous, although it will be taken as parametric by the SRO to choose $y$ optimally. Unlike the previous example, this will provide a continuous set of equilibrium levels of vigilance and fraud deterrence. The only constraint to $\pi^*$, as it will be established later, is that it does not exceed consumers' benefit for the expected (equilibrium) level of fraud.[10]

[10] This is to ensure non-negativity of consumers' surplus for the equilibrium expectation of quality.

Let $w(e)=\delta(1-\delta)\pi_L+\delta^{P+1}v(\delta,y)$ be the continuation payoffs if exposure occurs, and let $w(n)=\delta v(\delta,y)$ be the continuation payoffs when there is no exposure. Then, it is straightforward to verify the following result.

**Result 3**. *If $\pi_L < \pi^*$, then $w(e) < w(n)$, and therefore under these circumstances it is never optimal for the SRO to expose fraud if it is discovered. Therefore, $\alpha_e = \gamma\ p(\cdot)$.*

This result is equivalent to the one obtained in the previous model with infinite length of punishment and discrete choice. It states that if during the cooperative phases the SRO gets paid more than during the punishment phases, then the SRO will not have any incentives to expose fraud. Therefore, this result implies that it should be common-knowledge that voluntary exposure by the SRO will never occur, which implies that $\alpha_e = \gamma\ p(\cdot)$. Therefore, fraud exposure can happen only if it leaks to the public against the SRO's will.[11] This simple result allows redefining the appropriate maximization problem for the SRO such that now the SRO will choose a vigilance level to maximize the discounted present value;

$$v(\delta,y)=(1-\delta)(\pi-c(y))+\gamma p(\cdot)[\delta(1-\delta^P)\pi_L+\delta^{P+1}v(\delta,y)]+(1-\gamma p(\cdot))\delta v(\delta,y) \tag{11}$$

and then solving for $v(\cdot)$ yields,

$$v(\delta,y)=\frac{(1-\delta)(\pi-c(y))+\gamma p(\cdot)\delta(1-\delta^P)\pi_L}{1-\gamma p(\cdot)\delta^{P+1}-(1-\gamma p(\cdot))\delta} \tag{12}$$

Note that $v(\cdot)$ is decreasing in $\gamma p(x,y)$ and in $P$, confirming the intuition that a higher likelihood of exposure, and a longer punishment phase decrease the discounted present value of the SRO's payoffs.

**Consumers' Behaviour.** Consumers' behaviour is defined as follows. As long as they do not observe exposure, they will pay the SRO an amount equal to $\pi^*$. If they observe fraud exposure, then the transfers to the SRO will be cut to $\pi_L$ for $P$ periods. Then consumers' discounted present value will amount to;

$$\begin{aligned} v(\delta,y)=(1-\delta)(B(Ex)-\pi)+\gamma p(\cdot)\big[\delta(1-\delta^P)B(x_{max})+\delta^{P+1}v(\delta,y)\ \big] \\ +(1-\gamma p(\cdot))\delta v(\delta,y) \end{aligned} \tag{13}$$

where $B(Ex)$ is the total expected benefit of purchasing goods or services that have an expected fraud $x$. Therefore, $B(x)-\pi$ is the consumer's surplus. It is assumed that $B' > 0$ reflecting that quality is beneficial for consumers.

### 3.2. Equilibrium

The SRO decision problem can be analyzed as a discounted dynamic programming problem in which the SRO has to choose the current level of vigilance, given that the continuation payoffs and the future decisions are optimal.

---

[11] Consequently, in what follows the terms "leak exposure" and simply "exposure" will be employed interchangeably.

Then, because for all periods during the co-operative phase the decision problem is identical, the optimal choice will accordingly be the same every period in this phase. Recall that $w(e)=\delta(1-\delta^{p})\pi_L+\delta^{p+1}v(\delta,y)$ is the continuation payoff if exposure occurs, and $w(n)=\delta v(\delta)$ be the continuation payoff if no exposure is observed. Then, during a co-operative phase the SRO maximizes,

$$v(\delta,y)=(1-\delta)(\pi-c(y))+\gamma p(y)w(e)+(1-\gamma p(y))w(n) \tag{14}$$

There are two possible types of equilibria. First, the situation involving no vigilance and maximum fraud is always an equilibrium of the repeated game. Indeed, if consumers pay $\pi_L$ to the SRO during a co-operative phase then it does not pay off to engage in any vigilance, which reinforces consumers' initial decision. Therefore, $y^* = 0$, $x^* = x_{max}$ and $\pi^*=\pi_L$ is always a Nash equilibrium. The second possible equilibrium in this game will involve the alternation of co-operative and punishment phases, and the equilibrium conditions must be established for each case. During the co-operative phases, an equilibrium with positive vigilance and fraud deterrence will be characterized by;

$$\frac{dv(\delta,\pi^*)}{dy}=0 \tag{15}$$

$$\frac{du(y^*,T)}{dx}=0 \tag{16}$$

and

$$\pi^* \leq B(Ex = x^*(y^*)) \tag{17}$$

where $y^*$ satisfies $\frac{dv(\delta,\pi^*)}{dy}=0$ and $x^*$ satisfies $\frac{du(y^*)}{dx}=0$. Consumers must pay less than the benefit they extract from the expected level of quality in equilibrium to ensure that consumer's surplus is non-negative. Therefore, in equilibrium, the SRO, the SRO member and consumers are able to predict each other's behaviour and additionally all parties are behaving optimally in response. Differentiating $v(\cdot)$ with respect to $y$ and simplifying yields,

$$\frac{dv(\delta,\pi^*)}{dy}=\frac{dp(\cdot)}{dy}\gamma\delta[w(e)-w(n)]-(1-\delta)c_y \tag{18}$$

In order to have an equilibrium with positive vigilance, then $\frac{dv(\delta,\pi^*)}{dy}$ must be positive for at least some vigilance level. The term $[w(e) - w(n)]$ is certainly non-positive, as it was assumed that $\pi^*>\pi_L$. Therefore, for $\frac{dv(\cdot)}{dy}$ to be positive at least for some $y$, $\frac{dp(\cdot)}{dy}$ must be negative. If the latter is non-negative, then $v(\cdot)$ is maximised at $y^* = 0$. This confirms the result obtained in the previous example; if vigilance increases the chances of fraud discovery, then a Stackelberg SRO cannot do any better than choosing zero vigilance, which will lead to maximum fraud.

**Result 4**. *Necessary conditions for an equilibrium satisfying some vigilance and some fraud deterrence ($y^* > 0$ and $x^* < x_{max}$) are; i) $\gamma > 0$, ii) $\delta > 0$, iii) $\frac{dp(\cdot)}{dy}<0$. Otherwise, the unique equilibrium of the game is $y^* = 0$, $x^* = x_{max}$, and*

$\pi^*=\pi_L$. *However, these conditions are not sufficient because the latter is always an equilibrium*.

These necessary conditions are equivalent to those obtained in the previous simple model. The first two conditions confirm intuition. Fraud deterrence is possible only if there is a threat of leak exposure. Moreover, positive vigilance requires the SRO to be concerned about future payoffs. The third condition implies that positive vigilance will be desirable only if it reduces the probability of exposure. This, however, is not guaranteed because vigilance has a partial positive effect on the probability of fraud discovery. In order to characterise the set of equilibria with positive vigilance and fraud deterrence, the expression for consumers' surplus and the SRO's optimal choice of vigilance for a given value of $\pi$ must be investigate. This can be investigated by differentiating the SRO FOC totally, which yields,

(19) $$\frac{dy^*}{d\pi} = -\frac{\frac{dp}{dy}\gamma\delta\frac{d\left[w(e)-w(n)\right]}{d\pi}}{\frac{d^2p}{dy}\left[w(e)-w(n)\right]-c_{yy}}$$

Note that the numerator is always positive if $\frac{dp(\cdot)}{dy}<0$. The slope will necessarily be positive and will tend to infinity for a sufficient high level of vigilance because $c_{yy}>0$ tends to infinity as $y$ increases. Therefore, for sufficiently high values of $\pi$, SRO vigilance will no longer be responsive to changes in $\pi$. Recall that $B_x<0$ and therefore $\frac{dB}{dy}=B_x\frac{dx^*}{dy}>0$ because $\frac{dx^*}{dy}<0$. Therefore, the analysis of the SRO optimal vigilance and consumers' total benefit function imply that there will be an intersection of the two functions that will establish an upper limit to the set of Nash equilibria of the game. This situation is depicted in Figure 1.

FIGURE 1

SET OF EQUILIBRIA LIES ON $y^*(\pi)$

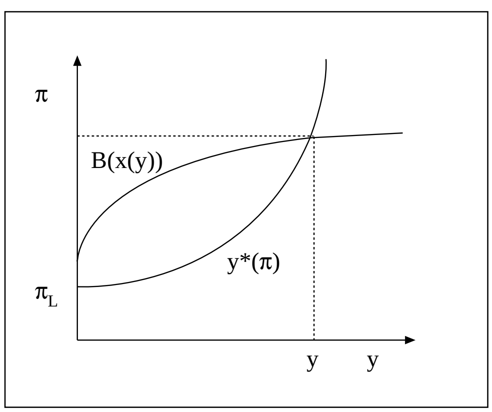

All the combinations of $\pi$ and SRO vigilance on the function y*($\pi$)that also satisfy the consumer's "participation constraint", namely $B(x)>\pi*$ will constitute the set of possible equilibria during the co-operative phases.[12]

### 3.3. Equilibrium Punishment Phases

In repeated games with observable actions, punishment phases are never observed in equilibrium, although it is precisely the credible threat of punishment what disciplines the players so that punishment need not be carried out. However, in this game, the non-observation of actions implies that punishment phases will eventually be observed in equilibrium. Indeed, if $\gamma p(x^*,y^*)>0$ so that exposure can happen even if the SRO is being vigilant, then consumers will start a phase of punishment as soon as fraud is exposed.[13] In equilibrium, however, consumers will not punish the SRO as a result of their believing that the SRO deviated from the equilibrium level of vigilance. Indeed, consumers can forecast perfectly the SRO's actions in equilibrium, and therefore they would know that if exposure is observed, it must be due to "bad luck", instead of deviation from equilibrium vigilance. In this model punishment provides the incentives for high vigilance by forcing the SRO to reduce the probability of bad luck by being more vigilant. This differs from repeated games with observable actions, where punishment is intended to happen only if deviation occurs, and consequently it is not observed in equilibrium.

The incentives towards vigilance are, as expected, increasing in the length of the punishment phase. This is analysed in the next derivative.

(20) $$\frac{d^2v(\cdot)}{dydP}=\frac{dp(\cdot)}{dy}\gamma\delta\frac{w(e)}{dP}>0$$

Clearly $\frac{w(e)}{dP}<0$ because the longer the punishment phase is, the lower the continuation payoffs after exposure would be. Therefore, $y^*$ is increasing in $P$ implying that longer punishment periods will provide enhanced incentives for the SRO to be vigilant.

### 3.4. Several Revision Periods

The model above has the disadvantage that a single exposure may not be persuasive enough about the misconduct of the SRO to make consumers trigger a costly punishment phase. To address this issue, the previous model can be modified to allow consumers to observe the outcome of several periods before deciding whether to engage in punishment or to continue purchasing the good. Radner (1985) developed a principal-agent model where the principal follows a "review strategy", in order to chose a money transfer to the agent. In a review strategy, the principal observes the agent's actions during a predetermined number of periods (namely a review phase) and compares the number of periods the agent co-operates to a fixed target, after which she decides on whether to agree

12 It is assumed that the lowest quality provides some consumers' surplus so that $B(\cdot)>\pi_L$.
13 This is also the case in Green and Porter (1984).

to pay the agent the co-operative transfer or start a punishment phase. This section adapts Radner (1985) to the self-regulation problem studied here. Let *R* denote the length of the review phase, in which consumers observe the outcomes of exposure or non-exposure, before deciding the next course of action. In this context, the probability that the agent fails the consumers' review is no longer $p(x,y)$. Now consumers will decide what to do based on the proportion of exposures during the review phase. Let $E_R$ denote the number of exposures observed during the review phase *R*. Then, passing the review will require that $E_R \leq R\gamma\, p(x(y),y)+M$, where $M > 0$ can be interpreted as a "margin of error" established by consumers. For example, if consumers followed a statistical inference analysis, then $M = -z\sqrt{R\gamma\, p(\cdot)(1-\gamma\, p(\cdot))}$ where *z* is the corresponding z-statistic table value corresponding to the desired level of confidence.[14] The term $R\gamma p(x(y),y)$ is the expected number of exposures over the review phase of length *R* if the SRO is co-operating. Let $\psi$ denote the probability that the SRO fails the review that is, $\psi = Prob(E_R > R\gamma\, p(x(y),y)+M)$. Therefore, the SRO will now choose *y* to maximise the present value of payoffs defined as,

$$v(\delta,y) = (1-\delta)\sum_{i=1}^{R}\delta^{y-1}(\pi - c(y)) + \psi\left[\delta^R(1-\delta^P)\pi_L + \delta^{R+P}v(\delta,y)\ \right] + (1-\psi)\delta^R v(\delta,y) \tag{21}$$

so that solving for $v(\cdot)$ yields,

$$v(\delta,y) = \frac{(1-\delta)\sum_{i=1}^{R}\delta^{i-1}(\pi - c(y)) + \psi\delta^R(1-\delta^P)\pi_L}{1-\psi\delta^{R+P} - (1-\psi)\delta^R} \tag{22}$$

Note that equation (22) has an identical structure to the present value in the previous case. In fact, the two cases are identical when $R = 1$ and if $\gamma\, p(x,y)$ is replaced by $\psi$. Therefore, by analogy the results derived in the previous case are also valid in this case, namely the necessary conditions for $y^* > 0$ and the effect of *P* on $y^*$.

The assumption that $R > 1$ has an odd property, however, namely that *R* is fixed and not sensitive to the sequence of events during the review phase. For example, it is plausible to expect that an unusually high proportion of exposures early in a review phase would make the SRO fail the review regardless of the subsequent outcomes in the review. In this context it is not clear why consumers would wait until the end of the review phase to carry out the punishment, neither it is clear why the SRO would not give up co-operating, knowing that the chances of passing the review are slim. Moreover, this fact reinforces consumers' incentives to engage in punishment as early as it becomes evident that the SRO will fail the review. It seems more appropriate to assume an endogenous length of the review phase so that punishment can be decided de-

---

14 However, as it is customary in statistical theory, this analysis would make sense only for a reasonable long review phase satisfying $R\gamma\, p(\cdot)\geq 5$, which may not be necessarily the case.

pending on the cumulative proportion of exposures. A problem with this alteration is that now the probabilities of exposure will no longer be independent, as the SRO will take into account the past events in the current review period.[15] This also imposes more complexity on consumers' statistical inference process and makes the model very complex. There is another unappealing property of $R > 1$, namely that consumers' demand and willingness to pay must follow anti-intuitive paths in time. For example, after an exposure the SRO will increase its vigilance effort to avoid further exposures in the review period, which consumers should realize. Then, assuming that consumers are not committed to a fixed reward strategy during the review phase, after an exposure $B(E(x^*))$ should *increase* and after non-exposure it should *decrease*. This fact seems to defy both intuition and empirical evidence. Following the arguments above, the case with $R = 1$ appears a simpler and more realistic modelling option.

## 4. Parallel Regulation

The main conclusion of the previous model is that there are several conditions for successful self-regulation. In addition to the requirement of a sufficiently high discount rate, positive vigilance requires the probability of fraud discovery to be decreasing in vigilance, once the effect of the agent's optimal reaction to increased vigilance is taken into account. Secondly, positive vigilance also requires that there is some minimum probability that evidence of fraud can leak out to the public. It is not clear the extent to which these conditions are likely to be satisfied in practice. However, even if they are satisfied, it is interesting to examine whether self-regulation will provide *sufficient* incentives for vigilance and fraud deterrence, and whether there is any room for public regulation to fight fraud and wrongdoing. In Núñez (2000 A, B) the main conclusions in this respect were that parallel public regulation very often enhances the SRO vigilance incentives. These conclusions were derived in a context where self-regulation was modelled as a signalling game in a hidden information context. It is interesting to analyse whether parallel regulation can improve self-regulation in the repeated interaction framework studied here, and if so, under what conditions. Parallel regulation can be fitted in the model above in two fundamental ways. First, parallel regulation can focus on investigating the vigilance effort of the SRO and follow the evolution of any fraud discovery made by the SRO. This can be done, for example, by having access to SRO internal reports, to SRO meetings and to interviewing its personnel. Second, parallel regulation may operate by monitoring the existence of fraudulent behaviour directly, and rather independently of the SRO vigilance effort. This can be done, for example by undertaking random checks on fraud and quality provision or by gathering feedback from consumers.[16] However, this form of

---

15 For example, suppose that exposure occurs early in a review period. Then in the next period, the SRO may increase its vigilance effort to reduce the chances of another exposure that would put at risk the fulfillment of the non-exposure quota.

16 This is not inconsistent with the assumption of credence goods. Consumers can provide evidence of malpractice even if they are not aware of it, for example if they were lead to believe that a bad outcome was due to bad luck, as opposed to negligence.

parallel regulation assumes that the public regulator is technically prepared to detect and expose fraud directly and independently of the SRO. Instead, in the former case, the regulator has to rely on the SRO's knowledge, and the regulator's main role is to make public any evidence of fraud within the SRO. Naturally, these forms of parallel regulation are not necessarily mutually exclusive.

The first form of regulation can be captured by the parameter g. The more the regulator investigates the SRO's vigilance activities, the more likely it is to find and expose evidence of fraud, for a given level of fraud. It is clear from the SRO maximisation problem and FOC that optimal vigilance is monotonically increasing in γ if $\frac{dp}{dy} < 0$.[17] In this case, parallel regulation does not increase the total probability of fraud discovery. The only party capable of finding fraud is the SRO, and parallel regulation can only make fraud known to the public. This provides additional incentives for the SRO to reduce fraud and minimise the probability of fraud exposure. It is interesting to note that in these circumstances, this form of parallel regulation would be desired by the SRO. Parallel regulation would signal to consumers that if fraud is not observed, it is because they "have nothing to hide". As a result, this reduces fraud and increases $\pi^*$, and the present value of SRO payoffs.

The second type of parallel regulation can be modelled by introducing an additional variable to represent the probability of fraud discovery by the public regulator. Let $q(x)$ be this probability with $q_x > 0$. At this point it is necessary to refer to the relationship between leak and regulator exposure. These two types of exposure may be seen by consumers as either mutually *exclusive* or *non-exclusive* events, depending on whether they believe that joint exposure can happen with positive probability. There exist justifications for both assumptions. Consider two possible scenarios in which exposure can take place. In the first, exposure may happen anytime within a period and there is no reason why fraud would not be exposed immediately after discovery. In these circumstances, if exposure occurs by either party, consumers can be certain that the other party has no evidence of fraud. In this case, both types of exposure will be exclusive events.[18] In the second scenario, either or both parties can be constrained to report their findings of fraud at the end of discrete periods, say once a year. In this case, multiple fraud discovery is certainly a possibility. For example, if there is leak exposure in a period, it can still be the case that the regulator may have had evidence of fraud already but did not disclose it.

### 4.1. Leak and Regulator Exposure as Exclusive Events

In this case, the SRO maximization problem becomes,

(23) $$v(\delta, y) = (1-\delta)\big(\pi - c(y)\big) + \big[\gamma p(\cdot) + q(x)\big]\big[w(e) - w(n)\big]$$

[17] If $\frac{dp}{dy} \geq 0$, then optimal vigilance is zero, regardless of the value of γ.

[18] If there were multiple exposure but at different times, the situation is still one of mutually exclusive events because the first exposure would have triggered the punishment phase already.

and the FOC is $\frac{dv(\delta,\pi^*)}{dy} = \left[\gamma \frac{dp(\cdot)}{dy} + q_x \frac{dx}{dy}\right][w(e) - w(n)] - (1-\delta)\, c_y = 0$. Now the agent chooses $x$ to maximize;

(24) $$x - \left[p(x,y) + q(z,x)\right]T$$

and the FOC and SOC become $1 - [p_x + q_x]T = 0$ and $-[p_{xx} + q_{xx}] < 0$. Differentiating the FOC totally, the slope of the agent's reaction function under parallel regulation is,

(25) $$\frac{dx}{dy} = -\frac{p_{xy}}{p_{xx} + q_{xx}}$$

The denominator is positive because it is the negative of the SOC. Therefore, $\frac{dx}{dy}$ will be unambiguously negative. However, note that if $q_{xx} > 0$, the slope would be less negative, implying that optimal fraud would be less sensitive to SRO vigilance than it is the case without parallel regulation. The next step is to investigate the position of the agent's reaction function relative to the case without parallel regulation. By inspection of the agent's FOC it follows that optimal fraud will necessarily be lower under parallel regulation for any given level of SRO vigilance. Indeed, since $q_x > 0$, then it follows that $p_x$ must be reduced to satisfy the FOC, and this must necessarily be done by reducing fraud because $p_{xx} > 0$. Therefore, although the slope of the agent's reaction function may be flatter under parallel regulation, optimal fraud must be lower for any level of SRO vigilance. This is shown in Figure 2.

Let $y^*_r$ and $x^*_r$ denote the equilibrium levels of SRO vigilance and fraud under parallel regulation, and given a fixed value of $\pi$. Also, let $x^r_{max} = x^*_r(y^*_r)$, i.e. $x^r_{max}$ is the optimal level of fraud with parallel regulation when SRO vigilance is zero. From the discussion above it follows that $x^r_{max} < x_{max}$.

**Result 5**. *Let leak exposure and regulator exposure be seen as exclusive events by consumers. Then, for a certain value of p, if* ($y^* = 0$, $x^* = x_{max}$)*, the equilibria under parallel regulation are* ($y^*_r > 0$, $x^*_r < x^r_{max}$) *if* $\left[\gamma \frac{dp(\cdot)}{dy} + q_x \frac{dx}{dy}\right][w(e) - w(n)] > (1-\delta)c_y$ *for y = 0, and* $(y^*_r = 0, x^*_r = x^r_{max} < x_{max})$ *otherwise.*

The interpretation of this result is as follows. For the cases where SRO vigilance was zero and fraud maximum, the introduction of parallel regulation would *always* lead to a lower level of fraud. If parallel regulation does not turn positive the SRO's FOC, then its vigilance will remain zero, but fraud will be lower because of the direct effect of regulator vigilance on the agent's behaviour. However, parallel regulation *can* turn the SRO's FOC positive. This will imply a positive level of vigilance in which case fraud will be reduced even further. It can be shown that this condition is equivalent to the requirement that the slope of the SRO's iso-profit contour at $y = 0$ has to be steeper than the slope of the agent's reaction function. This requirement is sufficient for an equilibrium with positive vigilance and some fraud deterrence. This is illustrated in Figure 2. Starting from an equilibrium $E$ characterised by ($y^* = 0$, $x^* = x_{max}$), parallel regulation would lead to $E^1_r$ or $E^2_r$ where fraud is unambiguously lower. In conclusion, parallel regulation will necessarily imply fraud deterrence in the cases where the unregulated SRO has no incentives to be vigilant.

FIGURE 2
CHANGE OF EQUILIBRIA UNDER PARALLEL REGULATION

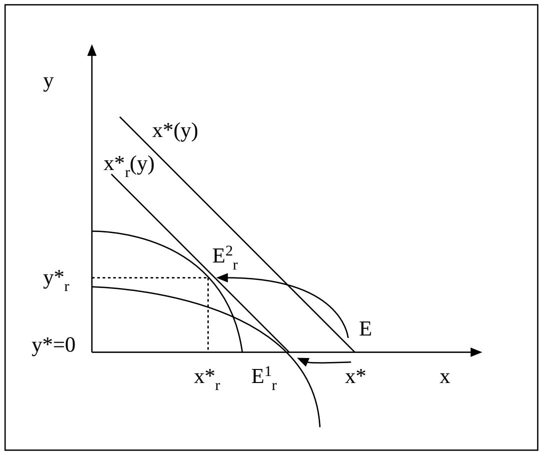


It is also interesting to examine whether parallel regulation can enhance the SRO's vigilance incentives in the cases where the unregulated SRO does undertake some vigilance voluntarily. Unfortunately, this possibility will not hold as a general case because parallel regulation provides conflicting incentives for vigilance. On the one hand, it increases the vigilance incentives because the SRO would want to reduce equilibrium fraud to diminish the probability of exposure by the regulator. This can be seen at work by inspecting the slope of the SRO's iso-profit contours, namely $\frac{dy}{dx}=-\frac{(\gamma p_x+q_x)[w(e)-w(n)]}{(1-\delta)c_y-\gamma p_y[w(e)-w(n)]}$. Clearly, the slope becomes more negative as term $q_x$ increases, which reflects the degree of responsiveness to fraud of the probability of public exposure. On the other hand, however, parallel regulation makes the agent less responsive to SRO vigilance if $q_{xx} > 0$, which follows from the slope of the agent's reaction function. This reduces the SRO's vigilance incentives. Finally, parallel regulation also decreases the equilibrium level of fraud for a given level of SRO vigilance. Because $p_{xy} > 0$, the positive effect of SRO vigilance on the probability of fraud discovery (determined by $p_y$) decreases with parallel regulation.

These conflicting effects make the total effect ambiguous. In fact, it is possible to construct equilibria where SRO vigilance can increase or decrease, and even situations where equilibrium fraud can increase.[19] However, the latter seems a rather unlikely situation. Indeed, inspecting the agent's FOC, it follows that

---

19 For example, suppose $q_{xx} = 0$, and $p_{xx}$ and $p_{xy}$ are constants, so that the slope of the agent's reaction function is unchanged. Then, it can be shown that the slope of the iso-profit contours can be increasing or decreasing in *y* (for a fixed *x*). If it is increasing in *y*, then a parallel shift of the agent's reaction function down and to the left would lead to a tangency point with higher fraud.

the optimal choice of fraud is determined by $p_x + q_x = 1/T$, where $1/T$ is a constant. By inspection of the FOC in the unregulated case, this implies that $p_x(y_r^*, x_r^*) = p_x(y^*, x^*) - q_x(x_r^*)$. Recall that $p_{xx} >$ and $p_{xy} > 0$. Therefore, lower SRO vigilance decreases the value of $p_x$, which helps re-establish equilibrium because $q_x > 0$. However, the fact that $x^*_r > x^*$ increases $p_x$. In other words, the impact of the reduction in SRO vigilance on $p_x$ has to be high enough not only to compensate for $q_x$ but also to compensate for the positive effect of increased fraud on $p_x$.[20] However, recall that when the unregulated SRO does have incentives to be vigilant (i.e. when $\frac{dp}{dy} < 0$), SRO vigilance is increasing in $\gamma$. Therefore, although this form of parallel regulation does not increase the total probability of fraud discovery, it would be sufficient to increase vigilance and reduce fraud further.

## 5. Conclusions

If consumers and the SRO interact repeatedly, positive vigilance, and consequently fraud deterrence may indeed happen in equilibrium. However, there are five obstacles for this to happen. The first two are the customary conditions for co-operation in repeated games, namely, that the SRO has to exhibit a sufficient concern for future payoffs to resist temptation to cheat, and secondly, that vigilance costs must be sufficiently smaller than the difference between high and low profits. However, this work unveils three other additional obstacles for positive SRO vigilance to occur. These conditions emerge from the assumption of non-observation of SRO vigilance or fraud levels by consumers that characterise the industries where self-regulation is most commonly found. First, since the SRO has no incentive to expose any fraud by itself, then positive SRO vigilance requires other mechanisms to make exposure happen, such as involuntary leakage of evidence of fraud to the public. Moreover, this must happen with a sufficiently high probability. Second, the probability of exposure has to be "sufficiently" decreasing in SRO vigilance. This may not be the case because fraud is endogenously determined in a principal-agent context, where higher vigilance has a positive impact on the probability of exposure, which may dominate over the negative effect of the agent's fraud choice. Third, the fulfillment of the conditions above does not guarantee positive vigilance because there is still the issue of multiplicity of equilibria and equilibrium selection because zero vigilance and maximum fraud remains as a possible equilibrium of the repeated game.

---

[20] The case of SRO and regulator exposure as non-exclusive events is more complex, and it is not addressed here. In particular, two complications emerge in this case. First, the slope of the agent's reaction function is not necessarily negative under parallel regulation. Second, optimal fraud may not be smaller for a given level of SRO vigilance under parallel regulation. These facts make it difficult to derive general results.

The nature of equilibrium in this game is substantially different from other repeated games. While in the repeated games with observable action punishment is not observed in equilibrium because there is no need to carry it out, in this model the exercise of punishment is observed in equilibrium because it is precisely its existence what enforces co-operation. However, it must be clarified that punishment by consumers does not exist because of actual deviation by the SRO. Instead, it would happen even though consumers would be *certain* that fraud exposure resulted from "bad luck" and not from actual deviation.

In conclusion, the conditions for positive vigilance and fraud deterrence described earlier impose further constraints for co-operation than the ones often found in repeated games. They suggest that co-operation may be less likely to emerge, or if it does, equilibrium vigilance, although positive, may not be high enough to make self-regulation an effective way to cope with fraud and wrongdoing in these industries. This raises the questions of whether and how parallel regulation can improve matters. This work distinguishes between two forms of parallel regulation. First, parallel regulation can investigate the SRO vigilance effort and the evolution of the cases of fraud discovered by the SRO. Therefore this form of regulation aims at informing the public about the activities of the SRO. Alternatively, it can investigate the existence of fraud *independently* of the SRO. The main conclusions in this respect are the following. First, under the assumption that leak exposure and regulator exposure are exclusive events, parallel regulation will always be beneficial in terms of reducing fraud if the starting situation involves no vigilance and maximum fraud, even if optimal SRO vigilance remains zero. In addition, equilibrium SRO vigilance may turn positive, which would reduce fraud even further. This would happen because the SRO would wish to reduce fraud to avoid exposure by the regulator. Second, if the unregulated case involves some positive vigilance and fraud deterrence, parallel regulation based on independent enforcement, is likely to increase vigilance and reduce fraud. However, in this case, parallel regulation in the form of investigation of the SRO's enforcement actions, is sufficient to increase vigilance and reduce fraud further.

In conclusion, this work reinforces the central messages of previous works about Self regulation, namely Núñez (2000 A, B), in which SROs also exhibit scant incentives to enforce quality and expose malpractice to the public. This would suggest that the fact that SR is found in practice in many industries need not imply that it actually works properly. On the contrary, there seem to be several sound theoretical arguments to claim that in fact it may not, although it would be difficult to assess that claim empirically due to the very nature of SR. Finally, this work also reinforces the general message in Núñez (2000 A, B) that SR can, however, contribute significantly to fight malpractice and fraud, if it is properly accompanied by public parallel regulation of product quality.

## References

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, 309-325.

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

Emons,W. (1997). "Credence Goods and Fraudulent Experts", *RAND Journal of Economics*, Vol. 28, Nº 1, 107-119.

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

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

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

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

Núñez, J.(2000 C). "La Autorregulación como Concepto Regulatorio" Departamento de Economía, Universidad de Chile.

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

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