Reputations, Investigations and Self Regulation

Peter M. DeMarzo[*], Michael J. Fishman[**], and Kathleen M. Hagerty[***]

September 2007

Preliminary and Incomplete

Abstract

Consider a principal who hires an agent, e.g., a broker is hired to fill stock orders. The outcome may be good or bad, e.g., the orders may be filled at favorable or unfavorable prices. A bad outcome may simply reflect bad luck or it may result from the agent defrauding the principal. How is the agent's incentive to defraud the principal restrained? With a dynamic costly-state-verification model we examine two possibilities: (i) reputational penalties, meaning that principals can choose not to transact with an agent who has a bad history even if there is no evidence of fraud; and (ii) investigations, meaning that at a cost, an agent's performance can be directly checked for fraud. Reputational penalties have no direct cost but they are imprecise in the sense that an agent can be penalized even if it was just bad luck. Investigations are costly but precise in the sense that an agent is penalized only if he is found to have committed fraud. We examine how principals would balance the two ways of controlling an agent's incentives. Then we examine the incentives of agents to form a self-regulatory organization to conduct its own investigations of agent performance and to penalize agents who are caught cheating. We show that the threat of SRO investigations can enhance agents' welfare but reduce the welfare of principals. In effect, SRO investigations can inefficiently – from the principal's perspective – crowd out reputational penalties.

DeMarzo and Fishman would like to thank the National Science Foundation for financial support.

[*] Graduate School of Business, Stanford University, 518 Memorial Way, Stanford CA 94305, demarzo_peter@gsb.stanford.edu
[**]Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston IL 60208, m-fishman@kellogg.northwestern.edu
[***] Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston IL 60208, k-hagerty@kellogg.northwestern.edu

1. Introduction

Why don't agents always cheat principals? One reason is that an agent who cheats might be caught and penalized. A second reason is that the agent's reputation may suffer. That is, an agent who delivers bad outcomes to principals can develop a bad reputation even if he is never actually caught cheating. The bad reputation will deter future principals from hiring him. With a dynamic costly-state-verification model we examine the tradeoff between these ways of deterring cheating.[1]

Reputational penalties have no direct cost but they are imprecise in the sense that an agent can be penalized even if it was just bad luck. Investigations are costly but precise in the sense that an agent is penalized only if he is found to have cheated. We examine how principals would balance the two ways of controlling an agent's incentives. Then we examine the incentives of agents to form a self-regulatory organization (SRO) to conduct its own investigations of agent performance and to penalize agents who are caught cheating. We show that the threat of SRO investigations can enhance agents' welfare but reduce the welfare of principals. In effect, SRO investigations can inefficiently – from the principal's perspective – crowd out reputational penalties.

Our interest in SROs comes from the securities industry where examples of SROs include the Financial Industry Regulatory Authority (which was formed in 2007 by combining the SRO operations of the NASD and the NYSE Group) and the Chicago Board Options Exchange. These SROs monitor their members, policing for fraud and other violations, imposing sanctions when

---

[1] A third reason why agents don't always cheat principals is that most agents are honest. Though honesty is obviously important in practice, we do not focus on that explanation here.

such behavior is identified.[2] An often-raised question is whether we should expect SRO enforcement efforts to adequately serve the needs of the investing public; see, for example, McCaffrey and Hart (1998). We examine this issue under the assumption that an SRO's objective is to maximize the utility of its members.[3]

Why are SROs doing enforcement? One possibility is that specialized securities industry organizations can police securities markets at lower cost than can the Securities and Exchange Commission (SEC) or the state securities regulators. This case was considered in the DeMarzo, Fishman, and Hagerty (2005) analysis of self regulation with government oversight. That analysis used a one-period costly-state-verification model; if the agent was caught cheating, a monetary sanction was imposed. There we found that an SRO has the incentive to be more lax in its enforcement efforts than would be preferred by principals. While an SRO would choose to impose the maximum penalty for fraud, it would not investigate agents with as high a frequency as principals would prefer.

In the current paper, we examine a dynamic version of the DeMarzo, Fishman, and Hagerty (2005) model. A sequence of principals wants to hire an agent to perform a transaction. A principal's choice of which agent to hire depends on agents' performance histories. In addition to monetary penalties, principals can penalize agents by taking their business elsewhere if agents have bad performance, even without direct evidence of fraud. In effect, the market imposes

---

[2] See National Association of Securities Dealers (1996), Phillips (1997) and McCaffrey and Hart (1998) for discussions of SRO enforcement and Frankhauser, et al. (1997) for details of SRO enforcement procedures. SROs also perform enforcement tasks that are not directed at their membership, e.g., the exchanges monitor their listed companies to ensure that they satisfy listing requirements.

[3] See Shuchman (1981) and Ameringer (1999) for a discussion of self regulation in accounting, law, and medicine.

discipline on agents. Since agents earn positive expected rents per transaction, total penalties can include the discounted value of the anticipated future rents.

Whether an SRO's enforcement is too lax is a more subtle issue in this dynamic setting. In some cases, an SRO's investigation policy may be so lax that it relinquishes all agent discipline to the market through reputation. In other cases, SRO enforcement might be too aggressive in that investigations may be more frequent than principals would prefer; but at the same time too lax in that penalties for fraud may be lower than principals would prefer. In such a case, the SRO's investigation policy may be so aggressive that it pre-empts any market discipline through reputation.

LITERATURE

The basic dynamic model with reputation is described in Section 2 and analyzed in Section 3. In Section 4 we add investigations to the model and analyze the optimal tradeoff between the use of incentives based on reputations vs. investigations. We present the optimal enforcement policy from the principal's perspective. Then in Section 5 we analyze the optimal enforcement policy from an SRO's perspective under the assumption that the SRO seeks to maximize agents' utility. Section 6 contains concluding remarks. Proofs are in the Appendix.

2. The Basic Model

We have a dynamic principal-agent model. It is an infinite-horizon model with an infinite number of identical risk-neutral agents. Agents have zero initial wealth and limited liability. Further, agents are assumed not to save over time; agents consume any cash flow in the period it is earned. These assumptions imply that agents will be unable to compete away their rents by paying principals to do business with them. Agents discount expected consumption streams according to the discount factor $\delta < 1$. An agent's alternative to being hired in period t is a period t payoff of 0.

There is an infinite sequence of risk-neutral principals, indexed by $t \in \{1, 2, \ldots\}$. In period t, principal t hires an agent for the purpose of generating the risky cash flow W_t, a binary random variable with support $\{w_1, w_2\}$, where $w_2 > w_1 \geq 0$. Let $\pi_i \equiv \text{Prob}(w_i)$. The principal's best alternative to hiring an agent entails a payoff of $\underline{\alpha}$, where $w_1 < \underline{\alpha} \leq E[W_t]$. The random variables $\{W_t\}$ are independent and identically distributed.

The agent hired by principal t privately observes the realization of the cash flow W_t and reports it to the principal. This is the source of the agency problem. The agent can misreport the cash flow and divert the remainder to himself. Specifically, if the realization of W_t is w_2, the agent can report a realization of w_1 and keep $w_2 - w_1$ for himself. If the realization of W_t is w_1, then there is no scope for misrepresentation. In the context of the brokerage industry, this model corresponds to several possibilities. Consider a customer who hires a broker to sell securities that have been trading in the range $[w_1, w_2]$. Realizations of w_2 and w_1 correspond to the sell orders being filled at high and low prices, respectively. Suppose the agent could fill the order at the more favorable price w_2. The agent could cheat the customer by allowing an associate (or perhaps himself for that matter) to take the other side of the sell order at a price of w_1 and then resell the securities at w_2 for an illicit gain of $w_2 - w_1$. The reverse strategy is of concern if the customer is buying. Alternatively, a customer may give a broker authority to trade in the customer's account on the customer's behalf. In this case, the random variable W_t corresponds to the value of the customer's portfolio. With a favorable realization w_2, the broker might churn the account and rack up excessive brokerage fees equal to $w_2 - w_1$. For a third example, if a customer's portfolio is worth w_2, a broker might simply steal $w_2 - w_1$ from the account.[4]

Let $r_t \in \{w_1,w_2\}$ denote the agent's report of the cash flow. A contract between the agent and principal t is given by the function $z(r_t)$ specifying the payment from the agent to the principal given an agent report of r_t. Reflecting the perspective of competitive agents (as in the brokerage industry), the contract is chosen by the principal to maximize his own utility. Each principal transacts once. Consequently these are short-term contracts – principal t's contract depends only on r_t.[5]

Principals can rely on agent reputation in choosing which agent to hire. Reputation considerations can induce agent honesty because a principal can choose not to hire an agent who has reported low cash flow realizations in the past. Hence agents may behave honestly today to avoid paying a reputation penalty in the future. Specifically, principals observe the experience of prior principals – who they hired and the outcome of the transaction. Based on this information, principal t can hire an agent who was hired by an earlier principal $s < t$ or can hire an agent who has not yet been hired.

Reputation penalties can be imposed in a large variety of ways. For instance, principals could follow a rule of never hiring an agent who has reported a low cash flow more than m times; or m times in a row or more than m% of the time; etc. Or an agent's performance over the last m periods could trigger the next n principals to not hire him, and so on. Further, these reputation penalties can be applied stochastically. Without loss of generality, however, we can restrict attention to the case in which agents receive either (i) no reputation penalty and are hired

---

[4] McCaffrey and Hart (1998) discuss the nature of securities industry enforcement cases.

[5] See DeMarzo and Fishman (2006) for an analysis of long-term contracts in a model with this sort of agency problem.

again by the next principal; or (ii) they are never hired again by any subsequent principal. That is, based solely on the outcome of transaction $t - 1$, principal t either hires the agent from transaction $t - 1$ or hires an agent that has not been hired before. To see that this is without loss of generality, note that an agent's payoff in a given period consists of current income plus a continuation payoff. The continuation payoff depends on future expected reputation penalties. Let $V(\alpha)$ denote the agent's highest possible continuation payoff (as of the beginning of the next period) given an expected principal payoff of at least α each period. The lowest possible continuation payoff for an agent is 0, which corresponds to never being hired again. Thus, after any history, the agent must receive a continuation payoff between $V(\alpha)$ and 0, which can always be implemented by randomizing between no reputation penalty and never being hired again. Moreover, since the imposition of reputation penalties in future periods may distort the agent's incentives in these future periods, randomizing between no reputation penalty and never being hired again is Pareto optimal.

So for an agent who has handled many consecutive transactions, only his performance on transaction $t - 1$ matters for whether he is hired by principal t. Suppose the agent hired by principal $t - 1$ reports r_{t-1}. Let $\theta(r_{t-1})$ denote the *reputation penalty*, the probability that this agent is not hired by principal t, which in turn implies that he is not hired by any subsequent principal either. With probability $1 - \theta(r_{t-1})$ this agent is hired by principal t.

Since agents are identical and don't save over time (by assumption) all agents are equally desirable from the perspective of principal t. Consequently an individual principal's choice of whether to hire the previously hired agent or choose a new agent affects the incentives of the agent when dealing with the *previous* principal. That is, it does not benefit the current principal directly since all agents are identical. Moreover principal t is willing to follow any proposed reputation penalty $\theta(r_{t-1})$. The choice of $\theta(\cdot)$ is therefore best interpreted as a "social norm" that specifies the extent to which principals use reputation when selecting an agent. We will consider the choice of $\theta(\cdot)$ that maximizes principals' utility. It is also the case that in any period, for a given $\theta(\cdot)$, the optimal contract for the principal is independent of the agent's history.

Since the principal's reservation payoff $\underline{\alpha}$ satisfies $w_1 < \underline{\alpha} < w_2$, we can immediately apply the revelation principle. That is, since $z(r(w_1)) \leq w_1$ because of agent limited liability, $z(r(w_2)) > w_1$ is required to induce the principal to do business with an agent (rather than take the reservation payoff $\underline{\alpha}$). Since the principal receives different payments in the two states we can associate truthful reports with the two states, i.e., $r(w_i) = w_i$. So in what follows we restrict attention (without loss of generality) to choices of (z, θ) that induce truthful agent reports. Note that inducing truthful reporting is only an issue for cash flow realizations of w_2; since a contract will entail $z(w_2) > w_1$, it is not feasible for an agent to fraudulently report a cash flow of w_2.

It is immediate that $\theta(w_2) = 0$ is optimal. There is no reason to impose a reputation penalty on an agent who reports the good outcome. With these results we can use the shorthand notation $z_i \equiv z(w_i)$ and $\theta \equiv \theta(w_1)$.[6] We now analyze the principal's optimization problem.

[6] With regard to the reputation threat, there are multiple equilibria. In one, principals ignore past outcomes and choose the agent who has handled the most transactions. In equilibrium this leads to $\theta(w_1) = 0$. In this equilibrium, switching agents is suboptimal for a given principal because the new agent would expect to have no future transactions (regardless of his report) and so would face no reputation threat. Another equilibrium is for principals to use the agent who has handled the most transactions without reporting a low outcome. This implements $\theta(w_1) = 1$. Following this rule would be optimal for each principal since an agent who reported a low outcome in the past expects no future principals. Presumably, evolutionary forces will lead to the equilibrium that is best for principals. We do not model equilibrium selection here; we simply identify which equilibrium is best for principals.

3. Reputation

The reputation penalty θ is chosen to maximize the expected payoff of principals. Taking the reputation penalty θ as given, each individual principal chooses a contract z to maximize his own expected payoff $E[z]$. Taking θ and z as given, the agent observes the cash flow realization and chooses a report to maximize his own expected payoff. Given that the problem is stationary, the reputation penalty will not change over time, each principal will offer the same contract, and each agent will follow the same reporting strategy.

Let V denote the agent's continuation payoff provided that he is hired by principal $t + 1$. Taking θ and V as given, principal t solves the following problem (suppressing the subscript t):

$$\mathrm{P}(\theta) = \max_z \pi_1 z_1 + \pi_2 z_2$$

$$\text{subject to (F)} \qquad z_i \leq w_i \tag{1}$$

$$\text{(IC)} \quad w_2 - z_2 + \delta V \geq w_2 - z_1 + \delta(1-\theta)V$$

Principals maximize the expected payment received from the agent. The feasibility constraint (F) reflects the agent's limited liability; he cannot pay the principal more than the cash flow. The incentive-compatibility constraint (IC) ensures that the agent truthfully reports a realization of w_2. The left side of (IC) is the agent's payoff if he reports w_2; there is a cash component, $w_2 - z_2$, plus the discounted continuation payoff, δV. The right side of (IC) is the agent's payoff if he reports w_1 when the actual realization is w_2; the cash component is $w_2 - z_1$ and the expected continuation payoff is $\delta(1 - \theta)V$. This reflects the reduced probability, $1 - \theta$, of being hired by the next principal. Of course if the maximized objective is less than principals' reservation payoff $\underline{\alpha}$, then principals will choose not to transact.

Anticipating the contract that principals will offer to agents, the optimal reputation penalty for principals solves:

$$\max_{\theta} \; P(\theta) \tag{2}$$

The principals' solution to (1) is straightforward:

$$z_1 = w_1 \quad \text{and} \quad z_2 = w_1 + \delta\theta V \tag{3}$$

The agent payment z_1 is set as high as possible subject to (F); this raises the objective and relaxes (IC). The agent payment z_2 is set as high as possible subject to (IC). The feasibility constraint does not bind for z_2.

The agent's continuation payoff V is given by

$$V = \pi_2[w_2 - z_2 + \delta V] + \pi_1[w_1 - z_1 + \delta(1-\theta)V] \tag{4}$$

which can be rewritten as

$$V = \frac{\pi_2[w_2 - z_2] + \pi_1[w_1 - z_1]}{1 - \delta\pi_2 - \delta\pi_1(1-\theta)}. \tag{5}$$

Maximizing $P(\theta_1)$ is equivalent to maximizing z_2 (z_1 does not depend on θ_1). Using (3) and (5) we have

$$z_2 = w_1 + \delta\theta \frac{\pi_2[w_2 - w_1]}{1 - \delta(1-\theta)} \tag{6}$$

which is increasing in θ. Hence $\theta = 1$ maximizes the welfare of principals. It is a severe reputation penalty. If an agent reports a low outcome then he is not hired by subsequent principals. We summarize these results:

PROPOSITION 1. The optimal contract and reputation penalty for principals – the solution to (1) and (2) – entails the following:

(i) a contract $z_1 = w_1$ and $z_2 = w_1 + \delta\pi_2(w_2 - w_1)$

(ii) a reputation penalty $\theta = 1$

Each principal's expected payoff equals

$$\alpha = E[W] - \pi_2(1-\delta\pi_2)(w_2 - w_1) \tag{7}$$

and the expected payoff of an agent who is hired in a given period (and may be hired in subsequent periods) equals

$$V = \pi_2(w_2 - w_1)\,. \tag{8}$$

From (7) and (8) we can see that holding $E[W]$ and π_i fixed, an increase in $w_2 - w_1$ (a mean-preserving spread) decreases the expected payoff of principals and increases the expected payoff of agents. This is because a mean-preserving spread aggravates the incentive problem and generates a higher rent for the agent. A higher discount factor, δ, increases the importance of an agent's reputation and hence improves incentives, raising principals' expected payoffs.

4. Reputations and Investigations

We now consider the possibility that the agent's performance can be investigated. We use a costly-state-verification model, adapted from Townsend (1979), Border and Sobel (1987), and Mookherjee and Png (1989). At a cost $c \geq 0$, an investigation reveals the realization of W. So if an agent reports w_1, an investigation will determine whether the agent is reporting truthfully.

Let p denote the probability that an agent is investigated if he reports a realization of w_1. There will be no reason to investigate an agent who reports w_2 (clearly he is reporting truthfully). Suppose the agent reports w_1, he is investigated, and the investigation finds that the realization is actually w_2. For this case, let x denote the monetary penalty imposed on the agent. We could also consider a penalty involving a suspension though this would be redundant given the possibility of a reputation penalty. That is, principals would optimally choose to never again hire an agent who has been investigated and found to have misreported. We do not consider other non-monetary penalties, e.g., incarceration. There will be no reason to impose any monetary or reputation penalty on an agent if he reports w_1, is investigated, and the investigation finds that the realization actually is w_1. As in the reputation-only case, the reputation penalty θ and the investigation/penalty schedule (p, x) are chosen to maximize principals' expected utility.

The addition of investigations/penalties leads to the following optimization problem for a principal. Taking (θ, p, x) and V as given, a principal solves:

$$\mathrm{P_I}(\theta, p, x) = \max_z \pi_1 z_1 + \pi_2 z_2 - \pi_1 pc$$

$$\text{subject to} \quad \text{(F)} \quad z_i \leq w_i \tag{9}$$

$$\text{(IC-I)} \quad w_2 - z_2 + \delta V \geq (1-p)[w_2 - z_1 + \delta(1-\theta)V] + p[w_2 - z_1 - \min\{x, w_2 - z_1\}]$$

The principal's objective now includes the expected investigation cost. Since we can still restrict attention to agent truth-telling, the principal collects no penalty x in equilibrium. The feasibility constraint is unchanged. The left side of the incentive-compatibility constraint (IC-I) is unchanged but the right side now reflects the possibility of investigation. Suppose the agent falsely reports w_1 (the right side of (IC-I)). With probability $1 - p$ there is no investigation and (as in the reputation-only case) with probability θ the agent is not hired again. With probability p the agent is investigated and the fraud is identified. The agent pays a monetary penalty, which equals $\min\{x, w_2 - z_1\}$ because of limited liability, and the agent is not hired again.

The optimal reputation penalty and investigation/penalty schedule is the solution to:

$$\max_{\theta, p, x} \mathrm{P_I}(\theta, p, x) \tag{10}$$

Like before, a principal's solution to (9) sets the payment z_1 as high as possible subject to (F) and sets the payment z_2 as high as possible subject to (IC-I). That is

$$z_1 = w_1 \text{ and } z_2 = w_1 + p\min\{x, w_2 - w_1\} + [1-(1-p)(1-\theta)]\delta V \tag{11}$$

The agent's continuation payoff V is given by

$$V = \pi_2[w_2 - z_2 + \delta V] + \pi_1[w_1 - z_1 + p\delta V + (1-p)(1-\theta)\delta V] \tag{12}$$

Substituting in for z_1 and z_2 and rearranging, (12) can be rewritten as

$$V = \frac{\pi_2[w_2 - w_1 - p\min\{x, w_2 - w_1\}]}{1 - \delta[p(1-\pi_2) + (1-p)(1-\theta)]} \tag{13}$$

Now consider the choice of θ, p, x that solves (10). Maximal monetary penalties for fraud are optimal for principals:

$$x = w_2 - w_1.$$

No penalty is collected in equilibrium, but a maximal penalty relaxes the incentive-compatibility constraint.

It remains to determine θ and p. Reputation penalties (θ) have no direct cost (since an agent would be replaced with an identical agent) but are imprecise in the sense that an agent can be penalized even if he honestly reports a low outcome. Investigations (p) have a direct cost but lead to penalties that are correctly assessed in the sense that an agent is only penalized if there is direct evidence of fraud. In other words, while agents can pay reputation penalties in equilibrium (even though they didn't cheat), investigations involve penalties that remain out-of-equilibrium threats. The imprecision of reputation penalties carries an indirect cost in that it reduces the value of maintaining a good reputation. See (12). Holding z_2 fixed, $\partial V / \partial\theta \leq 0$ with strict inequality for $p < 1$; and $\partial V / \partial p \geq 0$ with strict inequality for $\theta > 0$. That is, aside from the direct effect on z_2, a greater reputation penalty makes it more likely that an agent is penalized for truthfully reporting a low outcome and this reduces the agent's payoff V. This reduces the agent's incentive to maintain a good reputation. And aside from the direct effect on z_2, a greater investigation probability makes it less likely that an agent is penalized via reputation for truthfully reporting a low outcome and this increases the agent's payoff V. This increases the agent's incentive to maintain a good reputation. As stated above, investigation threats are precise but directly costly, whereas reputation threats are imprecise and thus indirectly costly. Of interest is the optimal tradeoff between the two.

The optimal policy is characterized by distinct regimes in which one or both of these instruments is at a boundary. In an "investigation" regime, $\theta = 0$ and $p > 0$; an agent is only penalized if caught cheating. In a "performance-based" regime, $p = 0$ and $\theta = 1$; an agent will not be investigated but will be penalized (meaning not hired again) simply for reporting a low outcome. In a "zero-tolerance" regime, $\theta = 1$ and $p_1 > 0$; an agent who reports a low outcome is not rehired unless he is investigated and cleared of fraud. Including our results on z, x, and p, we have:

PROPOSITION 2. The optimal contract, reputation penalty, and investigation/penalty schedule for principals – the solution to (9) and (10) – entails a:

(i) contract $z_1 = w_1$ and $z_2 = w_1 + p(w_2 - w_1) + [1 - (1 - p)(1 - \theta)]\delta V$.

(ii) monetary penalty $x = w_2 - w_1$.

(iii) reputation penalty and investigation schedule that satisfy either $\theta = 0$ and $p \geq (1 - \delta)/\delta\pi_1$; or $\theta = 1$ and $p = 0$; or $\theta = 1$ and $0 < p < (1 - \delta)/\delta\pi_1$.

The principal's expected payoff equals

$$\alpha = E[W] - \pi_2(w_2 - w_1)\frac{(1-p)[1-\delta(1-\pi_1(1-p)\theta)]}{1-\delta[p\pi_1 + (1-p)(1-\theta)]} - \pi_1 pc \qquad (14)$$

and the expected payoff of an agent who is hired in a given period (and may be hired in subsequent periods) equals

$$V = \frac{\pi_2(w_2 - w_1)(1-p)}{1-\delta[p\pi_1 + (1-p)(1-\theta_1)]}. \qquad (15)$$

The optimal reputation penalty is at a corner, $\theta = 0$ (no reputation penalty) or $\theta = 1$ (the most severe reputation penalty). Notably, if the investigation threat p is above a threshold, the cost of imposing (in equilibrium, undeserved) reputation penalties on agents outweighs the benefit and it is optimal to employ no reputation threat.

For example, suppose W is uniform on {100, 200}, the discount factor $\delta = .9$, and the investigation cost $c = 100$. Figure 1 plots the agent's and principal's expected payoffs for all combinations of p and θ. The investigation regime, $p > 0$ and $\theta = 0$, corresponds to the parabola-shaped segment with p increasing

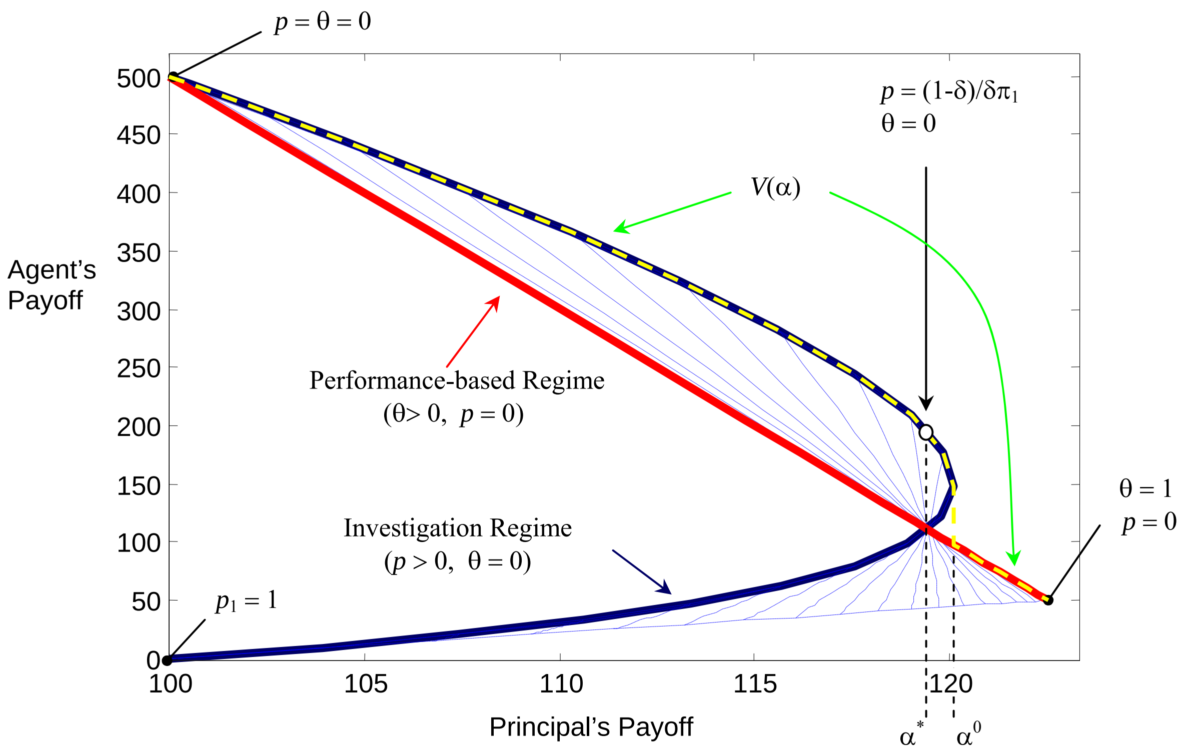

Figure 1. Payoff Possibilities for all (p, θ)

clockwise along it. Within the investigation regime, the optimal p for principals is interior with $p \geq (1 - \delta)/\delta\pi_1$. If this were a static example, with risk neutrality the principal's payoff would be monotone in p. This is not true in the dynamic case since the presence of future agent rents magnifies incentives. These rents are decreasing in p. So while initial increases in p help the principal, further increases lower the rents and diminish incentives.

The performance-based regime, $p = 0$ and $\theta > 0$, corresponds to the linear segment with θ increasing to the right along it. With no investigations, $\theta = 1$ is optimal for principals. This segment also corresponds to the case examined in Section 3 in which there were only reputation penalties and the result matches that from **PROPOSITION 1**.

The shaded region corresponds to interior (suboptimal) policies. In this example, the zero-tolerance regime is never optimal, which is true for all but the most extreme parameter values we considered.

The Pareto frontier $V(\alpha)$ is discontinuous, lying along the parabola for $\alpha \leq \alpha^0$ and jumping to the linear segment for $\alpha > \alpha^0$. As the investigation cost c declines, the parabola and thus α^0 shifts out while the linear segment is unchanged. So a small improvement in investigation technology can lead to a switch from the performance-based regime to the investigation regime.

5. Self Regulation

We considered the optimal investigation/penalty schedule (p, x) from the principals' perspective – and the solution was a Pareto optimum. This is a natural outcome if enforcement is performed by a government agency. But suppose enforcement is controlled by agents of the agents, for example, a self-regulatory organization (SRO). What investigation/penalty schedule would an SRO choose if its objective is to maximize the utility of the agents (rather than the principals)?

As in Section 4, here we assume that the principal pays the expected investigation cost, $\pi_1 pc$. This corresponds to the SRO charging the principal a fee to transact and this fee covers the SRO's enforcement cost. We will also discuss the case in which the agents pay the enforcement cost.

Our model of self regulation has the following timing and structure:

1. The SRO chooses the investigation/penalty schedule (p, x). The SRO's objective is to maximize the expected payoff of agents, V, subject to the principals' participation constraint, $\alpha = \pi_2 z_2 + \pi_1 z_1 - \pi_1 pc \geq \underline{\alpha}$.[7]
2. Taking (p, x) as given, the reputation penalty θ is chosen to maximize the expected payoff of principals, α.
3. Taking (p, x, θ) and V as given, principals choose the contract z to maximize their own expected payoff.

In maximizing agent expected utility the SRO anticipates the optimal response with regard to the reputation threat and contract. Like before, the revelation principle applies and we can restrict attention to truthful agent reporting strategies. We solve the self regulation model by starting with the third stage and working back.

A principal's optimization problem is unchanged; his choice of z is again determined as the solution to (9). Consequently principals again choose payments z according to (11) and the agent's expected payoff V again satisfies (13).

The reputation threat θ is determined as the solution to

---

[7] As in the analysis of Section 4, the SRO has no incentive to impose a penalty if the agent reports w_1, is investigated, and is found to have reported truthfully. If it were feasible, the SRO might choose to investigate and penalize agent reports of w_2. By dulling agent incentives to report w_2, this could potentially force principals to offer contracts with a lower payment z_2.

$$\max_{\theta} \ P_I(\theta, p, x) \tag{16}$$

That is, taking (p, x) as given, θ is chosen to maximize the principals' expected payoff. The choice of θ affects z_2 but not z_1 (see (11)) so maximizing the principals' expected payoff is equivalent to maximizing z_2. This leads to the following result regarding the reputation threat.

PROPOSITION 3. The optimal reputation threat – the solution to (16) – satisfies:

$$\theta = \begin{cases} 0 & \text{if} \quad p \geq \dfrac{1-\delta}{\delta\pi_1} \\ 1 & \text{if} \quad p < \dfrac{1-\delta}{\delta\pi_1} \end{cases} \tag{17}$$

The reputation threat depends on the probability of an investigation p but not the monetary penalty x. This is because there is no interaction between the monetary and reputation penalties; the monetary penalty x is only relevant in the event of an investigation while the reputation penalty θ is only relevant in the absence of an investigation. For the same reason as in the analysis of Section 4, the reputation threat is at a corner. If the SRO sets the investigation probability above a threshold, principals employ no reputation threat. If, however, the SRO sets the investigation probability below the threshold, principals employ the most severe reputation threat; if an agent reports w_1 and an investigation does not clear him of fraud, the agent will not be hired by the next principal (or any subsequent principal).

Anticipating how the reputation threat θ and the payments z will be chosen, how will the SRO choose the investigation/penalty schedule (p, x)? The SRO solves the following problem:

$$V(\underline{\alpha}) = \max_{p,x,\theta} \frac{\pi_2[w_2 - w_1 - p\min\{x, w_2 - w_1\}]}{1-\delta[p(1-\pi_2)+(1-p)(1-\theta)]}$$

$$\text{subject to} \quad (\text{R}) \quad \theta = \begin{cases} 0 & \text{if } p \geq \dfrac{1-\delta}{\delta\pi_1} \\ 1 & \text{if } p < \dfrac{1-\delta}{\delta\pi_1} \end{cases} \tag{18}$$

$$(\text{PC}) \quad E[W] - [1-\delta+\delta\pi_1\theta(1-p)]V(\underline{\alpha}) - \pi_1 pc \geq \underline{\alpha}$$

The SRO maximizes the expected payoff of the agent hired by principal *t*. The constraint (R) reflects the fact that the reputation penalty θ will be chosen to maximize the principals' expected payoff. The principal's participation constraint is (PC), where the contract *z*, which is chosen by the principal, is given by (11).

PROPOSITION 4. The optimal investigation/penalty schedule for the SRO and the resulting reputation penalty – the solution to **(18)** – entails:

(i) $p > (1 - \delta)/\delta\pi_1$, $x = w_2 - w_1$, and $\theta = 0$ and (PC) binds; or

(ii) $p = (1 - \delta)/\delta\pi_1$, $x \leq w_2 - w_1$, $\theta = 0$, and if $x > 0$ then (PC) binds; or

(iii) $p < (1 - \delta)/\delta\pi_1$, $x \leq w_2 - w_1$, $\theta = 1$, and if $p > 0$ and $x > 0$ then (PC) binds.

In cases (i) and (ii) of **PROPOSITION 4**, the SRO investigates aggressively enough to preempt any use of reputation penalties. In these cases, principals might do better with no investigations, $p = 0$. Refer back to the example in Figure 1. The dashed lines emanating from the curve indicate the payoffs resulting from an increase in θ. So with $p > (1 - \delta)/\delta\pi_1$, the dashed lines have a positive slope, meaning that an increase in θ decreases both the principal and agent payoff; and with $p = (1 - \delta)/\delta\pi_1$, the dashed line is vertical, meaning that an increase in θ decreases the agent payoff but leaves the principal payoff unchanged. Hence $p \geq (1 - \delta)/\delta\pi_1$ preempts the use of reputation penalties. Now suppose principals' reservation utility satisfies $\alpha^* \leq \underline{\alpha} \leq \alpha^0$. With $p \geq (1 - \delta)/\delta\pi_1$ the SRO satisfies a principal's participation constraint, keeping the solution on the far right side of the curve. This is optimal for agents. Principals would do better, however, with no investigations and just the use of reputation penalties, $p = 0$ and $\theta = 1$.

In case (ii), the SRO might choose a less than maximal monetary penalty, $x < w_2 - w_1$. This differs from typical enforcement analyses. Since such penalties are not incurred in equilibrium, the typical result is that these penalties are maximized. This relaxes the incentive-compatibility constraint without directly entering equilibrium payoffs. The difference is that here, the investigation/penalty schedule (p, x) and the contract z are chosen by different parties, the SRO and the principal, respectively. So for the SRO to raise agent utility at the expense of principal utility, the SRO cannot directly lower z; instead it will do so indirectly by lowering x. A case (ii) solution involving $x < w_2 - w_1$ would not be Pareto optimal. An alternative solution with $x = w_2 - w_1$ and a lower value of p would be preferred by both agents and principals. But such a solution would not be forthcoming in a setting where the principal chooses the contract and the SRO chooses the investigation/penalty schedule.

6. Concluding Remarks

7. Appendix

References

Ameringer, C. F., State Medical Boards and the Politics of Public Protection. Baltimore: Johns Hopkins University Press, 1999.

Border K. and Sobel, J., "Samurai Accountant: A Theory of Auditing and Plunder," Review of Economic Studies, 1987, 54, 525-540.

DeMarzo, P. M. and Fishman, M. J., "Optimal Long-Term Financial Contracting with Privately Observed Cash Flows," forthcoming Review of Financial Studies, 2006.

DeMarzo, P. M. and Fishman, M. J., and Hagerty, K. M., "Self Regulation and Government Oversight," Review of Economic Studies, 72, 2005, 687-706.

Frankhauser, M. M., Gardner, L. M., McNally, B. F., and Leatherwood, L. A., "Enforcement by Self-Regulatory Organizations: A Growing Impact on Broker-Dealers and their Personnel," in The Securities Enforcement Manual: Tactics and Strategies, ed. by R. M. Phillips. Chicago: American Bar Association, 1997.

Gehrig, T. and Jost, P-J., "Quacks, Lemons, and Self Regulation: A Welfare Analysis," Journal of Regulatory Economics, 1995, 7, 309-325.

Jarrell, G. A., "Change at the Exchange: The Causes and Effects of Deregulation," Journal of Law and Economics, 1984, 27, 273-312.

Kondo, J. E., "Self-Regulation and Enforcement in Financial Markets: Evidence from Investor-Broker Disputes at the NASD," MIT working paper, 2006.

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

Mahoney, P. G., "The Exchange as Regulator," Virginia Law Review, 1997, 83, 1453-1500.

McCaffrey, D. P. and Hart, D. W., Wall Street Polices Itself: How Securities Firms Manage the Legal Hazards of Competitive Pressures. Oxford: Oxford University Press, 1998.

Mookherjee, D. and Png, I. P. L., "Optimal Auditing, Insurance, and Redistribution," Quarterly Journal of Economics, 1989, 104, 399-415.

National Association of Securities Dealers, "Securities Regulation in the United States," NASD, 1996.

Phillips, R. M., "Overview: The Multilayered Securities Enforcement System," in The Securities Enforcement Manual: Tactics and Strategies, ed. by R. M. Phillips. Chicago: American Bar Association, 1997.

Pirrong, S. C., "The Self-Regulation of Commodity Exchanges: The Case of Market Manipulation," Journal of Law and Economics, 38, 141-206.

Saloner, G., "Self-Regulating Commodity Futures Exchanges," in The Industrial Organization of Futures Markets, ed. by R. W. Anderson. Lexington: D. C. Heath, 1984.

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

Shuchman, H. L., Self-Regulation in the Professions. Glastonbury: The Futures Group, 1981.

Townsend, R., "Optimal Contracts and Competitive Markets with Costly State Verification," Journal of Economic Theory, 1979, 22, 265-293.