Governance, Reputation, Crises and Recovery: An Experiment

Thomas H. Noe
Saïd Business School and Balliol College
University of Oxford
Oxford, England, OX1 1HP

Michael J. Rebello
School of Management
University of Texas at Dallas
Richardson, Texas 75080-3021

Thomas A. Rietz
Henry B. Tippie College of Business
University of Iowa
Iowa City, IA 52242

This version: January 2, 2026

Abstract

We model and experimentally test the relationship between a firm's reputation and its governance when firms can repair reputation damage through reform. Our model shows that conditions for establishing a reputation depend on whether control of firm operations is delegated to professional managers. Moreover, although the option to reform always dilutes firms' incentive to commit to reputable behavior, delegation ameliorates the dilution. We assess the external validity of our model through a laboratory experiment. The results of this experiment are directionally consistent with the model's predictions.

JEL Classification Codes: C7, C9, D82, G31, G32, L15
Keywords: corporate governance, stakeholders, adverse selection, reputation, experiment

---

We thank the Oxford University Centre for Corporate Reputation for financial support. For comments and suggestions, we thank participants at the 2021 American Finance Association Meeting, the 46th Annual Meeting of the European Finance Association, the 2020 Financial Management Association Conference, the 2019 Society for Experimental Finance Conference, Copenhagen, and the 2019 Econometric Society North American Summer Meetings. Previously this paper was titled *Governance, Stakeholder Welfare, Crises and Recovery: An Experiment*. All errors are our own.

Corresponding Author: Thomas Rietz, Thomas-Rietz@uiowa.edu, 319-335-0856

The incentives of individuals not to milk the firm's reputation has not been clarified; it must be the case that somehow the incentives of the stock-holding layer trickles down through the rest of the hierarchy.

–Holmstrom and Tirole (1989)

1 Introduction

By some estimates, reputation accounts for more than 60% of a firm's value and is its most valuable asset (Gaines-Ross, 2008). Economists have argued that the desire to maintain a (good) reputation underpins a range of firm policies including responses to competitive threats, product quality, dividend payments, and the financing mix (e.g., Kreps and Wilson, 1982a; Milgrom and Roberts, 1982; John and Nachman, 1985; Diamond, 1989; Maksimovic and Titman, 1991; Mailath and Samuelson, 2001; Cripps et al., 2004). However, the models on which these arguments rest rely on assumptions that do not match the governance structures of many real-world firms or firms' ability to repair damage to their reputations. We examine the implications of the two deviations and find that both have a significant effect on firm reputation. The ability to repair damage weakens the incentive to maintain a reputation and professional management can mitigate this effect.

Models of firm reputation typically assume that firms are "owner-managed," i.e., owners control reputation management because they directly set the policies that impact a firm's reputation.[1] This assumption fits firms like Bechtel, Cargill, Comcast, or Tesla that are controlled by hands-on "mavericks," i.e., non-financial or inside blockholders (Amel-Zadeh et al., 2022). Owner management clearly doesn't fit firms like Boeing, Mattel, and Wells Fargo. Like the majority of large firms in the U.S. and overseas, these firms have a very different governance structure. They are "professionally managed": Operations and reputation management are the purview of professional managers with small or negligible ownership stakes (McMillan, 2011; Economist Intelligence Unit, 2005). Ownership effectively resides with large institutional investors like BlackRock and Fidelity, referred to as "common" or "universal" owners. Such owners hold substantial stakes in many firms and are not intimately involved in individual firm operations or reputation management.

The extensive principal-agent literature suggests that, when owners delegate management to professionals, the policies that firms adopt can be quite different from the policies their owners would prefer (e.g., Myerson, 1982; Shleifer and Vishny, 1997; Cole and Kocherlakota, 2001). This raises a long-standing question, whose importance is evidenced by the opening quotation: Do typical reputation models, which focus on the reputations of owner-managed firms, accurately describe reputations when management is delegated to professionals?

[1] More generally, reputation models have been employed to model the actions of economic agents in incomplete information settings by many authors (e.g., Mailath and Samuelson, 2001; Cripps et al., 2004; Liu, 2011; Cremer, 1986; Tirole, 1996; Morrison and Wilhelm, 2004; Levin and Tadelis, 2005). See Bar-Isaac and Tadelis (2008) for a survey of reputation models. The framework has also been deployed to examine the sustainability of pro-social behavior (e.g., Bénabou and Tirole, 2006), which is consistent with a broader definition of quality in our model and experiment.

Reputation models also tend to assume that reputation damage permanently impairs firm profitability. However, real-world firms tend to "reform" after reputation damage, i.e., they act to repair the damage and restore profitability (Farber, 2005; Gaines-Ross, 2008; Chakravarthy et al., 2014; The Economist, 2018). We provide some examples in Table 1.[2] The reforms we describe are typical, and they tend to reflect guidance from an increasingly sophisticated multi-billion dollar reputation-management industry that has emerged to help firms navigate reputation damage.[3] Consistent with our examples, while some firms fail because of the reputation damage (e.g., Thernos and Takata), many completely recover (e.g., Mattel and Siemens, which have become case studies in reputation repair). Moreover, both owner-managed (e.g., Thernos and Benetton) and professionally-managed firms (e.g., Mattel and Wells-Fargo) experience both outcomes.[4] While this list is not exhaustive, the descriptions of reform are not complete, and the governance structures of all the firms do not fit our model exactly, the examples motivate the following questions: How does the possibility of successful corporate reform ex post affect ex ante incentives to maintain firm reputations? Are the ex ante effects of reform tied to firm governance?

Table 1. *Examples of reform attempts to repair reputation damage.*

Year	**Company**	**Event**	**Ownership Structure**	**Response**
2006-2008	Siemens	Corruption	Publicly traded, but with significant founding family influence.	Set up anticorruption task force. Created rules and compliance processes, training programs, disciplinary actions, terminated employees. Became the subject of a Harvard case study on how to recover from a reputation crisis. Ranked #63 in the Reputation Institute (2025) rankings of the world's 100 most reputable companies.

Continued on next page

[2]While Table 1 briefly describes reform efforts, readers can easily find more detailed descriptions. For example, the New York Times (Nov. 10, 2017) describes Kobe Steel's reform as follows: *Kobe Steel without input from regulators or other outside parties, concluded that the company had erred by elevating the pursuit of short-term profit over the maintenance of scrupulous quality standards. That failing, it said, was exacerbated by lax oversight by senior executives and an "insular" corporate culture that discouraged employees from questioning improper but long-established practices . . . The report published on Friday outlined several changes the company plans to make to prevent cheating, including automating record keeping for product tests and requiring multiple employees to verify that test results are accurate.*

[3]In their survey of 106 UK public relations consultancies, Bennett and Kottasz (2000) find 78% offer crisis-management consulting services to cope with reputation damage.

[4]Some firms in Table 1 have governance structures between these two extremes (e.g., Siemens). Additional examples include Kobe, Takata, and Toyota, which belong to keiretsus. Volkswagen has a particularly complex ownership structure. Effective voting control is owned by Porsche Automobil (in turn, owned entirely by the Porsche and Piech families). The State of Lower Saxony has a significant stake and Volkswagen is partially publicly traded.

Table 1 – *Continued from previous page*

Year	Company	Event	Ownership Structure	Response
2007	Mattel	Lead paint on toys.	Publicly traded with significant institutional and dispersed ownership.	Recalled toys. Created new testing procedures. Changed suppliers. Put Mattel personnel in supplier's factories. Became the subject of a Harvard case study on how to recover from a reputation crisis. Ranked #25 in the Axios and Harris (2024) reputation ranking of the 100 most visible brands in America and #98 in Reputation Institute (2025).
2007	RC2	Lead paint on toys.	Publicly traded with significant institutional and dispersed ownership.	Recalled toys (initially at owners' shipping expense). Takara Tomy purchased RC2 in April 2011.
2009-2010	Toyota	Acceleration problems.	Part of a keiretsu with significant founding family ownership and influence.	Recalled and redesigned systems. Developed the "Toyota Way." Restructured oversight under "Customer First Committee." Ranked #12 in Axios and Harris (2024) and #31 in Reputation Institute (2025).
2009-2016	Takata Air Bags	Manipulated test data on airbag inflators.	Founded, owned, and controlled by the Takada family.	NHTSA ordered revision of production and quality control procedures. NHTSA fined Takata for inadequate response. Takata filed for bankruptcy on 6/26/17.
2012	JP Morgan	"London Whale."	Publicly traded, but "maverick" CEO (Jamie Dimon).	Overhauled risk metrics. Replaced personnel. Dismantled trading arm that caused the problem. Ranked #37 in Axios and Harris (2024).
2012	Barclays	Manipulated LIBOR rates.	Publicly traded with significant institutional and dispersed ownership.	Created new "Brand and Reputation Committee." Replaced CEO. Adopted "zero tolerance" policy. Credit ratings fell after the scandal and have not fully recovered.
2013	Lululemon	"Too sheer" yoga pants.	Founder (Chip Wilson) and CEO (Christine Day) owned 30% of the company and had effective controlling interest.	Created new factory oversight system and new organizational structure. Replaced CEO.
2015	Volkswagen	Faked emission test results.	Effective voting control owned by Porsche Automobil, owned entirely by Porsche and Piech families.	Replaced CEO. Suspended some employees. Recalled cars. Refocused on electric and hybrid vehicles. Compensated owners. Regulators signed off on software updates.
2015-2018	Theranos	Falsified medical testing ability.	Private company with "Maverick" CEO (Elizabeth Holmes)	Theranos denied allegations. Partners stopped using Theranos. SEC charged Thernos with Fraud, removed CEO, and barred her from serving as an officer or director of public companies. Theranos closed 9/5/2018.

Continued on next page

Table 1 – *Continued from previous page*

Year	Company	Event	Ownership Structure	Response
2016	Wells Fargo	Fake accounts.	Publicly traded with significant institutional and dispersed ownership.	Fined by regulators. Terminated employees, ended sales goals, restructured sales practices, replaced CEO and other management. New CEO resigned in April 2019. Ranked #92 in Axios and Harris (2024).
2017	Kobe Steel	Falsified quality data.	Part of a keiretsu with significant founding family ownership and influence.	Changed quality control processes and reporting procedures. Established an "Independent Investigation Committee," a "Quality Governance Restructuring Deliberation Committee," and a "Quality Management Department."
2018	Benetton / Autostrade	Failure to maintain bridge resulting in collapse that killed 43 people.	Owned entirely by Benetton Family.	Italy nationalized some toll road and bridge concessions and restructured others held by Autostrade. CEO of subsidiary later fired. Autostrade was spun out of Benetton's holdings in 2022.
2024 -	Boeing	Repeated quality control issues	Publicly traded with significant institutional and dispersed ownership.	CEO stepped down. FAA demanded comprehensive safety improvement and training programs. Boeing streamlined operations and established a "Safety Management System" and an anonymous "Speak Up" reporting system.

In this paper, we attempt to answer these questions about reform along with the question about the ability of standard reputation models to describe reputations of professionally managed firms. To arrive at our answers, we use both a model and a laboratory experiment to compare reputations of owner and professionally managed firms with their customers in a "reform setting" where firms can attempt to repair reputation damage through reform. We also explore a "no-reform setting" in which the revelation of opportunistic behavior permanently impairs future reputation, a setting that mirrors classic reputation models.

Our model shows that both the option to reform and professional management profoundly affect the incentives of owners when it comes to managing their firm's reputations. Customers recognize the implications for the goods firms produce, so the price for a firm's goods depends on the option to reform and its governance structure. Importantly, in our model, professional management creates stronger incentives to support firm reputation in low trust environments, particularly when reform is possible. Thus, in these environments goods from professionally-managed firms command higher prices and these firms are more profitable. Our experiment confirms that (a) delegation to professional managers fundamentally changes the conditions under which firms act reputably and is better able to support firm reputation in low trust environments; (b) goods' prices vary with governance structure in the way our model predicts, and (c) the presence of an opportunity to reform dilutes the ex ante incentive to act reputably.

In our model, we adopt the incomplete-information reputation framework developed by Kreps and Wilson (1982a) and Milgrom and Roberts (1982) (henceforth KWMR) and frequently employed in experiments (e.g., Camerer and Weigelt, 1988; Neral and Ochs, 1992; Brandts and Figueras, 2003; Noe et al., 2012). We consider a firm whose value depends on its reputation with customers for producing high quality goods.[5] An operator picks the firm's operating policies. The operator may be a "committed type" who, like the "tough monopolist" type in Kreps and Wilson (1982a) or the "honest firm" type in Maksimovic and Titman (1991), always picks a "reputable" policy that always results in high quality goods.[6] If the operator is not the committed type, instead of always picking the reputable policy, the operator can opportunistically adopt a cheaper operating policy that sometimes results in low-quality goods, and internalize the cost savings.

Outsiders, including customers, do not know the operator's type and learn about the operator's commitment by observing the quality of the firm's goods. Outsiders' initial beliefs about the operator's type reflect their "base level of trust" in the firm's commitment to the reputable policy. The production of a low-quality good results in reputation damage, which we refer to as "revelation" because it reveals an absence of commitment to quality. Successful reform completely blocks operator opportunism and ensures high quality goods in the future. While reform can restore some outsider trust in the firm's commitment to the reputable policy, it is costly and its success is uncertain.

Under owner management, as is typical in the KWMR framework, ex ante, the operator is the owner and would like to commit to reputation-ensuring policies. However, the operating policies that affect the firm's reputation are set ex post. When setting these policies, ensuring the reputation is only optimal for an uncommitted operator when the benefits of current opportunism are less than the expected future cost of reduced goods' prices following revelation, which we will refer to as "the reputation cost." Price reductions are larger when the base level of trust is larger. Hence, a *high* base level of trust encourages the owner-operator to eschew opportunism.

Under professional management, the owner faces a very different problem. Like the principal in a typical principal-agent model, the owner (who is an outsider) only influences operating policies indirectly via the terms of the professional operator's compensation. To maintain the firm's reputation, the owner must bear the cost of incentivizing the operator to eschew opportunism. When the operator is committed, the operator will choose reputation-ensuring safe policies even without costly compensation. Thus, while an outsider owner avoids reputation costs by providing incentive compensation, the owner's expected gain from the avoidance is low when the base level of trust is high. Thus, in contrast to owner-management, under professional management a *low* base level of trust encourages owners to ensure reputable behavior and protect the firm's reputation through compensation.

[5]The features that determine "quality" depend on the firm's circumstances. For example, if the customers are other firms, the features could be the adherence to these firms' supplier codes of conduct. Alternatively, if the customers are retail consumers, the features might also include sustainable production practices or the goods' carbon footprint.

[6]The "committed type" is a feature of the KWMR framework that delivers analytical convenience and sharp results (Kreps, 1996). It is motivated by the possibility that, as they often publicly assert, some real-world firms may in fact be committed to only implementing reputable policies. In terms of environmentally responsible production, there is some empirical evidence consistent with more or less committed firm operators (Chatterji et al., 2009). The commitment can stem from operators' intrinsic motivations or social/corporate norms (e.g., Kreps, 1996; Bénabou and Tirole, 2006). Firms often point to some aspect of their culture or internal control systems, the sorts of systems that they overhaul when they reform (see Table 1), to support their assertions that they are committed to acting reputably. For example, Toyota publicized its "Toyota Way" corporate culture to underline its commitment to high quality (Liker, 2004). A lack of commitment to high quality is often revealed by evidence of low quality, and judging by the extensive value losses firms experience in these instances, preserving the perception that there is even a small probability that a firm is committed to reputable policies is valuable (Peltzman, 1981; Jarrell and Peltzman, 1985; Barber and Darrough, 1996; Karpoff and Lott, 1993; Alexander, 1999; Murphy et al., 2009; Karpoff et al., 2008; Karpoff, 2011).

Reform decreases the anticipated reputation cost under both ownership structures, and weakens the owner's ex ante incentive to prevent/eschew opportunism. However, the ex ante incentive effect varies across the ownership structures. Under owner management, because firm value flows entirely to the owner, the owner internalizes both the reputation cost and the benefits of reform. Under professional management, the professional operator is replaced when disreputable conduct is revealed, so the operator manager will not internalize the gains from reform. Consequently, delegating to a professional manager separates the agent whose actions affect reputation from the benefits of corporate reform, which mitigates the pernicious ex ante effects of corporate reform and promotes commitment to reputation assuring policies.

Testing our model's predictions with field data would require measuring firm policies, customer expectations, managerial commitment, and outsider trust, which is difficult, if not impossible, and, in addition, tackling thorny issues related to endogeneity.[7] Plott (1991) argues that a laboratory experiment can serve as a first step in assessing the external validity of a model.[8] In a laboratory setting, we can accurately make measurements that are important in our model, control for confounding factors, and we can set governance policy (e.g., managerial contracts) and reform to match our model. Moreover, it is not unreasonable to expect the decisions of individuals playing the roles of customers and operators in an experiment to approximate the policy decisions made by real-world customers and operators that are central to our model.[9] Therefore, we conduct a laboratory experiment and test whether (i) governance structure affects opportunism; (ii) reform weakens the ex ante incentive to prevent opportunism; (iii) customers recognize how governance structure affects opportunism; and (iv) whether customer responses create incentives for reputable behavior like those in our model.

Like prior reputation experiments using the KWMR framework, we find that subject behavior deviates from rational-expectations predictions. Experience weighted attraction models similar to Erev and Roth (1998) and Camerer et al. (2002) explain how subject behavior evolves with experience.[10] Despite the deviations from predicted behavior, we find that (i) reputable firm actions are commonly observed; (iii) customers anticipate firms' actions reasonably well; (iii) customers punish firms when opportunism is revealed; and (iv) reform raises the likelihood that firms will act opportunistically. The important insights of our model survive the behavioral deviations: When the model predicts that the likelihood of reputable actions is lower under owner management, owner management actually does result in significantly more opportunism. Moreover, when the model predicts ownership structure will not affect the level of reputable actions, there are not significant differences between opportunism levels under owner and professional management.

[7] Governance structures vary across firms, industries, and countries, and are not randomly assigned. The structures may be chosen for their reputational effects, but myriad other factors also matter. While some of these factors may be observable, others, such as the managerial human capital of owners, may not. Econometric fixes may not resolve ownership structure endogeneity (Coles et al., 2012).

[8] By "external validity" we mean that the major insights of the model hold up in spite of behavioral deviations optimizing behavior. In the laboratory.

[9] While we use student subjects, surveys of experimental economics research comparing students to business professionals (e.g., see surveys by Ball and Cech, 1996; Frechette, 2015) generally find little difference in outcomes.

[10] Such adaptive learning is commonly observed in experimental settings (e.g., Part 4.3 in Plott, 1982).

Related literature

Our model of owner management closely matches the KWMR framework on which much of the research on reputation is based. This literature includes many models (e.g., John and Nachman, 1985; Diamond, 1989; Maksimovic and Titman, 1991; Fudenberg and Levine, 1992). Our experimental implementation of owner management also closely matches experiments using the KWMR framework (e.g., Camerer and Weigelt, 1988; Neral and Ochs, 1992; Brandts and Figueras, 2003; Noe et al., 2012).[11]

One departure we make from this literature is that we introduce the option to reform after a firm's reputation is damaged. A more fundamental departure is that we model a setting in which the firm is managed by a professional manager rather than an owner. In the professional management setting, the firm's rents from reputation are owned by uninformed outsiders who can only use governance policy to indirectly affect the informed insider's actions, whose actions directly determine the firm's reputation. Second, managers have no personal reputation, no ownership claim to the rents produced by firm reputation, and can be costlessly replaced at any time. These model features reflect principal-agent models in which an uninformed principal (outside owner) sets policies to incentivize an informed agent (professional manager) to maximize the principal's welfare (e.g., Myerson, 1982; Shleifer and Vishny, 1997; Cole and Kocherlakota, 2001). However, in contrast to typical principal-agent models, the effects of the agent's actions are not exogenously specified. Rather, as in typical reputation models, the effects are produced by the behavior of outsiders (customers) who have incomplete information about the firm. Our model and experiment show that incorporating this principal-agent problem into a reputation framework fundamentally changes the calculus of reputation.

2 The model

We model a firm that operates for two periods.[12] Agents are risk-neutral and patient. The firm has an owner and an operator. At the start of each period, the owner supplies capital, denoted by $e > 0$. The operator uses the capital to produce a good and sets a policy that determines the quality of the good. The good is sold to customers for the numeraire good, cash. There is no storage technology, thus cash and any good produced in a period must be consumed during the period.

We consider reform and two governance structures: owner management and professional management. Under owner management, the owner and operator are the same agent. Under professional management, the owner and operator are different agents. The owner receives all of the firm's profits under both governance structures. Under professional management, the operator's payoffs are determined by a contract chosen by the owner. We will start by describing and analyzing model features common to both governance structures when the firm can reform. Then we will use this analysis to identify the effects of governance structure and reform.

[11] In these experiments, reputations and trust arise from supposedly "rational" behaviors. This contrasts with experimental "trust games" where behavior is not explained by self-interested rational choice of (e.g., Berg et al., 1995).

[12] As shown by Noe et al. (2025), in settings like the ones we consider, increasing the number of periods, by increasing the potential expected loss from opportunistic behavior, can make it easier to sustain equilibria in which firms produce high-quality goods due to reputational concerns.

Quality A good's quality in period $t = 1,2$, denoted by q_t, is either high, h, or low, l. All agents commonly observe q_t only at the end of period t.[13] Quality is neither verifiable nor contractible.

Prices Customers assign value $v_t = 1$ to a quality h good and $v_t = 0$ to a quality l good in period $t = 1,2$. The period t good's price, $p_t \in \mathbb{R}_+$, is set at the beginning of the period through Bertrand competition and equals customers' expectation of the good's value.[14] Goods' prices are verifiable and contractible.

Operating policy In period $t = 1,2$, the operator sets an operating policy, o_t. The policy is not observed by other agents, whom we refer to as "outsiders." There are two possibilities: "safe," denoted by s, and "mixed," denoted by m. Policy s costs e to implement and always produces a high-quality good. Policy m costs $e - c$, where $c \in (0,e)$. If the operator chooses policy m in period t, with probability δ the quality of the period t good is high, and with complementary probability the quality is low. If m is chosen in both periods, the quality draws across periods are independent.

We impose the following restriction on the mixed policy, which ensures that the firm cannot operate profitably if customers are certain that the operator will choose policy m.

Assumption 1 $\delta < e$.

Operator type There are two types of operators: "committed" and "normal." A committed operator is non-strategic and always chooses policy s. A normal operator chooses strategically between acting reputably, i.e., choosing policy s, and acting opportunistically, i.e., setting policy m and consuming the cost savings.[15] Operator type is the operator's private information. Outsiders assign prior probability $\rho \in (0,1)$ to the operator being committed. Thus, ρ measures outsiders' base level of trust in the firm's commitment to high quality and policy s.[16]

Revelation Only a normal operator can produce a low quality good, and we will refer to the operator as being "revealed" (to outsiders) if $q_1 = l$. Otherwise, the operator remains "unrevealed."

Reform If the operator is revealed, the firm reforms at the start of period 2. Reform imposes a cost on the owner equal to R and is publicly observable. Reform is an attempt to make the operator in period 2 commit to the safe operating policy s. The period 2 operator behaves like a committed type if and only if reform succeeds. Reform succeeds with probability r and fails with probability $1 - r$. Only the period 2 operator observes whether reform succeeds.

We impose the following restriction on the effectiveness of reform, which ensures that reform cannot completely restore customers' trust in the firm's commitment to high quality.

Assumption 2 $0 < r < \rho$.

[13]This assumption implies that both customers who do and do not purchase an item observe its quality. Information on quality, especially large lapses, can spread in many ways including word-of-mouth (in person or online, e.g., Yelp reviews or Amazon star ratings), publications and websites (e.g., *Consumer Reports* and JD Powers), through corporate recalls, etc.

[14]Intuitively, this means that the customers' willingness to pay determines the price for the firm's goods based on expected quality. See Resnick et al. (2006) for a real-world example linking reputation to prices based on willingness to pay.

[15]We tie the adoption of opportunistic policies to the operator's type to simplify the model. We would obtain identical equilibrium outcomes if we tied the operator's ability to act opportunistically to a firm monitoring system. In fact, a large literature models monitoring systems that increase the costs of opportunism (Johnson et al., 2000; Shleifer and Wolfenzon, 2002).

[16]One can think of this as an accurate Bayesian prior formed from whatever information outsiders have available at the start of period 1.

Payment structure Each period the operator receives a payment from the firm that depends on a "payment structure." The period 1 payment, ϕ_1, depends on the period 1 good's price, i.e., $\phi_1 : [0,1] \to \mathbb{R}_+$. The period 2 payment depends on the period 2 good's price as well as the period 1 good's quality, and we denote the payment by $\phi_h : [0,1] \to \mathbb{R}_+$ when $q_1 = h$ and by $\phi_l : [0,1] \to \mathbb{R}_+$ when $q_1 = l$. We denote the payment structure by the triple $\phi \equiv (\phi_1, \phi_h, \phi_l)$, and represent the set of payment structures with Φ.

We denote a normal operator's period 2 payoff which, in addition to the payment from the firm, also reflects any cost savings from the period 2 policy choice, by

$$v_2^{\mathrm{o}}(o_2, h) = \begin{cases} \phi_h(p_2) & o_2 = s, \\ \phi_h(p_2) + c & o_2 = m, \end{cases} \qquad v_2^{\mathrm{o}}(o_2, l) = \begin{cases} \phi_l(p_2) & o_2 = s, \\ \phi_l(p_2) + c & o_2 = m. \end{cases} \tag{1}$$

We denote the expected payoff to a normal period 1 operator by

$$v_1^{\mathrm{o}}(o_1) = \phi_1(p_1) + \begin{cases} v_2^{\mathrm{o}}(o_2, h) & o_1 = s. \\ c + \delta\, v_2^{\mathrm{o}}(o_2, h) + (1-\delta)\Big(r\, v_2^{\mathrm{o}}(s, l) + (1-r) v_2^{\mathrm{o}}(o_2, l)\Big) & o_1 = m, \end{cases} \tag{2}$$

The payment structure is set before any action is taken in period 1 and ultimately depends on the firm's governance structure in a fashion that we will describe later. At the time that the operator makes operating decisions, how the payment structure originated has no effect on the operator's welfare and thus on the operator's best replies.

Operator replacement Reform can impact the operator's ability to continue with the firm. As we will detail later, the impact tends to vary with governance structure. Therefore, we examine two settings with regard to the impact of reform operator replacement. In one setting the operator from period 1 continues to operate the firm in period 2 whether or not the firm reforms. In this setting, the payment structure is what we have just described. We also consider a setting in which the operator is replaced if and only if the firm reforms, and all agents know this. The following aspects of reform are also common knowledge in this setting: The replacement is a normal type, whose characteristics and preferences are the same as the period 1 normal operator's. The replacement behaves like a committed type if and only if reform succeeds.

In the setting in which the period 1 operator is replaced after reform, the replaced operator receives no payment, i.e., $\phi_l := 0, \forall\, p_2 \in [0,1]$, and no longer controls operating policy so the replaced operator's period 2 payoff is zero. As will be apparent later, because the period 2 goods' price depends on operator actions in period 1 but not actions in period 2, the replacement operator receives no payment, i.e., $\phi_q(p_2) := 0, \forall\, q \in \{h, l\}, \forall\, p_2 \in [0,1]$ for the replacement operator.

Normal operator history and strategy Figure 1 shows the sequence of events in model. As the figure shows, before setting the operating policy in period 1, the operator observes the payment structure and the period 1 good's price. In period 2, the operator has also observed o_1, q_1, and p_2. Since the policy-realized quality pair (s, l) is impossible when $o_t = s$, we can denote a normal operator's period t history by $h_t^o \in \mathcal{H}_t^o$, $t = 1, 2$, where

$$\begin{aligned} \mathcal{H}_1^o &:= \Phi \times \mathbb{R}_+, \\ \mathcal{H}_2^o &:= \mathcal{H}_1^o \times (\{s, m\} \times \{h, l\}) \backslash \{(s, l)\}) \times \mathbb{R}_+. \end{aligned} \tag{3}$$

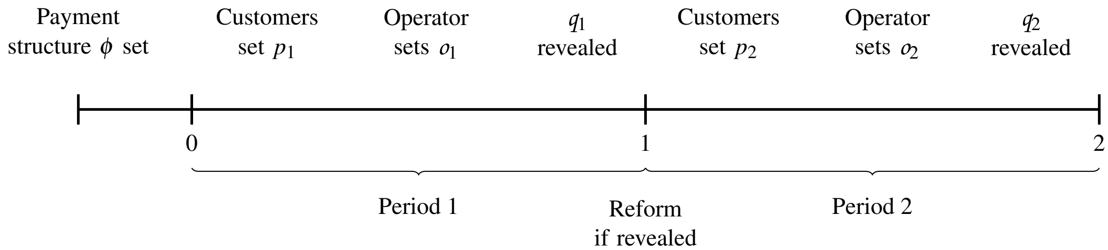

Figure 1. *Timeline.* This figure shows the sequence of events in the game.

Each period, the normal operator maximizes his expected payoff under the payment structure. Operators can take one of two actions at each of their histories, m and s. Let $O := \{s,m\}$. A normal operator pure strategy is an ordered pair $o := (o_1, o_2)$, where $o_t : \mathcal{H}_t^o \to O$, $t = 1,2$.

We permit operators to play mixed strategies, i.e., choose a probability distributions from the set of probability distributions over O, which we denote by ΔO. These distributions are conditioned on the operator's history. We represent the support of a strategy with Supp$[\Delta O]$. A normal operator's strategy is an ordered pair $\boldsymbol{\sigma} := (\sigma_1, \sigma_2)$, where $\sigma_t : \mathcal{H}_t^o \to \Delta O$, $t = 1,2$.

Customers' history and goods' prices Before the period 1 good's price is set, outsiders observe the payment structure. At the end of period 1 outsiders observe the quality of the period 1 good and whether the firm reforms. Thus, we can represent customers' histories at the start of each period, when they set prices, by $h_t^c \in \mathcal{H}_t^c$, $t = 1,2$, where $\mathcal{H}_1^c := \Phi$ and $\mathcal{H}_2^c := \mathcal{H}_1^c \times \{h, l\}$. Let $P = (P_1, P_2)$ represent the vector of pricing functions, where $P_t : \mathcal{H}_t^c \to [0,1]$, $t = 1,2$. Prices also depend on customers' belief about the operator's type. We denote these beliefs by μ_t, where $\mu_t : \mathcal{H}_t^o \to [0,1]$, $t = 1,2$.

Definition 1 [Equilibrium] Let SE(ϕ) represent the set of sequential equilibria (Kreps and Wilson, 1982b) in the subgame rooted in payment structure ϕ, i.e., the set of ordered triples $(P^*, \boldsymbol{\sigma}^*, \mu^*)$ that satisfy, under payment structure ϕ, the following conditions for being a sequential equilibrium:

(a) Pricing condition: $P_t^*(h_t^c) = \mathbb{E}[\tilde{v}_t | h_t^c]$, for all $h_t^c \in \mathcal{H}_t^c$, $t = 1,2$.

(b) Operator sequential rationality: for all $h_t^o \in \mathcal{H}_t^o$, $t = 1,2$, if $o \in$ Supp$[\sigma_t^*(h_t^o)]$, then o is a best response for the operator.

(c) Belief consistency: Customer beliefs, μ^*, are determined by Bayes rule for all histories reached on the equilibrium path. For histories off the equilibrium path, beliefs must be the limit of some sequence of purely mixed normal operator strategies that reach the history.

If the normal operator ever chooses policy s in equilibrium with probability 1, it will only happen in period 1. We refer to equilibria with this property as "reputation equilibria."

Definition 2 [Reputation equilibria] For a given payment structure $\phi \in \Phi$, a *reputation equilibrium* is an equilibrium in which $\tilde{q}_1^* = h$ with probability 1.

3 Foundational results

Now, to greatly simplify our analysis of governance structures and firm reputation, we will derive results using the common model elements we have just described. All derivations not in the body of the paper are provided in Appendix A.

We start with the following observation: The operator's contracted payment in period 1 depends only on the good's period 1 price, which is set before the operator moves, and the operator's period 2 payment does not depend on the period 1 price. Thus, the difference between the normal operator's payoff from operating decisions m and s is not affected by the period 1 price. Hence, the set of best responses for the operator are the same at all period 1 operator histories. For this reason, we restrict attention to equilibria in which the operator's period 1 strategy is invariant over operator period 1 histories.

Next, we derive the operator's period 2 strategy. Because a good's price reflects only its expected quality and the operator's actions are unobservable, operating policy in the current period has no effect on the payment received by the operator in the current period. In period 2 there are no future periods, so the only effect of choosing the mixed policy, m, over the safe policy, s, is that the operator captures the cost saving, c. Thus, a normal operator will always act opportunistically in period 2.

Lemma 1 *In period 2, it is strictly optimal for the normal operator to set operating policy $o_2 = m$ for all operator histories $h_2^o \in \mathcal{H}_2^o$.*

As Lemma 1 shows, a normal operator's unique best response in period 2 is $o_2 = m$ and the operator never plays a mixed strategy in period 2. So, to simplify notation, for all period 1 operator histories, we will represent, an operator strategy at period 1 simply by σ, where σ represents the probability that the operator chooses action s with probability σ at all period 1 operator histories.

In a sequential equilibrium, if customers anticipate that the operator's period 1 strategy is σ, the period 1 good's price will satisfy condition (a) of Definition 1, i.e., $P_1^*(h_1^c) = \mathbb{E}[\tilde{v}_1 | h_1^c] = P_1(\sigma)$, where

$$P_1(\sigma) = \rho + (1-\rho)\Big(\sigma + (1-\sigma)\delta\Big). \tag{4}$$

After observing the quality of the period 1 good, customers will update their beliefs. If $q_1 = l$, the firm will be reformed and thus the probability that the firm is committed equals r. If $q_1 = h$, because the period 1 price, P_1, equals the probability that good quality equaled h in period 1, Bayes rule implies that the updated belief that the firm is the committed type is $\rho/P_1(\sigma)$. Thus, the customers' posterior belief that the operator is the committed type at the start of period 2, will equal, for all $h_2^c \in \mathcal{H}_2^c$, $\mu_2^*(h_2^c) = \mu_2(\sigma)$, where

$$\mu_2(\sigma) = \begin{cases} \frac{\rho}{P_1(\sigma)} & \text{if } q_1 = h \\ r & \text{if } q_1 = l \end{cases}. \tag{5}$$

Lemma 1 implies that the period 2 good's equilibrium price will reflect the expectation that the normal type operator will act opportunistically in the period. Thus, the period 2 good's price will depend only on

μ_2, customers' posterior belief at the start of period 2. Consequently, $P_2^*(\hbar_2^c) = \mathbb{E}[\tilde{v}_2|\hbar_2^c] = P_2(\mu_2)$, where

$$P_2(\mu_2) := \mu_2 + \delta\,(1-\mu_2). \tag{6}$$

It follows that the period 2 good's price will equal $P_2(r)$ after customers observe $q_1 = l$. If customers anticipate that $\sigma = 1$ and they observe $q_1 = h$, the period 2 good's price will equal $P_2(\rho)$. Finally, if customers anticipate that $\sigma \in (0,1)$ and they observe $q_1 = h$, the period 2 good's price will equal $P_2(\rho/P_1(\sigma))$.

Lemma 2 *In an equilibrium in which the operator picks the safe policy with probability σ^*, $p_1^* = P_1(\sigma^*)$. If $q_1 = h$, then $\mu_2^* = \rho/P_1(\sigma^*)$ and $p_2^* = P_2(\rho/P_1(\sigma^*))$. If $q_1 = l$, then $\mu_2^* = r$ and $p_2^* = P_2(r)$.*

In each period, the operator's payment is tied to the good's price. Once the operator is revealed, the good's price is fixed at $P_2(r)$. Thus, to streamline notation, we will switch to denoting that operator's period 2 payment after revelation by $\phi_l \geq 0$. It follows that, if there is no operator replacement, using Lemmas 1 and 2, if customers expect the operator to set $o_1 = s$ with probability σ, we can rewrite expression (2) to represent the expected payoff to a normal period 1 operator as

$$v_1^{\mathrm{o}}(o_1) = \phi_1(P_1(\sigma)) + \begin{cases} \phi_h(P_2(\rho/P_1(\sigma))) + c & o_1 = s \\ c + \delta\left(\phi_h(P_2(\rho/P_1(\sigma))) + c\right) + (1-\delta)\left(r\phi_l + (1-r)(\phi_l + c)\right) & o_1 = m \end{cases}. \tag{7}$$

If the operator is replaced, the period 1 operator's expected payoff is given by expression (7) after dropping all the terms associated with $1-\delta$ when $o_1 = m$.

Expression (7) shows that, when picking the period 1 operating policy, o_1, the operator trades off the immediate gain from opportunism, c, against the cost of revelation, which will depend on the payments $\phi_h(P_2)$ and ϕ_l. In equilibrium the customers' conjecture about σ must be correct. The following result describes equilibrium conditions based on this tradeoff.

Lemma 3 *(a) In a setting in which the operator remains with the firm after reform, an equilibrium in which the operator picks the safe policy with probability σ^* in period 1 exists if and only if*

$$\begin{aligned} \sigma^* = 1 \quad &\textit{and} \quad (1-\delta)\Big(\phi_h(P_2(\rho)) - \phi_l\Big) &&\geq \delta c + (1-\delta)(1-r)c, \\ \sigma^* \in (0,1) \quad &\textit{and} \quad (1-\delta)\Big(\phi_h\left(P_2\left(\rho/P_1(\sigma^*)\right)\right) - \phi_l\Big) &&= \delta c + (1-\delta)(1-r)c, \\ \sigma^* = 0 \quad &\textit{and} \quad (1-\delta)\Big(\phi_h\left(P_2\left(\rho/P_1(0)\right)\right) - \phi_l\Big) &&\leq \delta c + (1-\delta)(1-r)c. \end{aligned}$$

(b) In a setting in which the operator is replaced after reform, an equilibrium in which the operator picks the safe policy with probability σ^* in period 1 exists if and only if

$$\begin{aligned} \sigma^* = 1 \quad &\textit{and} \quad (1-\delta)\phi_h P_2(\rho) &&\geq \delta c, \\ \sigma^* \in (0,1) \quad &\textit{and} \quad (1-\delta)\phi_h P_2\left(\rho/P_1(\sigma^*)\right) &&= \delta c, \\ \sigma^* = 0 \quad &\textit{and} \quad (1-\delta)\phi_h P_2\left(\rho/P_1(0)\right) &&\leq \delta c. \end{aligned}$$

(c) *Every equilibrium is a reputation equilibrium if the applicable condition for $\sigma^* = 1$ in Part* (a) *or* (b) *is satisfied and ϕ_h is a strictly increasing function of the period 2 good's price.*

Lemma 3 shows that reputation equilibria exist when the normal operator expects revelation to trigger a sufficiently large reputation cost. It also shows that the operator's period 1 payment has no effect on the operator's equilibrium behavior. If the operator expects to remain with the firm after revelation, the operator's reputation cost is proportional to $\phi_h(P_2(\rho)) - \phi_l$. For reputation equilibria to exist this difference should be large relative to the operator's anticipated gain from opportunistically choosing $o_1 = m$ instead of $o_1 = s$, which equals c in period 1 less $c(1-\delta)r$ in period 2 because successful reform blocks period 2 opportunism. If the operator expects to be replaced following reform, the operator anticipates both a larger reputation cost, $\phi_h(P_2(\rho)) - 0$, and a smaller gain from opportunistically choosing $o_1 = m$, $c(1-(1-\delta))$. Thus, it is easier to satisfy the condition for reputation equilibria in part (b) of Lemma 3 than in part (a).

4 Firm reputation, governance structures, and reform

We will now examine whether a firm's reputation is different under professional management than it is under owner management, the governance structure typically employed in reputation models. We will start by describing the payment structure and operator replacement under owner management and derive conditions for reputation equilibria under it. Then we will turn to professional management. We will end the section by examining how reform, another unique feature of our model, affects equilibrium conditions under both professional and owner management.

4.1 Owner management and reputation equilibria

As is typical in models of firm reputation, under owner management, the owner is the operator and, each period, the owner-operator's payment equals the firm's profit, i.e., revenue (good's price) minus the investment cost, e. The owner-operator also internalizes both the cost of reform, R, and its benefit. To capture the operator's benefit from reform in the simplest possible way, we assume that the owner-operator remains with the firm in period 2 whether or not it reforms.[17]

Thus, under owner management, the operator's period 2 payment is $\phi_h(p_2) = p_2 - e$ after $q_1 = h$ and $\phi_l = P_2(r) - e - R$ after $q_1 = l$. In period 1, the operator's payment is $\phi(p_1) = p_1 - e$. Because this payment structure is fixed exogenously, an equilibrium under owner management is simply an equilibrium (Definition 2) for the payment structure we have just described. Let π_1 and π_r denote the firm's profit when the good's price is 1 and profit after reform, respectively, i.e., $\pi_1 = 1 - e$ and $\pi_r = P_2(r) - e - R = \phi_l$. The next proposition describes the condition for the existence of a reputation equilibrium under owner management. It also shows that all equilibria are reputation equilibria when this condition is satisfied.

Proposition 1 *If*

$$(1-\delta)\Big(P_2(\rho) - [1 - \pi_1 + \pi_r]\Big) \geq \delta c + (1-\delta)(1-r)c, \tag{8}$$

[17]Clearly, putting a professional manager in place as the operator instead of the owner after reform, while depriving the owner-operator of the rents from failed reform, will change the governance structure to professional management. It is possible that the owner could sell the firm to another owner-operator as part of the reform. However, since the replacement owner-operator would anticipate capturing the rents from diversion after failed reform, the sale price would reflect these rents and transfer them to the original owner through the sale price. Incorporating the sale option would thus offer little new insight.

a unique equilibrium exists and that equilibrium is a reputation equilibrium.

Since $P_2(\rho)$ is increasing in ρ, Proposition 1 shows that, as is typical in models of firm reputation, under owner management, reputation equilibria exist when outsiders have a sufficiently high base level of trust in the operator's commitment. The reason is that the owner-operator's reputation cost, given by the left-hand side of expression (8), is increasing in outsiders' base level of trust.

4.2 Reputation under professional management

Under professional management, the owner and operator are different agents. The payment structure is codified in a compensation contract for the professional operator. The owner, who is an outsider and is entitled to the firm's profits, chooses the compensation contract before the period 1 good's price is set.

The professional operator has no natural claim on the benefit of reform. Lemma 3 shows that it is easier to satisfy the conditions for a reputation equilibrium if the operator is replaced when the firm reforms. It also shows that, if a revealed operator is replaced, only the period 2 payment contingent on $q_1 = h$ matters for the period 1 operating policy choice. Thus, to keep the analysis of professional management simple and transparent, we assume that the owner replaces the professional operator if and only if the firm reforms, and the operator receives no payment in period 1, i.e., $\phi_1(p_1) = 0$ for all p_1. Note that, as we have described previously in Section 2, in period 2, a replaced operator receives his reservation wage of zero and a replacement operator receives a payment of zero, i.e., $\phi_l = 0$.

Lemma 1 shows that, regardless of the compensation contract, a normal professional operator will pick policy m in period 2. Thus, the contract can only influence the professional operator's period 1 policy choice. The owner, in each period, receives the firm's profit net of the operator payment. Thus, the owner's welfare depends both on the operator's actions and the cost of the contract, i.e., expected compensation payments to the operator. Consequently, across all contracts that implement the same operator choice in period 1, the owner's welfare is maximized by minimizing the contract cost.

Suppose that the owner wants to implement $o_1 = m$. Lemma 3 shows that setting the operator's period 2 payment after $q_1 = h$ to zero, i.e., $\phi_h(p_2) = 0$ for all p_2, will be sufficient to ensure this policy is implemented. Moreover, making no payments to the professional operator clearly minimizes the cost of the contract. Since customers observe the operator's contract before setting the period 1 good's price, they will anticipate that the firm will implement $o_1 = m$, and they will set $p_1 = P_1(0)$ and $p_2 = P_2(\mu_2(0))$ if $q_1 = h$. Therefore, noting that $\pi_r = P_2(r) - R - e$, the owner's expected payoff from choosing to make no payments to the professional operator will equal

$$P_1(0) - e + P_1(0)\Big(P_2(\mu_2(0)) - e\Big) + \Big(1 - P_1(0)\Big)\pi_r. \tag{9}$$

Alternatively, the owner may want a contract that implements $o_1 = s$. In this case, the contract must satisfy the condition in Part b of Lemma 3 that ensures that $\sigma^* = 1$. Since the operator receives no payment in period 1 or in period 2 when $q_1 = l$, the period 2 payment when $q_1 = h$ must satisfy this condition when $\phi_l = 0$. Let us denote the "bonus payment" in period 2 when $q_1 = h$ by $b \geq 0$. After observing the contract, customers will set $p_1 = P_1(1)$ and $p_2 = P_2(\rho)$ if $q_1 = h$. Thus, noting that $\pi_1 = 1 - e$, the optimal contract is a solution to the following problem for the owner

$$\begin{aligned} \max_{b} \quad & \pi_1 + P_2(\rho) - e - b \\ \text{subject to} \quad & (1-\delta)b \geq \delta c \\ & b + c \geq 0, \end{aligned} \tag{10}$$

where the first constraint reflects the equilibrium conditions for $\sigma = 1$ in Part b of Lemma 3 and ensures that $o_1 = s$ is incentive compatible for the operator, and the second constraint ensures operator participation. Since $c > 0$ the operator's participation constraint is always satisfied, so the binding constraint is the incentive compatibility constraint. Solving for the smallest payment that satisfies this constraint yields

$$b^* := c \frac{\delta}{1-\delta}, \tag{11}$$

which we will refer to as "reputation-assuring compensation."[18]

In equilibrium the contract will maximize the owner's welfare. Thus, the owner will choose to implement $o_1 = s$ whenever the resulting payoff,

$$\pi_1 + P_2(\rho) - e - b^*, \tag{12}$$

exceeds the payoff from implementing $o_1 = m$, given by expression (9).

Proposition 2 *Under professional management, the owner will prefer to implement $o_1 = s$ when*

$$P_2(\rho) \leq 1 - \frac{b^*}{\pi_1 - \pi_r + \delta}. \tag{13}$$

When condition (13) *is satisfied and the compensation contract the owner chooses implements a payment structure with $\phi_h(p_2)$ that is strictly increasing in p_2 and $\phi_h(P_2(\rho)) = b^*$, where b^* is described in* (11)*, $\sigma^* = 1$ in all sequential equilibria.*

Proposition 2 shows that, under professional management, satisfaction of equation (13) is sufficient to ensure that equilibria exist in which the manager will choose $o_1 = s$. Since $P_2(\rho)$ is increasing in ρ, under professional management, these equilibria exist when outsiders have a sufficiently low base level of trust in the operator's commitment. This requirement contrasts with the condition for reputation equilibria under owner management that outsiders have a sufficiently high base level of trust.

The proposition also shows that, when condition (13) is satisfied any bonus payment that is strictly increasing in the period 2 good's price ensures that equilibria with $o_1 = m$ or $\sigma \in (0,1)$ are not possible. Making the bonus payment strictly increasing has no cost to the owner because the only period 2 price on the equilibrium path when the owner implements $o_1 = s$ is $P_2(\rho)$, so only the payment to the manager when

[18] If a professional operator's reservation value were positive, the operator's reservation constraint might be the binding constraint and the firm's owner could ensure high-quality output in the first period without conceding rents to the operator. Because conceding rents to the operator is the only thing deterring the owner from offering reputation-assuring compensation, a positive reservation payoff would never discourage and sometimes encourage the owner to adopt reputation-assuring compensation. Thus, the zero reservation assumption, militates *against* reputable behavior under professional management and thus against the superiority of professional management relative to owner management with respect to assuring reputation.

$p_2 = P_2(\rho)$, i.e., b^*, affects the owner's payoff. However, a strictly increasing bonus payment does ensure that equilibria that support policies other than $o_1 = s$ do not exist.

Interpreting Proposition 2 is straightforward: reputation-assuring compensation yields two benefits. First, it eliminates the possibility of revelation. This results in a direct gain to the owner that is inversely related to the profit after reform, π_r. Hence, low values of the profit after reform favor equilibria with reputable firm behavior. Second, reputation-assuring compensation has an informational effect: when customers observe such compensation in period 1, customers rationally infer that the operator will choose policy s and thus the period 1 good's price equals one, and the firm's period-one profit equals $\pi_1 := 1 - e$. If the owner opts for a no compensation contract, customers rationally infer that the normal operator will choose policy m and thus the period 1 good's price equals $P_1(0) = P_2(\rho)$ and period 1 profit equals $P_2(\rho) - e$. Thus, the period 1 gain from paying reputation-assuring compensation is inversely related to $P_2(\rho)$, and hence inversely proportional to ρ, the base-level of outsider trust. Hence, under professional management, low base levels of outsider trust favor equilibria in which the firm acts reputably.

4.3 Governance structure, reform and firm reputation

To illustrate the effect of the base level of outsider trust on the viability of equilibria in which the firm acts reputably under professional and owner management, we present an example rooted in our experiment. Consider Table 2. Panel A presents the parameter sets we use in our experiment, and Panel B presents calculated equilibrium operator payoffs for each parameter set. Note that the two parameter sets vary only with respect to the base level of outsider trust, ρ, which is higher in Parameter Set I.

First consider Parameter Set I under professional management. Reputation-assuring compensation costs the owner $b^* = 0.0032$ (equation (11)) resulting from the payment when the operator is unrevealed and thus $p_2 = P_2(\rho)$. A normal professional operator's payoff from choosing $o_1 = s$ is $b^* + c = 0.0632$, produced entirely by the cost savings, c, from choosing policy m in period 2 and the bonus payment b^*. Choosing $o_1 = m$, results in the same payoff, in this case, produced by the cost savings, $c = 0.06$, from choosing policy m in period 1, and receiving $(b^* + c) = 0.0632$ in period 2 contingent on remaining unrevealed, which has a 5% probability, i.e., the operator's payoff from choosing policy m is $0.06 + 0.05 \times 0.0632 = 0.0632$. Thus, $o_1 = s$ is a best response for the professional operator.

For the owner, the reputation-assuring compensation produces prices $p_1 = P_1(1) = 1$ and $p_2 = P_2(\rho) = 0.2875$, thus the owner's payoff equals $(1 - e) + (P_2(\rho) - e - b^*) = 1.0623$ according to expression (12). In contrast, if the owner chooses no compensation, then, in period 1, $p_1 = P_1(0) = 0.2875$ (equation (4)); in period 2, $p_2 = P_2(\rho/P_1(0)) = 0.8761$ (equations (5) and (6)) if the operator is unrevealed, and $p_2 = P_2(r) = 0.1450$ if the operator is revealed. The resulting owner payoff is 0.4136 (equation (9)). Hence, the owner will strictly prefer paying reputation-assuring compensation and $o_1 = s$ in all equilibria.

Now consider Parameter Set II/professional management. The cost of the reputation-assuring compensation, b^*, does not depend on outsider trust, and so it is unchanged. However, because the base level of trust, ρ, is lower, conditional on no compensation, the period 1 good's price will be lower than under Parameter Set I. This price drop makes the cost of not assuring reputation even larger. Performing the same calculations as we performed for Parameter Set I, we see that the payoff from reputation assuring compensation is 0.9436 and the payoff from no compensation is 0.1862, so, again, the owner prefers paying reputation-assuring compensation and $o_1 = s$ in all equilibria. This is not surprising given the lower base level of trust in Parameter Set II compared with Parameter Set I and Proposition 2, which shows that low base levels of trust favor paying reputation-assuring compensation under professional management.

Table 2. Example

A. Parameters

Parameter	ρ	e	c	δ	R	r	Value l	Value h
Parameter Set I	0.250	0.111	0.060	0.05	0.010	0.1	0	1
Parameter Set II	0.125	0.111	0.060	0.05	0.010	0.1	0	1

B. Detailed operator payoff calculations

	Professional Management		Owner Management		
	$q_1 = h$	$q_1 = l$	$q_1 = h$	$q_1 = l$	
				Reform succeeds	Reform fails
Normal operator payoffs					
	$b^* + c$	0	$p_2^* - (e - c)$	$F_r - e - R$	$F_r - (e - c) - R$
Parameter Set I	0.0632	0	0.237	0.024	0.084
Parameter Set II	0.0632	0	0.118	0.024	0.084
Committed operator payoffs					
	b^*	NA	$p_2^* - e$	NA	NA
Parameter Set I	0.0032		0.177		
Parameter Set II	0.0032		0.058		

Now consider Parameter Set I/owner management. In a reputation equilibrium, $p_2^* = P_2(\rho) = 0.2875$ if $q_1 = h$ and $p_2^* = P_2(r) = 0.145$ if $q_1 = l$ (Lemma 3). Since the operator receives the firm's profits, as we show in Table 2.B, a normal operator's expected payoff in period 2 is 0.237 ($= 0.288 - (0.111 - 0.060)$) if unrevealed, 0.024 ($= 0.145 - 0.111 - 0.010$) if revealed and reform succeeds, and 0.084 ($= 0.024 + 0.060$) if reform fails. Thus, the operator's expected payoff in period 2 is 0.237 if $o_1 = s$ versus 0.086 ($= 0.237 \times 0.05 + (0.024 \times 0.1 + 0.084 \times 0.9) \times 0.95)$) if $o_1 = m$. This gap is large enough to deter the operator from defecting to $o_1 = m$ to capture the cost saving $c = 0.060$ in period 1, so the equilibrium is a reputation equilibrium.

Under Parameter Set II/owner management, $P_2(\rho) = 0.169$ and $P_2(r)$ is unchanged. Given these prices, following the arguments we just used to establish a reputation equilibrium under Parameter Set I, the operator's period 2 expected payoff is 0.118 if $o_1 = s$ and 0.080 if $o_1 = m$. The cost of revelation implied by these payoffs is too small to make up for the 0.060 the owner-operator can capture in period 1 by choosing $o_1 = m$, so owner management does not support a reputation equilibrium. Because the base level of trust is lower in Parameter Set II, this example illustrates the logic behind Proposition 1, which shows that high base levels of trust favor reputation equilibria. In Parameter Set II, the equilibrium under owner management is a mixed strategy equilibrium in which the owner-operator randomizes between operating policies in period 1. The expressions defining the operating strategies are readily derived from the condition for mixed strategy equilibria in Lemma 3 by substituting in the owner management payment structure. We provide the equilibrium condition and strategies in Lemma A-1 in Appendix A. This example highlights the contrasting implications of owner and professional management for firm reputation in a low trust environment.

4.4 Reform and reputation

To show how reform impacts reputation, we will now describe equilibrium conditions in a no-reform setting in which reform is not possible. Because reform can only occur if the operator is normal, along any history which includes reform, rational customers must believe that the operator is certainly normal, i.e., $\mu_2^* = 0$. This belief combined with the fact that $o_2 = m$ is the strictly optimal period 2 policy for the operator (Lemma 1) implies that, in the no-reform setting, after revelation, the period 2 good's price will equal $P_2(0) = \delta$ and the firm cannot possibly operate profitably in any equilibrium. Thus, to simplify the analysis of the no-reform setting we assume that, following revelation, the firm shuts down after the good's price is set in period 2. This truncates at p_2, all period 2 histories that include $q_1 = l$ and does not add any new histories.

Under owner management, shutting down the firm after the operator is revealed, from the perspective of the operator's payment structure, is equivalent to replacing the operator after revelation because the operator's payoff in either case is zero after revelation. As a result the conditions for equilibria under owner management switch to those in Part (b) of Lemma 3 after accounting for the fact that the owner-operator receives all the firm's profits.

Under professional management, the operator's expected payoff remains zero after revelation, as it is in the baseline setting. Thus, the optimal reputation-assuring compensation does not change when we switch to the no-reform setting, and the owner's expected payoff from picking reputation-assuring compensation continues to be represented by expression (12). Because reform is no longer possible, the owner's expected payoff from a no expected compensation policy changes from expression (9). Instead, the owner's expected payoff from choosing to make no payments to the professional operator will equal

$$P_2(\rho) - e + P_1(0)(P_2(\rho/P_1(0)) - e). \tag{14}$$

A little bit of algebra produces the following proposition.

Proposition 3 *In the no-reform setting*

a. *Under owner management, reputation equilibria exist if and only if*

$$(1-\delta)\Big(P_2(\rho) - [1-\pi_1]\Big) \geq \delta c. \tag{15}$$

b. Under professional management, an equilibrium with $\sigma^* = 1$ exists if and only if

$$P_2(\rho) \leq 1 - \frac{b_1^*}{\pi_1 + \delta}. \tag{16}$$

c. Under either owner management or professional management, the set of parameter values that support equilibria in which $\sigma^* = 1$ in the setting with reform is a proper subset of the set of parameter values that support equilibria in which $\sigma^* = 1$ in the no-reform setting.

A comparison of Proposition 3 with Propositions 1 and 2 shows that, under each governance structure, the firm's ability to sustain a reputation is greater in the no-reform setting. The underlying logic is straightforward: Reform lowers the owner's reputation cost from revelation. Hence, under professional management, reform shrinks the interval over which the owner is willing to bear the cost of reputation-assuring compensation. Under owner management, the profit generated via reform directly lowers the owner-operator's reputation cost. Moreover, a comparison of Part a of Proposition 3 with Proposition 2 shows that reputation in a professionally managed firms with the option to reform, a typical situation for many large public firms, can be quite different from that in traditional reputation models featuring owner-managed firms with no reform.

Because the move to the no-reform setting expands the set of parameters that supports reputable period 1 behavior, professional management continues to support reputable behavior under both parameter sets in Table 2.A, and owner management does the same under Parameter Set I. Under Parameter Set II switching to the no-reform setting makes the owner-operator's penalty for defecting from a reputation equilibrium sufficiently large to deter defection. Thus, once we switch to the no-reform setting, both governance structures support reputable behavior for both parameters sets in Table 2.A.

5 An experiment to examine the link between governance structure and reputation

In our model, professional management creates strikingly different incentives to maintain firm reputation than owner management. Importantly, under optimal governance policies, professional management better supports firm reputation in low initial trust environments. This is particularly true when reform is possible.

These differences reflect both the difference in the owner's information endowment and policy alternatives under two governance structures, as well as the different incentives faced by operators as illustrated by the example in Table 2. It is unclear (i) whether human owners of professionally managed firms will pick compensation contracts that deliver the same operator incentives as our model predicts, (ii) whether human operators will respond to the different incentives they face under the two governance structures and reform in the way our model predicts, and (iii) whether customers set prices that produce the operator incentives predicted by our model of Bayesian price-setting.

Adapting a design from previous studies, we conduct a laboratory experiment focused on issues (ii) and (iii). While this design choice narrows the scope of our experiment relative to that of our model, it permits us to test our model's predictions about (a) differences in operator behavior arising from reputation cost differences for operators under owner and professional management *that exactly match the differences* *in our model*; (b) whether outsiders (customers) recognize the differences in operator incentives; (c) the penalties customers impose for revelation; and (d) the effects of reform.

We run an owner management with reform treatment and professional management with reform treatment for each parameter set in Table 2.A. Each treatment consists of two sessions and, in each session, operators' period 2 payoffs are fixed according to Table 2.B so that, under each governance structure, operators face the same reputation cost tradeoff as in our model. To assess the effect of reform under each parameter set, we run an owner management treatment in which reform was not possible. For this treatment, in accordance with our model, we fix the operator's payoffs at zero after revelation and those in Table 2.B if the operator remains unrevealed. We do not run professional management without reform treatments because, as we have shown, the optimal contract, the operator's reputation cost and optimal policy do not change when we switch to the no-reform setting.

Our design is effectively a 2 (parameter set) $\times$ 2 (governance structure) $\times$ 2 (reform setting) design relative to operator choices and goods' prices. For each parameter set, by comparing OMR with OM treatments we can assess the impact of reform, by comparing OMR with PMR treatments we can assess the effect of governance structure, and by comparing OM with PMR treatments we can assess whether professional management mitigates the reform-induced erosion of operator incentives to act reputably. Table 3 summarizes our design. We will use the labels in the table to identify treatments, dropping the parameter set suffix when the discussion applies to both parameter sets.

Table 3. *Experimental treatments.*

	Parameter Set I			Parameter Set II		
Governance structure	Label	Sessions	Obs.	Label	Sessions	Obs.
Owner Management	OM-I	2	264	OM-II	2	264
Owner Management w/ Reform	OMR-I	2	264	OMR-II	2	264
Professional Management (w/ or w/o Reform)	PMR-I	2	264	PMR-II	2	228

5.1 Subject pool, instructions, and payments

Subjects in our experiment were drawn from a volunteer pool of undergraduate business and MBA students at the University of Iowa. There were 18 to 24 subjects in each session. Sessions lasted at most two hours and subjects were paid a $5 show-up fee. The experimental currency was "francs," which were converted to dollars at known exchange rates (depending on the treatment and subject's role). Subject payments (including the show-up fee) ranged from $10.31 to $32.00. They averaged $21.45 with a standard deviation of $4.37. Expected profits across roles were equalized by allowing customers to keep some of an endowment they received each period and setting different exchange rates for operators across the treatments.

The experiments were programmed in Z-Tree (Fischbacher, 2007). Upon arrival, subjects sat at separate computer terminals and received instruction sets (provided in Appendix B), forms to record profits by period, and receipts to be filled in during the session. The instructions clearly explained the choices that subjects and the experimenter would make in the experiment. The instructions also clearly showed the implications of these choices for subject payoffs. The instructions were read aloud and all questions were answered in public before each session.

At the beginning of each session, each subject was randomly assigned a role (“Green” player (customer) or “Blue” player (operator)) and remained in their roles throughout the session.[19] Each session consisted of 24 “rounds” and each round consisted of two-periods to match our model. Table 4 describes the sequence of events in each round of the experiment, with “|”s separating differences across ownership treatments and “*”s indicating unique features of our experiment that we will describe shortly.

Table 4. *Flow of each two-period game in the experiment.*

<table>
<tr><th rowspan="2">Period</th><th rowspan="2">Mover</th><th colspan="2">Action/Outcome</th></tr>
<tr><th>Owner Management</th><th>Professional Management</th></tr>
<tr><td rowspan="5">1</td><td>Experimenter</td><td colspan="2">Determines whether operator is committed or normal by setting production methods.</td></tr>
<tr><td>Stage I: Operator</td><td colspan="2">Chooses operating policy.</td></tr>
<tr><td>Stage II: Customer</td><td colspan="2">Sets price according to a modified BDM procedure.*</td></tr>
<tr><td>Experimenter</td><td colspan="2">Randomly determines item quality if operator chose mixed policy. Reveals item quality.</td></tr>
<tr><td></td><td colspan="2">Customer paid based on good’s quality, purchase price, and whether the good is purchased.
Operator paid firm’s period 1 profit. | Operator paid 0 or c based on policy.</td></tr>
<tr><td>2</td><td>Experimenter</td><td colspan="2">Randomly determines whether reform succeeds after revelation if reform is available.
Replaces revealed operator in no-reform treatments.* | Replaces revealed operator.*</td></tr>
<tr><td rowspan="4">2</td><td>Stage I: Operator</td><td colspan="2">Uses equilibrium policy.*</td></tr>
<tr><td>Stage II: Customer</td><td colspan="2">Sets price according to a modified BDM procedure.*</td></tr>
<tr><td>Experimenter</td><td colspan="2">Randomly determines item quality if operator chose mixed policy. Reveals item quality.</td></tr>
<tr><td></td><td colspan="2">Customer paid based on good’s quality, purchase price, and whether the good is purchased.
Operator paid according to Table 2.B.</td></tr>
</table>

* denotes a unique design feature explained later.

At the start of each round, subjects were randomly assigned to groups consisting of one customer and one operator (with random re-matching to start each new round). Then each operator was randomly assigned a set of operating policies or “methods,” which effectively fixed the operator’s type. A committed type could only use a safe policy (called “Method 1: Sure” in the experiment) that always produced a high quality good (a “round item” in the experiment) worth 1,000 francs to the customer. A fraction of operators (7/8 or 3/4 depending on the treatment) were assigned a normal type and could alternatively employ a mixed policy (called “Method 2: Mixed” in the experiment) that cost 60 francs less but produced a high quality good only 5% of the time and a low quality good (a “square” item worth 0 francs to the customer) 95% of the time. All subjects knew the assignment rules and fractions of each operator type. While each operator knew his/her type, customers did not. Customers could only learn the operator’s type indirectly if the operator produced a low quality item and, thus, was revealed.[20]

[19]While we refer to subjects as customers and operators, following the standard in the literature (see e.g., Plott, 1982), these terms were not used during the experiment to avoid value-laden connotations. The neutral language allows us to test our hypotheses without bringing in external preconceptions and framing effects that can make preferences deviate from those induced by the payoffs and contracts in the game. For example, if we refer to producing a low-quality good as “reputation damage,” subject behavior may be influenced by their experiences with reputation damage in other contexts in unpredictable ways, which can complicate strategic decision making by other subjects.

[20]Throughout each session, computer screens at the front and sides of the room displayed (1) the probability the Blue players

In each period, each operator made a production policy choice and the paired customer set the price for the good produced by the operator. Customers could not observe operators' choices. Prices were set using the modified Becker et al. (1964) procedure that we will describe shortly.

If the operator picked the safe policy the good's quality for the period was high. If a normal operator picked the mixed policy, the quality of the good for the period was determined via a random draw conducted by the experimenter with a 5% probability of high quality. At the end of each period, the good's quality was revealed to the two paired subjects, but not others.

If the first period good was low quality (the operator was revealed), then, in treatments with reform, the experimenter determined the outcome of reform via a random draw with a 10% probability of success. The outcome was disclosed to the operator, but not the customer. In the OMR treatments, the operator remained with the firm in period 2 after reform. In the PMR treatments, the operator was effectively replaced. (We will describe how this was implemented shortly.) In the no-reform OM treatments, a revealed operator's period 2 payoff was set to zero, consistent with the firm shutting down.

In period 1, the operator received the firm's profit in the OM and OMR treatments. The operator was paid zero or c in the PMR treatments if the chosen policy was safe or mixed, respectively. In period 2, the operators in treatments with reform (OMR and PMR) were paid according to Table 2.B. In the no-reform OM treatments, revealed operators received zero payoffs in period 2. Customers' payments in each period were determined according to the BDM procedure: they received the value of the goods they purchased less the price they paid.

5.2 Unique experiment features

Our experiment has several unique features. They were clearly explained to subjects in the instructions (shown in Appendix B). Before turning to the features we flagged in Table 4, we describe the one that is not flagged: ensuring salience (i.e., that subjects' payoffs vary meaningfully with their choices (e.g., Plott, 1982)). Each cash flow in Table 2 was scaled up by a factor of 1,000, so a cash flow of 1 in the model equaled 1,000 francs in the experiment. To create a meaningful difference across period 1 production policy choices, we further scaled up the bonus b^* from 3.15 to 30.[21] The added scaling does not change optimal actions: The scaled up bonus continues to provide the operator a strong incentive to pick the safe policy in period 1. While the owner's compensation cost rises, because the owner's gain from a reputation-assuring bonus is sufficiently large under both parameter sets in Table 2, it is optimal for the owner to pay the scaled up bonus and ensure that the operator acts reputably.

5.2.1 The procedure for setting goods' prices

Since embedding a competitive goods market would be time consuming and costly, we follow Noe et al. (2012) and use a modified Becker et al. (1964) BDM procedure to set goods' prices. This procedure requires

(operators) were a committed or normal type, (2) the probability that each policy produces each item type, and (3) the period 2 equilibrium operating policy conditioned on whether the operator was revealed or not.

[21] Setting the bonus to 30 also roughly equalizes the expected cost of managerial defection from the equilibrium supporting reputable behavior across treatments. The bonus of 30 leads to a net expected cost of 25.50 when a professional manager defects from the equilibrium. This is comparable to the 22.24 expected cost to defection for an owner-manager under Parameter Set II.

just one customer to simulate a competitive price and is designed to elicit the highest price a customer is willing to pay for a good.[22]

In the BDM procedure, the customer first specifies the most she is willing to pay. Once this "limit price" is set, the experimenter randomly draws a "discounted price" between 0 and 1,000 francs from a uniform distribution. The experimenter then buys the good from the firm at the limit price. The experimenter resells the good to the customer at the discounted price only if the limit price exceeds the discounted price. Otherwise, the experimenter keeps the good. When a customer purchases the item, she receives a payoff equal to her endowment plus the value of the item minus the discounted price. When the experimenter keeps the good, the customer receives only her endowment. In either case, she receives an ex post report showing the quality of the item and the discounted price. In the following analysis, we report customer limit (a.k.a., "established") prices.

5.2.2 Period 2 operator strategies

While normal operators were free to pick either policy in period 1, they were required to use the mixed policy in period 2 of each round. This is the optimal strategy as described in Lemma 1. This simplification, which was known to all subjects, allowed operators to focus all their attention on their period 1 choices. This is the only determinant of reputation in our setting. It also reduced strategic uncertainty faced by customers in period 2, which allows us to more cleanly assess the reason for price drops following revelation.

This simplification follows a long history of experimental research that fixes some acts or decisions in games to correspond to specific (often equilibrium) strategies to isolate and study other decisions that are of specific interest (e.g., Lieberman, 1962; Messick, 1967; Fox, 1972; Shachat and Swarthout, 2008). By reducing strategic uncertainty it should promote convergence of behavior (Lucas, 1986; Van Huyck et al., 1990). By lowering subject decision costs, it should lower subject errors (Smith and Walker, 1993). Moreover, in prior reputation experiments with settings like ours, experienced subjects nearly always follow the opportunistic strategy in the final period when there is no incentive to avoid opportunism (e.g., Brandts and Figueras, 2003, p. 96). Therefore, we expect that our subjects would have generally behaved opportunistically in (the final) period 2, just as subjects in past experiments have done.[23]

5.2.3 Revealed operator replacement and owners under professional management

In our model under professional management, the owner picks the governance policy and a revealed operator is replaced. In our experiment, governance policy is fixed according to Table 2.B. So, the owner's one remaining decision in the PMR treatments is choosing a replacement period 2 operator. In the experiment, the experimenter took on the role of the replacement operator and, matching our model, no payment was made to either replaced or replacement operators. This simplification effectively eliminated the need for additional subjects to play owners and period 2 replacement operators.

[22]On average, the BDM procedure elicits risk neutral valuations (Berg et al., 2005) which correspond to competitive prices in our context. The procedure was fast to implement, it avoided complications from auction procedures (e.g., overbidding, as in Kagel and Levin, 1993), and did not require pre-specifying a limited set of allowable prices (e.g., Forsythe et al., 1999). Recent experiments on financial decision making that use the BDM procedure include Frydman and Nave (2017) and Fuster et al. (2020).

[23]We also note that subjects in our experiment were quite opportunistic even in period 1, when they had an incentive to avoid opportunism.

In the no-reform OM treatments, a revealed operator's period 2 payoff was set to zero. This allows us to observe an integral part of the incentive for reputable behavior: the prices customers are willing to pay for period 2 goods produced by revealed but unreformed firms. After an operator was revealed the experimenter set the (dominant) mixed policy in period 2 and customers priced the period 2 goods.

Because normal operators were required to use the mixed policy in period 2 in the experiment, they made no strategic choices in period 2. These simplifications come at little cost but yield several benefits. First, in all treatments, our design matches prior experiments based on the KWMR reputation framework in that subjects only play two roles—insider (operator) and outsider (customer)—and take only one type of action in each role.[24] Second, the design reduces strategic uncertainty and equalizes it across treatments. This reduces complexity and lowers subject decision costs, which should make it easier to diagnose differences between model predictions and experimental behavior (Davis and Holt, 1993). Third, we can implement the professional and owner management treatments in essentially the same setting. This maximizes the likelihood that governance structure-based incentive differences drive outcomes rather than differences in the design of the treatments or subject instructions (Zizzo, 2010).

6 Evidence from our experiments

Recall that our model predicts reputable behavior in period 1 under both professional and owner management, both with and without the option to reform, under Parameter Set I. Thus, the treatments using Parameter Set I are effectively placebos: we should observe no differences across treatments. Under Parameter Set II, our model predicts a reduced frequency of Period 1 reputable behavior only in treatment OMR-II, when reform is possible and the firm is owner-managed. In ORM-II, the model predicts the operator will randomize between policies, reducing the predicted frequency of reputable behavior. This should decrease period 1 prices relative to other treatments. Period 2 prices should depend on the parameter set, whether the operator has been revealed, and whether reform is possible. Table 5 summarizes the equilibrium predictions.

Table 5. *Equilibrium predictions for parameter sets in Table 2.A* .

	Parameter Set I		Parameter Set II		
	OM-I	OMR-I & PMR-I	OM-II	PMR-II	OMR-II
Probability of period 1 reputable behavior	1.000	1.000	1.000	1.000	0.802
Period 1 Price	1,000	1,000	1,000	1,000	835
Period 2 Price ǀ Unrevealed	288	288	169	169	192
Period 2 Price ǀ Revealed	50	145	50	145	145

Prior research strongly suggests that subject behavior in experiments will frequently deviate from predictions of reputation models (e.g., Camerer and Weigelt, 1988; Neral and Ochs, 1992; Brandts and Figueras, 2003; Noe et al., 2012). Brandts and Figueras (2003) suggest that the deviations increase when games are

[24]Except for the number of periods in the game and the reform option, our experiment is comparable to the experiments in Camerer and Weigelt (1988), Neral and Ochs (1992), Brandts and Figueras (2003), and Noe et al. (2012). In Camerer and Weigelt (1988), Neral and Ochs (1992), and Brandts and Figueras (2003) outsiders make lending decisions which serve as an *indicator of sufficient reputation* rather than set a precise measure, a good's price, as they do in our experiment and in Noe et al. (2012).

shorter. Thus, we do not expect subject behavior to conform exactly to our two-period (i.e., shortest possible) reputation game predictions. However, we expect the predictions about reputation effects of governance structure and reform to hold, *at least qualitatively*, *in spite of* expected behavioral deviations.

Therefore, for each parameter set, we will compare the OMR with OM treatments to assess the impact of reform, compare OMR with PMR treatments to assess the effect of governance structure, and compare OM with PMR treatments to assess whether professional management mitigates reform-induced erosion of incentives to act reputably. By examining how outcomes vary across the two parameter sets we will assess whether the interaction between the level of outsider trust, governance structure, and reform impacts operator incentives as we expect. In summary, we will focus on addressing the following hypotheses that reflect our expectations:

Hypothesis 1 In period 1, normal operators will use the safe operating policy at similar rates in all the placebo treatments using Parameter Set I. In treatments using Parameter Set II, normal operators will use the safe policy more frequently in period 1 in treatments PMR-II and OM-II than in treatment OMR-II.

Hypothesis 2 Customers will set similar prices for period 1 goods in all the placebo treatments using Parameter Set I. In treatments using Parameter Set II, period 1 goods' prices will be higher in treatments PMR-II and OM-II compared with treatment OMR-II.

Hypothesis 3 Customers will set similar prices for period 2 goods from *unrevealed operators* in all the placebo treatments using Parameter Set I. In treatments using Parameter Set II, period 2 goods' prices from *unrevealed operators* will be lower in treatments PMR-II and OM-II compared with treatment OMR-II.

Hypothesis 4 Customers will set lower period 2 prices for goods from revealed operators in all treatments. For revealed operators, period 2 goods' prices will be higher in treatments with reform than in no-reform treatments.

Hypothesis 5 When we compare treatments using different parameter sets, but the same governance structure and reform protocol, we expect:

a. similar rates of use of the safe policy in period 1 in PMR and OM treatments but lower usage rates in treatment OMR-II compared with OMR-I,

b. similar period 1 goods' prices in PMR and OM treatments but lower period 1 goods' prices in treatment OMR-II compared with OMR-I.

We note that, while we do not expect reform to affect professional manager incentives in our experiment, we expect reform to erode the incentives for owner-operators to maintain reputations under both parameter sets. Under Parameter Set II, we expect the erosion to be sufficient to change the outcomes. While we do not expect the erosion to be sufficient to change outcomes under Parameter Set I, we anticipate that the reduced incentives to maintain reputations may affect observed behavior.

We will describe the evidence from our experiment in the following order: operator policy choices in period 1, period 1 goods' prices, and period 2 goods' prices. To simplify the exposition, we will refer to operators' period 1 policy choices, simply as policy choices.

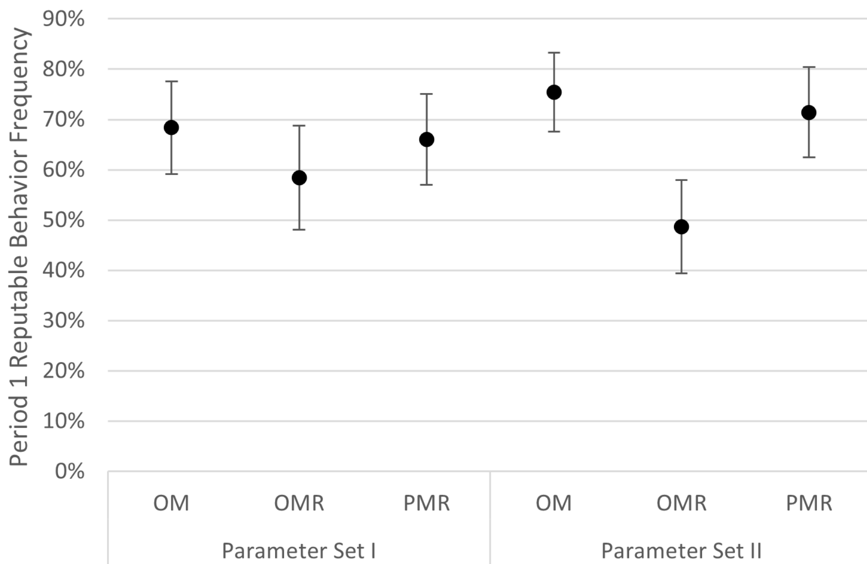

Figure 2. *Frequency of reputable behavior in the second half of each session.* This figure presents the frequency with which normal-type operators make reputable policy choices (in period 1).

6.1 Adoption of the reputable policy

The period 1 operating policy indicates whether firms behave reputably. We expect the operators who choose these policies to respond to monetary incentives. These incentives are tied to firms' governance structures. The availability of reform matters as well. Therefore, we expect period 1 operating policies to vary systematically with governance structure and reform protocol according to Hypotheses 1 and 5.

Figure 2 shows the percentage of times normal-type operators make reputable policy choices by treatment. The figure also shows 95% confidence interval bars. To limit confounding effects arising from subject learning, the figure shows prices in the second half of the experiment.[25]

Normal operators often choose reputable policies, but do so significantly less often than predicted in equilibrium. In Treatment OMR-II, the equilibrium predicts that normal operators will choose reputable policies about 80% of the time. However, normal operators pick the safe policy about 50% of the time in our experiment. In the remaining treatments, equilibrium predicts reputable policy choices 100% of the time. We observe reputable policy choices 73% of the time (in OM-II) or less.

In comparison, prior experiments, which closely match our OM treatments, also produce low rates of reputable behavior when the equilibrium prediction is 100% reputable behavior. For example, in Noe et al. (2012) the frequency of reputable behavior is around 60%. In longer games, Camerer and Weigelt (1988) report reputable behavior with a frequency of under 80% when subjects are relatively inexperienced. Thus, while reputable policies are chosen in our experiment with a low frequency, this is not unexpected. Moreover, in Appendix C, we show that, over time, operators in our experiment gravitate toward choosing reputable policies when our model predicts they will.

[25]Using the entire data set leaves these results as well as the remaining results we report in this section essentially unchanged. Because we find evidence that subjects learn about the game through experience, we are more comfortable using statistical tests based on the later periods in the experiment. In Appendix C we provide insights into the impact of subject learning.

Consider differences across treatments, which speak to our hypotheses about the impact of governance structure and reform. Table 6 presents the frequencies with which normal operators make reputable policy choices in each treatment along with formal tests of differences across treatments. To minimize the confounding effects of subject learning, we present and compare policy choices in the second half of each session (see footnote 25). The observed variation across treatments is consistent with our model's predictions about the effect of governance structure, and the impact of reform in weakening commitment to maintain stakeholder trust, particularly under owner management.

Result 1 The variation in reputable policy choices across treatments conforms with Hypotheses 1 and 5 about the effects of governance structure and reform, i.e., operators respond in the expected manner to the different incentives for reputable behavior created by governance structures and the opportunity to reform.

Table 6. *Incidence of reputable policy choices.* In this table we present the frequency with which normal-type operators choose the reputable policy (in period 1) during the second half of each session. ***, **, and * denote t-test significance at the 99%, 95%, and 90% levels of confidence respectively. DoF indicates Degrees of Freedom = the number of combined observations - 2.

		Parameter Set		Parameter Set I vs II			
Governance structure		Parm. I	Parm. II	Difference	t-Stat.	p-Value	DoF
OM		68.37%	75.42%	0.0706	1.15	0.25	214
OMR		58.43%	48.67%	-0.0975	-1.38	0.17	200
PMR		66.04%	71.43%	0.0539	0.83	0.41	202
OM	Difference	-0.0994	-0.2675***				
vs	t-Stat.	-1.41	-4.35				
OMR	p-Value	0.16	0.00				
	DoF	185	229				
OM	Difference	-0.0233	-0.0400				
vs	t-Stat.	-0.35	-0.66				
PMR	p-Value	0.73	0.51				
	DoF	202	214				
OMR	Difference	0.0761	0.2276***				
vs	t-Stat.	1.09	3.43				
PMR	p-Value	0.28	0.00				
	DoF	193	209				

We expect the incentives enforcing reputable behavior to be weaker under owner management with reform under Parameter Set II than in the other treatments. Hence, we expect that operators under OMR-II will behave less reputably than in other treatments, while behavior in other treatments should be similar to each other (Hypothesis 1). The results are largely consistent with our expectation. We find no significant differences between the frequencies of reputable period 1 policy choices between the placebo treatments using Parameter Set I. In treatment OMR-I, the frequency of reputable policy choices is (58.4%) which is not statistically significantly different than the frequencies in treatments OM-I (68.4%) or PMR-I (66.0%). Thus, as expected, governance structure and reform do not significantly impact the frequency of reputable behavior when the base level of trust is $\rho = 0.250$.

For the treatments using Parameter Set II, behavior is consistent with Hypothesis 1. Operators act reputably only 48.7% of the time in treatment OMR-II, which is statistically significantly lower than in treatment OM-II (t=4.35, p-value=0.00) and in treatment PMR-II (t=3.43, p-value=0.00). Hence, reform erodes incentives for reputable behavior under owner management, whereas professional management is more supportive of reputable behavior when the base level of trust is lowered to $\rho = 0.125$.

We expect operators to behave similarly across Parameter Sets I and II in OM and PMR treatments, but reputable behavior should fall in OM-II relative to OM-I (Hypothesis 5). We find no significant difference between treatments PMR-I and PMR-II, or between OM-I and OM-II. While there is reduced use of the safe policy in treatment OMR-II relative to OMR-I (48.67% versus 58.43%) as we expect, the change is not significant.

Neither governance structure nor the option to reform impacts the tendency to act reputably across our placebo treatments using Parameter Set I, while reputable behavior is less frequent in OMR-II that in OMR-I and the other treatments using Parameter Set II. This pattern of period 1 policy choices matches our expectation and shows the impact of the initial level of trust, reform, and governance structure we expect. As initial outsider trust ρ falls, operator incentives that benefits owners under professional management are more likely to result in reputable policy choices than operator incentives under owner management. Moreover, reform tends to undermine the ex ante incentive for owner-operators to make policy choices that maintain customer trust in a low trust environment while professional management mitigates the reform-induced erosion of incentives to act reputably.

6.2 Period 1 goods' prices

Period 1 goods' prices represent customer conjectures about period 1 quality. They measure firm reputation. The gap between period 1 goods' prices and $P_1(0)$, the price customers set in our model when they expect the normal operator to act opportunistically, is a measure of the reward firms receive for the level of reputable period 1 behavior customers anticipate. We expect firm reputation and period 1 prices to vary systematically with governance structure and reform according to Hypotheses 2 and 5.

Under each treatment, Figure 3 illustrates period 1 prices ($\times$s), average actual values of the goods ($\bullet$s), and the predicted values of the goods ($+$s). The figure also shows 95% confidence interval bars for goods' prices and actual values. Table 7 shows tests for differences between period 1 prices and (1) $P_1(0)$, (2) predicted prices, and (3) actual values. To limit confounding effects arising from subject learning, Figure 3 and Table 7 are based on prices in the second half of the experiment (see footnote 25).

Average period 1 prices are well below levels predicted by our model (835 francs in treatment OMR-II and 1,000 francs in the remaining treatments). T-statistics from formal tests of these differences range from -9.52 to -13.78. Thus, it appears that customers anticipated the lower than predicted level of reputable operator behavior in the experiment.

These discounts relative to predicted period 1 prices are not unprecedented nor unanticipated. They mirror outsider responses in prior experiments where the insiders are predicted to act reputably 100% of the time but fail to do so. For example, in the first period of the Noe et al. (2012) experiment, which closely matches our owner management treatment without reform, the average price for goods is 651 francs compared to the predicted price of 1,000 francs. While outsiders' willingness to lend in Camerer and

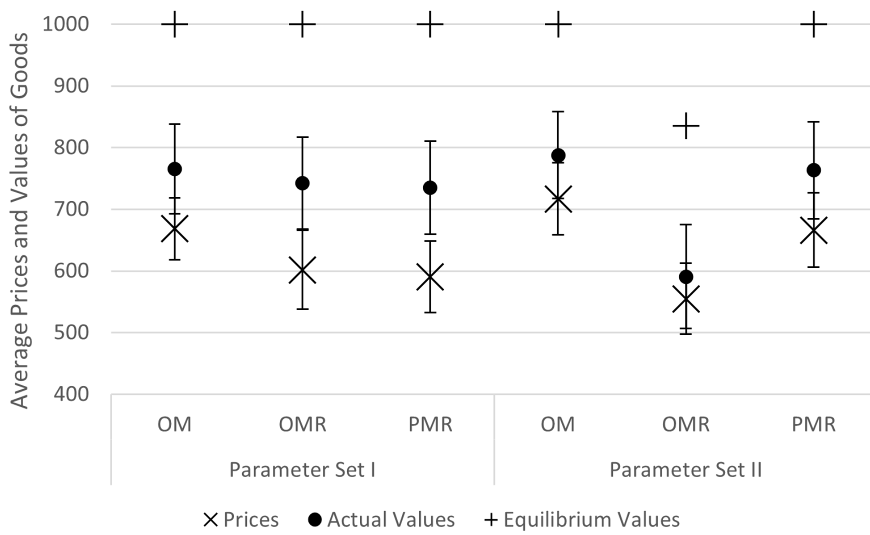

Figure 3. *Period 1 values and limit prices in the second half of each session.*

Weigelt (1988) serves as an indicator of sufficient reputation rather than a precise reputation measure like goods' prices, as few as 70% of lenders make loans compared to a 100% predicted rate in early periods of their experiment.

Comparing prices to $P_1(0)$ in the experiment, customers apparently anticipate reputable behavior from normal operators a significant fraction of the time. Specifically, period 1 prices are significantly higher than $P_1(0)$ for both Parameter Set I (287.5 francs) and Parameter Set II (168.75 francs), as shown in Table 7 (with t-statistics ranging from 9.64 to 18.41). Notably, the prices customers set in the experiment are high enough to make firms profitable in period 1.

However, customers generally under-price period 1 goods, suggesting that they expect operators to behave more opportunistically than they actually do. Average underpricing ranges from 36.16 francs in treatment OMR-II to 144.38 in treatment PMR-I. Notably, these differences are smaller in absolute terms than the differences with predicted prices or $P_1(0)$. The differences are also of less, if any, statistical significance. In treatments using Parameter Set II the difference between period 1 prices and the value of goods is either not statistically significant (OM-II and OMR-II) or only marginally significant (PMR-II). Only in the treatments using Parameter Set I do we find consistent evidence of a significant difference. The relatively small differences between prices and the value of goods suggest that customers anticipate actual operator behavior relatively well.

To formally examine Hypotheses 2 and 5, and assess the effects of governance structure and reform on prices, we compare period 1 goods' prices across treatments in the second half of the experiments (see footnote 25). The results are presented in Table 8.

We expect prices to be lower in under owner management with reform in Parameter Set II than in the other treatments, while period 1 prices in other treatments should be similar to each other (Hypothesis 2). For the placebo treatments using Parameter Set I, we expect the period 1 good's price to be unaffected by governance structure or reform. Consistent with this expectation, prices in treatment OMR-I (601.7

Table 7. *Customer conjectures of period 1 good quality versus $P_1(0)$, equilibrium prices, and values.* In this table we present period 1 prices for goods in the second half of the experiment. ***, **, and * denote t-test significance at the 99%, 95%, and 90% levels of confidence respectively. DoF indicates Degrees of Freedom = the number of combined observations - 2.

	Governance Structure and Parameter Set					
	OM-I	**OMR-I**	**PMR-I**	**OM-II**	**OMR-II**	**PMR-II**
Mean Est. Price	668.27	601.66	590.47	716.73	554.75	666.39
Equilibrium Price	1000	1000	1000	1000	853.3	1000
Mean Est. Price minus Eq. Price	-331.73	-398.34	-409.53	-283.27	-298.55	-333.61
t-Stat.	-13.07***	-12.22***	-13.78***	-9.52***	-10.17***	-10.85***
p-Value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
DoF	131	131	131	131	131	113
$P_1(0)$	287.5	287.5	287.5	168.75	168.75	168.75
Mean Est. Price minus $P_1(0)$	380.77	314.16	302.97	547.98	386	497.64
t-Stat.	15.00***	9.64***	10.19***	18.41***	13.14***	16.18***
p-Value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
DoF	131	131	131	131	131	113
Mean Value	765.15	742.42	734.85	787.88	590.91	763.16
Mean Est. Price minus mean Value	-96.88	-140.76	-144.38	-71.15	-36.16	-96.77
t-Stat.	-2.03	-2.76	-3.08	-1.63	-0.7	-1.82
p-Value	0.0447**	0.0067***	0.0025***	0.1059	0.4864	0.0709*
DoF	131	131	113	131	131	113

francs) are not significantly different from prices in OM-I (668.3 francs) or PMR-I (590.5 francs). However, contrary to our expectation, the difference between prices in OM-I and PMR-I of 77.8 francs is statistically significant at the 95% level of confidence.

Under Parameter Set II, as we expect, prices do not differ significantly between treatments OM-II and PMR-II. Moreover, as we expect, the average price in OMR-II is 554.8 francs, which is statistically significantly lower than the average price of 716.7 francs in OM-II (t=3.87, p-value=0.00) as well as the average price of 666.4 francs in PMR-II (t=2.62, p-value=0.01).

Across parameter sets within a treatment, we only expect to see a difference in period 1 prices between treatments OMR-I and OMR-II (Hypothesis 5). Consistent with this, prices in treatments OM-I and OM-II do not differ significantly. However, contrary to expectations, prices in treatment PMR-I are significantly lower than in PMR-I. While prices in treatment OMR-I are higher than in treatment OMR-II the difference is not statistically significant.

While period 1 prices in PMR-I are lower than we expect in comparison to other treatments, the remaining price differences across treatments tend to match our expectations, particularly when comparing treatments using the same parameter set. Under Parameter Set I, customers appear to recognize that neither reform (OM-I vs OMR-I) nor governance structure alone (OMR-I vs PMR-I) markedly erodes incentives for reputable behavior. Under Parameter Set II, they anticipate both effects to be consequential and also anticipate that professional management can mitigate the undesirable effects of reform. We formally state this in the following result:

Result 2 The variation in period 1 prices across treatments generally conforms with Hypothesis 2 about customer behavior, i.e., customers do anticipate that governance structure and reform affect incentives to act reputably in a manner that aligns with our expectations.

Table 8. *Customer conjectures of period 1 good quality.* In this table we present period 1 prices for goods in the second half of the experiment. ***, **, and * denote t-test significance at the 99%, 95%, and 90% levels of confidence respectively. DoF indicates Degrees of Freedom = the number of combined observations - 2.

		Parameter Set		Parameter Set I vs II			
Governance structure		Parm. I	Parm. II	Difference	t-Stat.	p-Value	DoF
OM		668.27	716.73	48.46	1.24	0.22	262
OMR		601.66	554.75	-46.91	-1.07	0.29	262
PMR		590.47	666.39	75.92*	1.77	0.08	244
OM	Difference	-66.61	-161.98***				
vs	t-Stat.	-1.61	3.87				
OMR	p-Value	0.11	0.00				
	DoF	262	262				
OM	Difference	-77.80**	-50.35				
vs	t-Stat.	-1.99	-1.17				
PMR	p-Value	0.05	0.24				
	DoF	262	244				
OMR	Difference	-11.19	111.64***				
vs	t-Stat.	-0.25	2.62				
PMR	p-Value	0.80	0.01				
	DoF	262	244				

6.3 Period 2 goods' prices

Period 2 prices reflect customer trust after observing the quality of period 1 goods and reform. The gap between the period 2 prices following high and low quality period 1 goods is a key determinant of the reputation cost for the firm's owner. We expect the gap between these prices to be smaller when firms can reform. We also expect the gap to depend on what customers infer from the period 1 goods' quality when they do not expect operators to behave reputably in period 1 as reflected in Hypotheses 3 and 4.

Figure 4 shows period 2 prices established by customers ($\times$s) , goods' actual values ($\bullet$s), and goods' predicted values ($+$s) in the second half of each session (see footnote 25). Panel A shows prices when the operator is unrevealed, and Panel B shows prices after revelation, both after reform and when reform is not possible. The figure also shows 95% confidence interval bars for prices and goods' actual values.

Period 2 prices tend to be higher than both predicted values and goods' actual values. Consider the difference between predicted values and prices. For unrevealed firms, the average gap ranges from 124 to 332 francs and is highly statistically significant with the exception of PMR-I. In PMR-I, the gap is only 39 francs. Prices for goods from revealed firms are also uniformly higher than predicted values. With the exception of OM-I and OM-II where the gap is close to 50 francs, the gap ranges from 123 to 141 francs and is statistically significant. In contrast to period 1, customers apparently expect operators to behave *more*

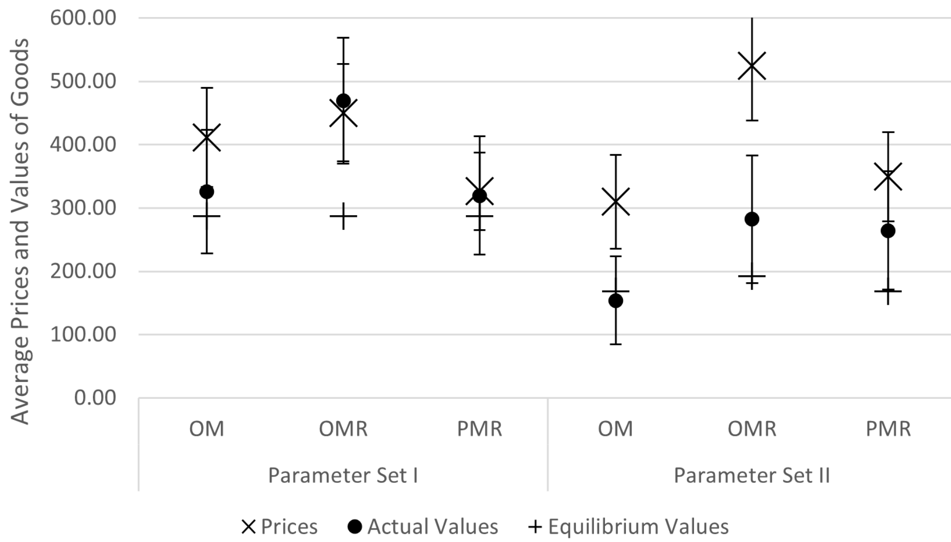

A. Unrevealed operators

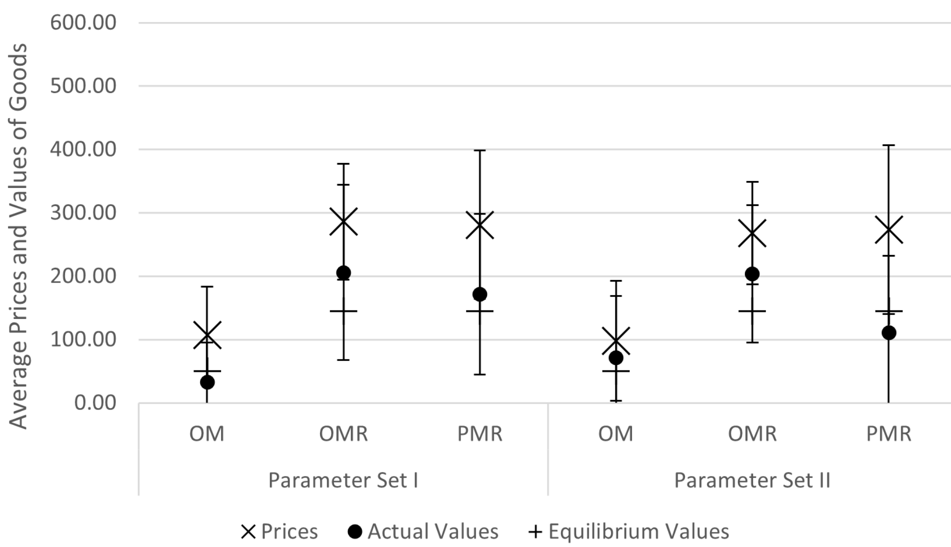

B. Revealed operators

Figure 4. *Period 2 values and limit prices established by customers in the second half of each session.*

reputably than predicted in period 2.

Next, compare prices and goods' actual values.[26] For unrevealed firms, average prices exceed average values except in OMR-I, where goods are underpriced by 19 francs. The overpricing is most marked and is statistically significant in OM-II (156 francs) and OMR-II (243 francs). Otherwise it ranges from 7 to 86 francs. Goods from revealed firms are also uniformly overpriced. The overpricing ranges from 26 to 162 francs but is never statistically significant. Thus, customers tend to expect unrevealed operators to act more reputably than they actually do. However, customers tend not to make significant pricing errors once operators are revealed or firms have reformed.

Table 9. *The penalty for being revealed.* This table presents the differences between period 2 prices for revealed and unrevealed firms in the second half of the experiment as well as the predicted values of the difference. ***, **, and * denote significance at the 99%, 95%, and 90% levels of confidence respectively.

		Parameter Set I			Parameter Set II		
		OM	OMR	PMR	OM	OMR	PMR
Unrevealed Prices		411.60	450.45	326.25	309.75	524.68	349.28
Revealed Prices		106.71	286.06	281.00	97.86	268.17	273.26
Predicted Difference		-237.5	-142.5	-142.5	-118.75	-47.20	-23.75
Observations		120	132	132	132	132	114
Revealed vs	Difference	-304.89***	-164.39*	-45.25	-211.89**	-256.51***	-76.02
Unrevealed Prices	t-statistic	-4.26	-2.29	-0.71	-2.75	-4.06	-1.01
Actual Difference vs	Dif. vs Pred.	-67.39	-21.89	97.25	-93.14	-209.31***	-52.27
Predicted Difference	F-statistic	0.89	0.09	2.36	1.46	11.00	0.49

In Table 9, we compare period 2 prices for revealed and unrevealed firms to see if prices fall with revelation and reform mitigates the drop as we expect (Hypothesis 4). Prices are always higher for unrevealed firms. The difference ranges from 45 francs in PMR-I to 305 francs in OM-I. Only in PMR-I and PMR-II is the difference not statistically significant. Thus, as we expect, firms are generally rewarded by customers with higher goods' prices in period 2 when they are likely to have acted reputably. Interestingly, with the exception of PMR-I, the average reward in the experiment equals or exceeds that predicted by our model, though the reward is significantly larger only in treatment OMR-II.

Result 3 Consistent with Hypothesis 4, customers set higher period 2 goods' prices when operators remain unrevealed.

We present additional tests of Hypothesis 3 as well tests of Hypothesis 4 in Table 10 by comparing period 2 prices across treatments within each revelation state.[27] The pattern is generally consistent with the hypotheses though there are some unexpected differences.

First, consider unrevealed firms under Parameter Set I ("Unrevealed" column, second row in each comparison set). We expect that period 2 prices will be similar across the placebo treatments, OM-I, OMR-I,

[26]Because of relatively small sample sizes, the average realized values of goods in the experiment do not exactly match their expected values under the model parameters.

[27]We use the second half of the data for consistency with the analysis above. Every significant relationship in Table 10 remains significant and in the same direction when using all the data.

Table 10. *Customer responses across treatments and parameter sets for revealed and unrevealed firms.* In this table we present period 2 prices for goods in the second half of the experiment. ***, **, and * denote t-test significance at the 99%, 95%, and 90% levels of confidence respectively. DoF indicates Degrees of Freedom = the number of combined observations - 2.

			Revelation Status	
Treatment	Parameter Set		Unrevealed	Revealed
OM	I		411.60	106.71
	II		309.75	97.86
OMR	I		450.45	286.06
	II		524.68	268.17
PMR	I		326.25	281.00
	II		349.28	273.26
OM vs OMR	I	Diff.	38.85	179.35***
		t-stat	0.69	-2.92
		p-value	0.49	0.00
		DoF	185	63
	II	Diff.	214.93***	170.31***
		t-stat	3.71	-2.54
		p-value	0.00	0.01
		DoF	180	80
OM vs PMR	I	Diff.	-85.35*	174.29**
		t-stat	-1.70	2.37
		p-value	0.09	0.02
		DoF	184	64
	II	Diff.	39.53	175.40**
		t-stat	0.75	-2.12
		p-value	0.46	0.04
		DoF	189	53
OMR vs PMR	I	Diff.	-124.20**	-5.06
		t-stat	-2.48	-0.07
		p-value	0.01	0.95
		DoF	193	67
	II	Diff.	-175.40***	5.09
		t-stat	-3.11	0.07
		p-value	0.00	0.95
		DoF	163	79

and PMR-I (Hypothesis 3). While the difference between OM-I and OMR-I is not statistically significant, PMR-I has significantly lower prices than either OM-I (-85, p-value=0.09) and OMR-I (-124, p-value=0.01). Thus, only one of the three differences matches Hypothesis 3.

Second, consider unrevealed firms under Parameter Set II ("Unrevealed" column, first row in each comparison set). We expect period 2 prices will be higher in OMR-II than OM-II and PMR-II, while OM-II and PMR-II should have similar prices (Hypothesis 3). Indeed, OMR-II prices are 215 higher than OM-II and 175 higher than PMR-II prices. Both differences are statistically significant (p-values = 0.00 for both). The difference between OM-II and PMR-II is 40 and is not statistically significant. All three comparisons support Hypothesis 3.

Finally, consider the effect of reform for revealed firms ("Revealed" column). We expect reform to increase prices for revealed firms relative to revealed, but unreformed, firms (Hypothesis 4). This is indeed the case. Comparing revealed firms in treatment OM to treatments OMR and PMR (both of which allow reform), prices are always significantly higher with reform (170 to 179 higher with p-values all 0.04 or below). Further, there is no significant difference between the reform treatments (OMR and PMR). Thus, all three comparisons are consistent with Hypothesis 4.

Overall, the period 2 prices tend to align with our expectations. The only departures from our expectations involve the PMR treatments. While subjects overprice goods in the PMR treatments just like they do in the remaining treatments, the level of overpricing in the PMR treatments falls short of the levels in OM and OMR treatments. This could be an indication of less misplaced trust when firms are professionally managed.

Result 4 While there are some exceptions, comparisons of period 2 goods' prices across treatments and reform protocol support Hypotheses 3 and 4. The mixed evidence comes primarily from the professional management (PMR) treatments in which we observe lower levels of mispricing than in the remaining treatments.

7 Discussion

This paper considers the question of how governance structure—the identity of the agents setting operating policy—affects the sustainability of firm reputation with outsider stakeholders and their trust in firms' production standards. We characterize sequential Nash equilibria in a setting where stakeholders have imperfect information about the commitment of firms to reputable behavior. In this framework, we show that governance structure matters; both firms controlled by informed insider owners and firms controlled by uninformed outsider "common/universal" owners may opt for opportunistic policies promising short-term gains that potentially destroy stakeholder trust and firm reputation. However, the conditions for maintaining trust and reputation under the two governance structures are fundamentally different. Factors that favor trust/reputation sustainability under informed insider control reduce sustainability under uninformed outsider control and vice versa.

We also consider corporate reform and restructuring after detected opportunism and consider how reform affects the sustainability of trust/reputation. We show that, under both governance structures, the option to reform, by lowering the cost of losing trust/reputation, makes it more difficult for firms to commit to maintaining trust/reputation. However, uninformed outsider control, by separating the agents making the operating decisions that affect trust/reputation from the value engendered by reform, mitigates the adverse effects of corporate reform on trust/reputation sustainability.

Through a laboratory experiment, we investigate the robustness of our findings to behavioral deviations. In the experiment, responses to firm behavior are produced by laboratory subjects' choices rather than by Bayesian updating. Although subject behavior is generally inconsistent with forward-looking rational expectations, the qualitative conclusions of our model are nevertheless supported by the experiment. Thus, our experiment provides evidence for the external validity of our theoretical predictions under behavioral deviations.

References

Alexander, Cindy R, “On the nature of the reputational penalty for corporate crime: Evidence,” *Journal of Law and Economics*, 1999, *42* (S1), 489–526.

Amel-Zadeh, Amir, Fiona Kasperk, and Martin Schmalz, “Mavericks, Universal, and Common Owners-The Largest Shareholders of US Public Firms,” Technical Report, CESifo 2022.

Axios and Harris, “Axios-Harris Poll: America’s Top 100 Companies,” 2024. Accessed: 2025-08-18.

Ball, Sheryl B and Paula-Ann Cech, “Subject pool choice and treatment effects in economic laboratory research,” *Research in Experimental Economics*, 1996, *6* (3), 239–292.

Bar-Isaac, H. and S. Tadelis, “Seller reputation,” *Foundations and Trends in Microeconomics*, 2008, *4* (4), 273–351.

Barber, Brad M and Masako N Darrough, “Product reliability and firm value: The experience of American and Japanese automakers, 1973–1992,” *Journal of Political Economy*, 1996, *104* (5), 1084–1099.

Becker, Gordon M, Morris H DeGroot, and Jacob Marschak, “Measuring utility by a single-response sequential method,” *Behavioral Science*, 1964, *9* (3), 226–232.

Bénabou, Roland and Jean Tirole, “Incentives and prosocial behavior,” *American Economic Review*, 2006, *96* (5), 1652–1678.

Bennett, Roger and Rita Kottasz, “Practitioner perceptions of corporate reputation: an empirical investigation,” *Corporate Communications*, 2000, *5* (4), 224–235.

Berg, Joyce, John Dickhaut, and Kevin McCabe, “Trust, reciprocity, and social history,” *Games and Economic Behavior*, 1995, *10* (1), 122–142.

_ , _ , and _ , “Risk preference instability across institutions: A dilemma,” *Proceedings of the National Academy of Sciences of the United States of America*, 2005, *102* (11), 4209–4214.

Brandts, Jordi and Neus Figueras, “An exploration of reputation formation in experimental games,” *Journal of Economic Behavior & Organization*, 2003, *50* (1), 89–115.

Camerer, Colin and Keith Weigelt, “Experimental tests of a sequential equilibrium reputation model,” *Econometrica*, 1988, *56* (1), 1–36.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, “Sophisticated experience-weighted attraction learning and strategic teaching in repeated games,” *Journal of Economic Theory*, 2002, *104* (1), 137–188.

Chakravarthy, Jivas, Ed deHaan, and Shivaram Rajgopal, “Reputation repair after a serious restatement,” *Accounting Review*, 2014, *89* (4), 1329–1363.

Chatterji, Aaron K, David I Levine, and Michael W Toffel, “How well do social ratings actually measure corporate social responsibility?,” *Journal of Economics & Management Strategy*, 2009, *18* (1), 125–169.

Cole, Harold L and Narayana Kocherlakota, "Dynamic games with hidden actions and hidden states," *Journal of Economic Theory*, 2001, *98* (1), 114–126.

Coles, Jeffrey L, Michael L Lemmon, and J Felix Meschke, "Structural models and endogeneity in corporate finance: The link between managerial ownership and corporate performance," *Journal of Financial Economics*, 2012, *103* (1), 149–168.

Cremer, Jacques, "Cooperation in ongoing organizations," *Quarterly Journal of Economics*, 1986, *101* (1), 33–49.

Cripps, Martin, George Mailath, and Larry Samuelson, "Imperfect monitoring and impermanent reputations," *Econometrica*, 2004, *72* (2), 407–432.

Davis, Douglas D and Charles A Holt, "Experimental economics: Methods, problems, and promise," *Estudios Economicos*, 1993, pp. 179–212.

Diamond, Douglas W, "Reputation acquisition in debt markets," *Journal of Political Economy*, 1989, *97* (4), 828–862.

Economist Intelligence Unit, "Reputation: Risk of risks," *Economist Intelligence Unit Global Risk Briefing*, 2005, (4), 1–23.

Erev, Ido and Alvin E Roth, "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, 1998, *88* (4), 848–881.

Farber, David B., "Restoring trust after fraud: Does corporate governance matter?," *Accounting Review*, 2005, *80* (2), 539–561.

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

Forsythe, Robert, Russell Lundholm, and Thomas Rietz, "Cheap talk, fraud, and adverse selection in financial markets: Some experimental evidence," *Review of Financial Studies*, 1999, *12* (3), 481–518.

Fox, John, "The learning of strategies in a simple, two-person zero-sum game without saddlepoint," *Behavioral Science*, 1972, *17* (3), 300–308.

Frechette, Guillaume R., "Laboratory experiments: Professionals versus students," in Guillaume R. Frechette and Andrew Schotter, eds., *Handbook of Experimental Economic Methodology*, Oxford University Press, 02 2015.

Frydman, Cary and Gideon Nave, "Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism," *Management Science*, 2017, *63* (7), 2340–2352.

Fudenberg, Drew and David K Levine, "Maintaining a reputation when strategies are imperfectly observed," *The Review of Economic Studies*, 1992, *59* (3), 561–579.

Fuster, Andreas, Ricardo Perez-Truglia, Mirko Wiederholt, and Basit Zafar, "Expectations with endogenous information acquisition: An experimental investigation," *The Review of Economics and Statistics*, forthcoming 2020, pp. 1–54.

Gaines-Ross, Leslie, *Corporate Reputation: 12 Steps to Safeguarding and Recovering Reputation*, Hoboken, NJ: John Wiley & Sons, Inc., 2008.

Holmstrom, Bengt R. and Jean Tirole, "Chapter 2 The theory of the firm," in "Handbook of Industrial Organization," Vol. 1 of *Handbook of Industrial Organization*, Elsevier, 1989, pp. 61–133.

Huyck, John B Van, Raymond C Battalio, and Richard O Beil, "Tacit coordination games, strategic uncertainty, and coordination failure," *American Economic Review*, 1990, *80* (1), 234–248.

Jarrell, Gregg and Sam Peltzman, "The impact of product recalls on the wealth of sellers," *Journal of Political Economy*, 1985, *93* (3), 512–536.

John, Kose and David C Nachman, "Risky debt, investment incentives, and reputation in a sequential equilibrium," *Journal of Finance*, 1985, *40* (3), 863–878.

Johnson, Simon, Peter Boone, Alasdair Breach, and Eric Friedman, "Corporate governance in the Asian financial crisis," *Journal of Financial Economics*, 2000, *58* (1-2), 141–186.

Kagel, John H and Dan Levin, "Independent private value auctions: Bidder behaviour in first-, second-and third-price auctions with varying numbers of bidders," *Economic Journal*, 1993, *103*, 868–879.

Karpoff, Jonathan, "Does reputation work to discipline corporate misconduct?," in Michael Barnett and Timothy Pollock, eds., *The Oxford Handbook of Corporate Reputation*, Oxford University Press, 2011, chapter 18, pp. 361–382.

Karpoff, Jonathan M and John R Lott Jr, "The reputational penalty firms bear from committing criminal fraud," *Journal of Law and Economics*, 1993, *36* (2), 757–802.

_ , D Scott Lee, and Gerald S Martin, "The cost to firms of cooking the books," *Journal of Financial and Quantitative Analysis*, 2008, *43* (3), 581–611.

Kreps, David M, "Corporate culture and economic theory," in Peter Buckley and Jonathan Michie, eds., *Firms, Organizations, and Markets*, Oxford University Press, 1996, pp. 221–275.

_ and Robert Wilson, "Reputation and imperfect information," *Journal of Economic Theory*, 1982, *27* (2), 253–279.

_ and _ , "Sequential equilibria," *Econometrica*, 1982, *50* (4), 863–894.

Levin, Jonathan and Steven Tadelis, "Profit sharing and the role of professional partnerships," *The Quarterly Journal of Economics*, 2005, *120* (1), 131–171.

Lieberman, Bernhardt, "Experimental studies of conflict in some two-person and three-person games," in Joan H. Criswell, Herbert Solomon, and Patrick Suppes, eds., *Mathematical Methods in Small Group Processes: Papers*, Stanford, Calif.: Stanford University Press, 1962, chapter 14, pp. 203–220.

Liker, Jeffrey, *The Toyota-Way: 14 Management Principles from the World's Greatest Manufacturer*, McGraw-Hill Companies, Inc., 2004.

Liu, Qingmin, "Information acquisition and reputation dynamics," *The Review of Economic Studies*, 2011, *78* (4), 1400–1425.

Lucas, Robert E, "Adaptive behavior and economic theory," *Journal of Business*, 1986, *59* (4, Part 2), 401–426.

Mailath, George J and Larry Samuelson, "Who wants a good reputation?," *The Review of Economic Studies*, 2001, *68* (2), 415–441.

Maksimovic, Vojislav and Sheridan Titman, "Financial policy and reputation for product quality," *Review of Financial Studies*, 1991, *4* (1), 175–200.

McMillan, Charles, "From applause to notoriety: Organizational reputation and corporate governance," in Ronald J. Burke, Graeme Martin, and Cary L. Cooper, eds., *Corporate Reputation: Managing Opportunities and Threats*, Gower: Burlington, Vt. 2011, pp. 161–180.

Messick, David M, "Interdependent decision strategies in zero-sum games: A computer-controlled study," *Behavioral Science*, 1967, *12* (1), 33–48.

Milgrom, Paul and John Roberts, "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, 1982, *27* (2), 280–312.

Morrison, Alan D. and William J. Wilhelm, "Partnership firms, reputation and human capital," *American Economic Review*, 2004, *94*, 1682–1692.

Murphy, Deborah L, Ronald E Shrieves, and Samuel L Tibbs, "Understanding the penalties associated with corporate misconduct: An empirical examination of earnings and risk," *Journal of Financial and Quantitative Analysis*, 2009, *44* (1), 55–83.

Myerson, Roger B, "Optimal coordination mechanisms in generalized principal–agent problems," *Journal of Mathematical Economics*, 1982, *10* (1), 67–81.

Neral, John and Jack Ochs, "The sequential equilibrium theory of reputation building: A further test," *Econometrica*, 1992, pp. 1151–1169.

New York Times, "Kobe Steel blames plant managers for quality control scandal," *New York Times*, Nov. 10, 2017.

Noe, Thomas H., Michael J. Rebello, and Thomas A. Rietz, "Product market efficiency: The bright side of myopic, uninformed, and passive external finance," *Management Science*, 2012, *58* (11).

— , — , and — , "Firm reputation and agency: Information environments, corporate governance and its optics," 2025.

Peltzman, Sam, "The effects of FTC advertising regulation," *Journal of Law and Economics*, 1981, *24* (3), 403–448.

Plott, Charles R, "Industrial organization theory and experimental economics," *Journal of Economic Literature*, 1982, *XX*, 1485–1527.

— , "Will economics become an experimental science?," *Southern Economic Journal*, 1991, pp. 901–919.

Reputation Institute, "2025 Global RepTrak 100," 2025. Accessed: 2025-08-18.

Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood, "The value of reputation on eBay: A controlled experiment," *Experimental Economics*, 2006, *9* (2), 79–101.

Shachat, Jason and Todd Swarthout, "How do people play against Nash opponents in games which have a mixed strategy equilibrium?," *Working Paper*, 2008.

Shleifer, Andrei and Daniel Wolfenzon, "Investor protection and equity markets," *Journal of Financial Economics*, 2002, *66* (1), 3–27.

— and Robert W Vishny, "A survey of corporate governance," *The Journal of Finance*, 1997, *52* (2), 737–783.

Smith, Vernon L and James M Walker, "Monetary rewards and decision cost in experimental economics," *Economic Inquiry*, 1993, *31* (2), 245–261.

The Economist, "Schumpeter: Getting a handle on a scandal," *The Economist*, 2018, *March 28.*

Tirole, Jean, "A theory of collective reputations (with applications to the persistence of corruption and to firm quality)," *The Review of Economic Studies*, 1996, *63*, 1–22.

Zizzo, Daniel John, "Experimenter demand effects in economic experiments," *Experimental Economics*, 2010, *13* (1), 75–98.

For online publication only

Appendix
Governance, Stakeholder Welfare, Crises and Recovery: An Experiment

A Proofs of results

Proof of Lemma 1. Since customers cannot observe the period 2 good's quality and the good's price reflects customers' expectation about its quality, conditioned on quality $q \in \{h, l\}$, the payoff to a normal operator who selects operating policy $o_2 \in \{m, s\}$ in period 2 at any $h_2^o \in \mathcal{H}_2^o$ is

$$\begin{aligned} \phi_q(p_2) \quad & \text{if } o_2 = s, \\ \phi_q(p_2) + c \quad & \text{if } o_2 = m. \end{aligned} \tag{A.1}$$

The result follows directly from inspecting the operator payoffs described in this equation. □

Proof of Lemma 2. Suppose that σ^* is the normal operator's mixed strategy in period 1 in a sequential equilibrium. The period 1 good's price is set before customers observe anything other than the payment structure. The committed operator always chooses policy s, and under this policy s, $q_1 = h$. So the probability that customers assign to the good being high quality depends only on their prior belief that the operator is the committed type, ρ, and their hypothesis about the operator's strategy if the operator is the normal type, σ. Thus, the expected value of the period 1 good equals $P_1(\sigma^*) = \rho + (1-\rho)(\sigma^* + (1-\sigma^*)\delta)$, where P_1 is defined by equation (4). Hence, the pricing condition for a sequential equilibrium (Definition 1.a) implies that $p_1^* = P_1(\sigma^*)$.

Now consider the claim about the period 2 customer belief and price when $q_1 = h$. It is possible for $q_1 = h$ when the (i) operator is committed and when the (ii) operator is normal. The probability of (i) is ρ. The probability that $q_1 = h$ if the operator is normal, when the operator chooses policy s with probability σ^* is $\sigma^* + (1-\sigma^*)\delta$. Thus, Bayes rule implies that customer's posterior belief,

$\mu_2^*(h_1^c = (\Phi, h))$, is given by

$$\frac{\rho}{\rho + (1-\rho)(\sigma^* + (1-\sigma^*)\delta)} = \rho / P_1(\sigma^*).$$

which equals the upper branch of expression (5), the case where $q_1 = h$.

Lemma 1 shows that $o_2^* = m$. Using the same arguments as we have just employed to compute expected quality in the context of the period 1 good's price, we see that the pricing condition (Definition 1.a) implies that, if $q_1 = h$,

$$p_2^* = 1 \times \frac{\rho}{P_1(\sigma^*)} + \delta \times \left(1 - \frac{\rho}{P_1(\sigma^*)}\right) = P_2(\rho / P_1(\sigma^*)), \tag{A.2}$$

where P_2 is defined in (6).

Finally, consider the claim about the period 2 customer belief and price when $q_1 = l$. Note that $q_1 = l$ is only possible when the operator is normal. Thus, Bayes rule implies that customer's posterior belief that the operator will be committed in period 2 is given by the lower branch of expression (5), i.e., if $q_1 = l$

$$\mu_2 = r.$$

Using the same arguments we used to compute expected quality in the context of the period 1 good's price, we see that the pricing condition (Definition 1.a) implies that, if $q_1 = l$,

$$p_2^* = 1 \times r + \delta \times (1 - r) = P_2(r). \tag{A.3}$$

□

Proof of Lemma 3. For fixed payment structure, ϕ, histories in $\mathcal{H}_1^o$ vary only with respect to period 1 prices, $p_1 \in \mathbb{R}_+$. Because prices reflect customers' expectations and customers cannot observe the period 1 operating policy, the difference between the payoff to a normal operator from the s and m policies is not affected by p_1. Hence, any policy that is a best response for the operator at a history h_1^o, is a best response for all $h_1^o \in \mathcal{H}_1^o$.

Proof of Part a. From Lemmas 1 and 2, it follows that if customers expect the operator to pick the safe policy in period 1 with probability σ^* and the operator remains with the firm after reform, equation (7) describes the operator's expected payoff. In equilibrium, σ^* must be a best response for the operator. Thus, because $\rho/P_1(\sigma^*) = \rho$ when $\sigma^* = 1$, $\sigma^* = 1$ in equilibrium when

$$\begin{aligned}
&\phi_1(P_1(1)) + \phi_h(P_2(\rho)) + c \\
&\geq \phi_1(P_1(1)) + c + \delta\Big(\phi_h(P_2(\rho)) + c\Big) + (1-\delta)\Big(r\phi_l + (1-r)(\phi_l + c)\Big) \\
\iff &(1-\delta)\Big(\phi_h(P_2(\rho)) - \phi_l\Big) \geq \delta c + (1-\delta)(1-r)c.
\end{aligned} \tag{A.4}$$

Similarly, $\sigma^* = 0$ in equilibrium when

$$\begin{aligned}
&\phi_1(P_1(0)) + \phi_h(P_2(\rho/P_1(0))) + c \\
&\leq \phi_1(P_1(0)) + c + \delta\Big(\phi_h(P_2(\rho/P_1(0))) + c\Big) + (1-\delta)\Big(r\phi_l + (1-r)(\phi_l + c)\Big) \\
\iff &(1-\delta)\Big(\phi_h(P_2(\rho/P_1(0))) - \phi_l\Big) \leq \delta c + (1-\delta)(1-r)c.
\end{aligned} \tag{A.5}$$

Finally, $\sigma^* \in (0,1)$ in equilibrium when the operator is indifferent between policies s and m in period 1, or equivalently

$$\begin{aligned}
&\phi_1(P_1(\sigma^*)) + \phi_h(P_2(\rho/P_1(\sigma^*))) + c \\
&= \phi_1(P_1(\sigma^*)) + c + \delta\Big(\phi_h(P_2(\rho/P_1(\sigma^*))) + c\Big) + (1-\delta)\Big(r\phi_l + (1-r)(\phi_l + c)\Big) \\
\iff &(1-\delta)\Big(\phi_h(P_2(\rho/P_1(\sigma^*))) - \phi_l\Big) = \delta c + (1-\delta)(1-r)c.
\end{aligned} \tag{A.6}$$

Proof of Part b. From Lemmas 1 and 2, it follows that if the operator picks the safe policy in period 1 with probability σ^* and is replaced after reform, we can describe the operator's expected payoff by rewriting equation (7) as

$$v_1^o(o_1) = \phi_1(P_1(\sigma^*)) + \begin{cases} \phi_h(P_2(\rho/P_1(\sigma^*))) + c & o_1 = s. \\ c + \delta\left(\phi_h(P_2(\rho/P_1(\sigma^*))) + c\right) & o_1 = m \end{cases}. \tag{A.7}$$

The rest of the proof is identical to that of Part a once we account for the terms that are dropped from equation (7) to arrive at equation (A.7).

Proof of Part c. From expressions (4) and (5) it is clear that $\rho/P_1(\sigma)$ is strictly decreasing in σ. Moreover, from (6) it follows that the period 2 price is increasing in $\mu_2 = \rho/P_1(\sigma)$. Thus, if ϕ_h is strictly increasing in P_2, the difference $\phi_h P_2(\rho/P_1(\sigma^*)) - \phi_l$ is decreasing in σ. It follows that, conditions for equilibria with $\sigma^* \leq 1$ cannot be satisfied when the condition for the reputation equilibrium with $\sigma^* = 1$ is satisfied. □

Proof of Proposition 1. The proof follows directly from Lemma 3. To see this note that, Part a of Lemma 3 shows that, if the operator remains with the firm after reform, a reputation equilibrium exists if and only if

$$(1-\delta)\big(\phi_h(P_2(\rho))-\phi_l\big) \geq \delta c+(1-\delta)(1-r)c.$$

Note that $\pi_1 = 1-e$ and under the owner management payment structure, $\phi_h(P_2(\rho)) = P_2(\rho)-e$ and $\phi_l = P_2(r)-e-R = \pi_r$, so a reputation equilibrium exists if and only if

$$\begin{aligned}
&(1-\delta)(P_2(\rho)-e-\pi_r) \geq \delta c+(1-\delta)(1-r)c \\
\Longleftrightarrow &(1-\delta)(P_2(\rho)-(1-\pi_1)-\pi_r) \geq \delta c+(1-\delta)(1-r)c \\
\Longleftrightarrow &(1-\delta)P_2(\rho) \geq (1-\delta)(1-\pi_1+\pi_r)+\delta c+(1-\delta)(1-r)c
\end{aligned} \tag{A.8}$$

satisfies the hypotheses of Part a of Lemma 3. Since $\phi_h(P_2(\rho))$ is strictly increasing in P_2, Part c of Lemma 3 shows that the reputation equilibrium is unique. □

Proof of Proposition 2. From equations (5) and (6) it follows that, when $q_1 = h$

$$\begin{aligned}
P_2(\mu_2(0)) =&\frac{\rho}{P_1(0)}+\left(1-\frac{\rho}{P_1(0)}\right)\delta = \delta+\frac{\rho(1-\delta)}{P_1(0)} \\
=&\delta+\frac{\rho(1-\delta)}{P_1(0)}+\frac{\delta}{P_1(0)}-\frac{\delta}{P_1(0)} \\
=&1+\delta-\frac{\delta}{P_1(0)}
\end{aligned} \tag{A.9}$$

where the last expression follows since $P_1(0) = \rho+(1-\rho)\delta = \delta+\rho(1-\delta)$.

Noting that $\pi_1 = 1-e$, $\pi_r = P_2(r)-e-R$, $P_1(0) = \rho+(1-\rho)\delta = P_2(\rho)$, and comparing expected payments described by equations (9) and (12) shows that the owner will implement $q_1 = s$ when

$$\begin{aligned}
&1-e+P_2(\rho)-e-b^* \geq P_1(0)-e+P_1(0)\Big(P_2(\mu_2(0))-e\Big)+\Big(1-P_1(0)\Big)\pi_r \\
\Longleftrightarrow\ &\pi_1+P_2(\rho)-e-b^* \geq P_2(\rho)-e+P_2(\rho)\left(1+\delta-\frac{\delta}{P_2(\rho)}-e\right)+(1-P_2(\rho))\pi_r \\
\Longleftrightarrow\ &\pi_1-b^* \geq P_2(\rho)\Big(\pi_1+\delta\Big)+\delta+(1-P_2(\rho))\pi_r \\
\Longleftrightarrow\ &b_1^* \leq (1-P_2(\rho))(\pi_1-\pi_r+\delta).
\end{aligned} \tag{A.10}$$

Rearranging the last inequality in expression (A.10) produces condition (13) in the proposition. When condition (13) is satisfied, if the implementing contract, $\phi_h(p_2)$, is strictly increasing in p_2 then by Part c of Lemma 3, $q_1 = s$ is the manager's unique best response. This establishes the final claim in the proposition about all equilibria producing $q_1 = s$ when the owner chooses reputation-assuring compensation. □

Proof of Proposition 3. The proof of this result follows the same approach as used in the proofs of Propositions 2 and 1. First apply the payment structures for owner and professional management, and then use Lemma 3 (Part b) to characterize the conditions for a reputation equilibrium.

Proof of Part a. Given the payment structure under owner management, Part b of Lemma 3 shows that picking $o_1 = s$ in period 1 is a best response for the operator if and only if

$$\begin{aligned} & (1-\delta)\,\phi_h(P_2(\rho)) \geq \delta c \\ \iff & (1-\delta)\big(P_2(\rho)-e\big) \geq \delta c \\ \iff & (1-\delta)P_2(\rho) \geq (1-\delta)(1-\pi_1)+\delta c \end{aligned} \tag{A.11}$$

since $\pi_1 = 1-e$.

Proof of Part b Note that $\pi_1 = 1-e$, $P_1(0) = \rho+(1-\rho)\delta = P_2(\rho)$ and $P_2(\mu_2(0)) = 1+\delta-\delta/P_1(0)$. Given the payment structure under professional management and the optimal bonus payment for the manager (equation (11)), the owner will choose reputation-assuring compensation when

$$\begin{aligned} 1-e+P_2(\rho)-e-b^* &\geq P_2(\rho)-e+P_1(0)(P_2(\mu(0))-e) \\ \iff \pi_1-b^* &\geq P_2(\rho)\left(1+\delta-\frac{\delta}{P_2(\rho)}-e\right) \\ \iff b_1^* &\leq (1-P_2(\rho))(\pi_1+\delta). \end{aligned} \tag{A.12}$$

Proof of Part c The proof of the result under owner management follows directly from comparing condition (A.8) for reputation equilibria when reform is possible with condition (A.11) in the Proof of Part a and noting that the left hand sides of both expressions are identical while the right hand side of condition (A.11) is always smaller.

The proof of the result under professional management follows directly from comparing condition (A.10) which ensures that $o_1 = s$ when reform is possible with condition (A.12) in the Proof of Part b and noting that the left hand sides of both expressions are identical while the right hand side of condition (A.12) is always larger. □

Lemma A-1. *When the firm is controlled by an owner-operator, in a mixed equilibrium:*

1. $p_1^* = \frac{P_2(\rho)-\delta}{p_2^*-\delta}$.
2. *If the owner-operator remains unrevealed,* $p_2^* = 1-\pi_1+b^*+\pi_r+g$.

3. *If the owner-operator is normal, the owner picks policy s with probability* $\sigma^* = 1 - \frac{1-p_1^*}{(1-\rho)(1-\delta)}$.

Proof of Lemma A-1. Beliefs and, thus, prices conform to Bayes' rule. Thus, the period 2 price conditioned on the owner-operator remaining unrevealed in period 1 equals:

$$P_2(\mu(\sigma^*)) = \frac{\rho}{P_1(\sigma^*)} + \left(1 - \frac{\rho}{P_1(\sigma^*)}\right)\delta = \delta + \frac{\rho(1-\delta)}{P_1(\sigma^*)} = \delta + \frac{P_2(\rho)-\delta}{P_1(\sigma^*)}$$
$$\iff P_1(\sigma^*) = \frac{P_2(\rho)-\delta}{P_2(\mu(\sigma^*))-\delta}. \quad \text{(A.13)}$$

To be willing to randomize, the owner-operator must be indifferent between policies s and m. According to Lemma 3, the condition for $\sigma^* \in (0,1)$ is $(1-\delta)\left(\phi_h\left(P_2\left(\frac{\rho}{P_1(\sigma^*)}\right)\right) - \phi_l\right) = c + (1-\delta)(1-r)c$. Once we account for the owner management payment structure, this condition can be rewritten as

$$(1-\delta)\left(P_2\left(\mu(\sigma^*)\right) - e - \pi_r\right) = c + (1-\delta)(1-r)c$$
$$\iff (1-\delta)(P_2(\mu(\sigma^*)) - (1-\pi_1) - \pi_r) = \delta c + (1-\delta)(1-r)c$$
$$\iff (1-\delta)P_2(\mu(\sigma^*)) = (1-\delta)(1-\pi_1+\pi_r) + \delta c + (1-\delta)(1-r)c$$
$$\iff P_2(\mu(\sigma^*)) = (1-\pi_1+\pi_r) + b^* + (1-r)c, \quad \text{(A.14)}$$

since $b^* = \delta c/(1-\delta)$.

The probability σ^* follows from noting that

$$p_1^* = \rho + (1-\rho)(\sigma^* + (1-\sigma^*)\delta)$$
$$= \rho + (1-\rho)(\delta + \sigma^*(1-\delta))$$
$$= 1 - 1 + \rho + (1-\rho)(\delta + \sigma^*(1-\delta))$$
$$= 1 - (1-\delta)(1-\rho) + (1-\delta)(1-\rho)\sigma^*$$
$$\iff (1-\delta)(1-\rho) - (1-p_1^*) = (1-\delta)(1-\rho)\sigma^*. \quad \text{(A.15)}$$

□

B Subject instructions

This appendix contains instructions for all three treatments under Parameter Set I. Differences are set off by "((" followed by **OM:** for Owner Management, **OMR:** for Owner Management with Reform or **PMR:** for Professional management; and closed by "))." Parameter Set II instructions differ in describing the fraction of managers who have Method 2 (the vulnerable production technology) available and changes manager payoffs in accordance with the parameter set.

INSTRUCTIONS

General

You are about to participate in an experiment on decision making. If you follow these instructions carefully and make good decisions, you might earn a considerable amount of money, which will be paid to you in cash at the end of the experiment.

At the beginning of the experiment, all players are assigned a type: Blue or Green. You will remain Blue or Green throughout the experiment. Next, one Blue and one Green Player are randomly assigned to a 2-player group. These players will interact with each other for two periods. Then, players are randomly re-assigned to new groups for two more periods. This process repeats until the end of the experiment.

The currency used in these games is francs. Each period, you receive franc payoffs that are yours to keep. At the end of the experiment, you will exchange francs for dollars. Your individual exchange rate will be displayed on your computer terminal after the experiment starts. Do not reveal this number to anyone. At the end of the experiment, your francs will be converted to dollars at this rate, and you will be paid in dollars. Note that the more francs you earn, the more dollars you earn regardless of the exchange rate.

Stages of the Game

Each period is divided into two stages.

Briefly, in Stage I, the Blue Player makes decisions about producing an item and sells it to the experimenter. In Stage II, the Green Player will have an opportunity to buy the item from the experimenter. There are two types of items: Round, worth 1,000 francs to the Green Player and Square, worth 0 francs to the Green Player.

Because it will make it easier to understand, we describe the Stages in reverse order.

Stage II Instructions

Green Player Decision

In Stage II, the Green Player sets a "Limit Price" for an item and *may* buy the item. If he or she does buy the item, it will be at a "Discounted Price" that is less than or equal to the Limit Price.

If the Green Player buys the item, he or she will receive a "Redemption Value" of 1,000 francs for a Round item and 0 francs for a Square item.

The item type is determined in Stage I but the type will not be known by the Green Player until after he or she establishes the Limit Price in Stage II.

Procedures

The Limit Price and the Discounted Price for the item will be determined as follows.

1. Limit Price: The Green Player will be asked to indicate the *highest* price he or she is willing to pay for the item. This determines the ***Limit Price***. The Limit Price must be between 0 and 1,000 (inclusive). The Green Player is asked to enter the price and press "OK" to continue.

2. Discounted Price: The computer program will determine a ***Discounted Price*** by drawing a random number between 0 and 1,000 (inclusive) to be the Discounted Price.

If the random draw is ***less than or equal to*** the Limit Price, the Green Player will buy the item at the Discounted Price from the experimenter and receive the Redemption Value.

If the random draw is ***greater than*** the Limit Price indicated by the Green Player, then the Green Player will not buy the item.

Thus, the Limit Price defines the highest price that the Green Player will pay for the item. If he or she buys the item it will be at a Discounted Price less than or equal to the Limit Price.

Payoff Determination

The Green Player starts each period with working capital of 1,000 francs. This allows the Green Player to pay up to 1,000 francs for an Item. Three quarters (75%) of the working capital will be deducted at the end of the period. The Green Player's payoffs are determined as follows:

If the Discounted Price is greater than the Limit Price:

Period Payoff = 1,000 – 750 = 250.

If the Discounted Price is less than the Limit Price:

Period Payoff = 1,000 + Redemption Value – Discounted Price – 750.

Notice that buying the item increases your payoff whenever the Redemption Value is higher than the Discounted Price AND you buy the item (that is, the Limit Price is higher than the Discounted Price).

Notes on this Procedure

If you are a Green Player, it is in your best interest to be accurate; that is, the best thing you can do is state truthfully the highest price you would be willing to pay for the item as the Limit Price. If the price stated is too high or too low, then you pass up opportunities that you would prefer.

Suppose you are a Green Player and you are certain that the Item will be Round. Then, you should set a Limit Price of 1,000. If you set it less, say 800, then whenever the random draw is between 800 and 1,000, you will not purchase the item even though the item is worth more to you than the discounted price. For example, if the Discounted Price is 900, you are better off buying the item (Payoff = 1000+1000-900-750 = 350) than not buying the item (Payoff = 1000-750 = 250). The only way to insure that you will always buy the item when the Discounted Price is less than the value is to set the Limit Price at 1,000.

Similarly, if you are certain that the Item will be Square, then you should set a Limit Price of 0. If you set it more, say 200, then whenever the random draw is between 0 and 200, you end up buying the item even though it is worthless to you. For example, if the Discounted Price is 100, you are worse off buying the item (Payoff = 1000+0-100-750 = 150) than not buying the item (Payoff = 1000-750 = 250). The only way to insure that you will never buy the item when the Discounted Price is more than the value is to set the Limit Price at 0.

Similar arguments imply that, if, on average, you expect the item value to be X francs (e.g., 500), then you should set a Limit Price equal to X (e.g., 500). If so, you always buy the item when the Discounted Price is less than what you think it is worth on average and never buy it if the Discounted Price is more than what you think it is worth.

We will discuss Stage I next. Before doing that, are there any questions about the Green Player's actions in Stage II and the Green Player's payoffs?

Stage I Instructions

In Stage I, the Blue Player selects a "production method." Methods available and the earnings received by the Blue Player vary as described below.

Determining Available Production Methods

In the first period of a group interaction, the computer randomly determines the production methods available to the Blue Player.

In the first period, 1/4 (25%) of the Blue Players MUST produce according to the following method in BOTH periods:

- **Method 1: Sure**, which always produces a Round Item

In the first period, the other 3/4 (75%) of the Blue Players can CHOOSE between producing according to:

- **Method 1: Sure**, which always produces a Round Item.
- **Method 2: Mixed**, which produces a Round Item 1/20 (5%) of the time and a Square Item 19/20 (95%) of the time.

In the second period, these Blue Payers will be restricted to a single production method as follows:

- **((OM:** They MUST produce according to **Method 2 Mixed** in the second period of the group interaction.**))**
- **((OMR & PMR:** If these Blue Players produce a Round Item in the first period of a group interaction, they MUST produce according to **Method 2 Mixed** in the second period of the group interaction.**))**
- **((OMR & PMR:** If these Blue Players produce a Square Item in the first period of a group interaction, the computer will randomly re-draw production methods for the second period of the group interaction. In this case, 1/10 (10%) of the Blue Players MUST produce according to **Method 1: Sure** and 9/10 (90%) MUST produce according to **Method 2: Mixed** in the second period of the group interaction.**))**

Payoff Determination

The Blue Player's payoffs are determined by a "Sale Price" set by the experimenter and a method payment in each period.

In the first period of a group interaction, the Sale Price equals **((OM & OMR:** the Limit Price set by the Green Player**))** **((PMR:** 111**))**. Thus, the payoff is ((OM & OMR: the Limit Price)) ((PMR: 111)) minus a method payment that depends on the method chosen. Specifically:

Period 1 Payoff		
Method:	Method 1: Sure	Method 2: Mixed
Availability:	Always	7/8 (87.5%) of the time
+ Sale Price - Method Payment	((**OM & OMR:** Limit Price)) ((**PMR:** 111)) -111	((**OM & OMR:** Limit Price)) ((**PMR:** 111)) -51
Total Payoff	((**OM & OMR:** Limit Price - 111)) ((**PMR:** 0))	((**OM & OMR:** Limit Price - 51)) ((**PMR:** 60))

For example, suppose the Limit Price was 500 in Period 1. Your payoff would be:

- ((**OM & OMR:** 500-111=389)) ((**PMR:** 111-111=0)) if **Method 1: Sure** was available and chosen.
- ((**OM & OMR:** 500-51=449)) ((**PMR:** 111-51=60)) if **Method 2: Mixed** was available and chosen.

In the second period of the group interaction, the Blue Player is always restricted to one production method. The Sale Price is set by the experimenter depending on the Item Type produced in Period 1 ((**OMR & PMR:** and the redraw outcome. In addition, if the Period 1 Item was Square, the experimenter will charge a cost to redraw production methods)).

The Blue Player's payoffs are summarized in the following table:

((**OM:**

Period 2 Payoff			
Period 1 Item:	Round		Square
Method:	Method 1: Sure	Method 2: Mixed	Method 2: Mixed
Availability:	If only Method 1 was available in Period 1	If both Methods were available in Period 1	Always
+ Sale Price - Method Payment	288 -111	288 -51	51 -51
Total Payoff	177	237	0

))

((**OMR:**

Period 2 Payoff				
Period 1 Item:	Round (No Redraw)		Square (Redraw)	
Method:	Method 1: Sure	Method 2: Mixed	Method 1: Sure	Method 2: Mixed
Availability:	If only Method 1 was available in Period 1	If both Methods were available in Period 1	1/10 (10%) of the time	9/10 (90%) of the time
+ Sale Price	288	288	145	145
- Method Payment	-111	-51	-111	-51
- Redraw Cost	0	0	-10	10
Total Payoff	177	237	24	84

))

((PMR:

Period 2 Payoff				
Period 1 Item:	Round (No Redraw)		Square (Redraw)	
Method:	Method 1: Sure	Method 2: Mixed	Method 1: Sure	Method 2: Mixed
Availability:	If only Method 1 was available in Period 1	If both Methods were available in Period 1	1/10 (10%) of the time	9/10 (90%) of the time
+ Sale Price	141	141	121	121
- Method Payment	-111	-51	-111	-51
- Redraw Cost	0	0	-10	70
Total Payoff	30	90	0	0

))

In Period 2 after a Round Item in Period 1, your payoff would be:

- **((OM:** 177)) **((OMR:** 177)) **((PMR:** 30)) if **Method 1: Sure** becomes the only available method in Period 2 because it was the only available method in Period 1.
- **((OM:** 237)) **((OMR:** 237)) **((PMR:** 90)) if **Method 2: Mixed** becomes the only available method in Period 2 because both methods were available in Period 1 and a Round item was produced.

In Period 2 after a Square Item in Period 1 **((OMR & BMR:** and a redraw in Period 2**))**, your payoff would be:

- **((OM:** 0 with **Method 2: Mixed** becoming the only available method.**))**
- **((OMR:** 24 if **Method 1: Sure** becomes the only available method as a result of the redraw.**))**

- **((PMR:** 0 if **Method 1: Sure** becomes the only available method as a result of the redraw.**))**
- **((OMR:** 84 if **Method 2: Mixed** becomes the only available method as a result of the redraw.**))**
- **((PMR:** 0 if **Method 2: Mixed** becomes the only available method as a result of the redraw.**))**

We will discuss how player types and groups are determined next. Before doing that, are there any questions about the Blue Player's actions in Stage I and the Blue Player's payoffs?

<u>Group, Player Type and Technology Determination</u>

At the beginning of the experiment, you are assigned a player type: "Blue" or "Green." You remain this type of player for the entire experiment. Every 2 periods, one Blue and one Green Player are matched randomly to play the game for two periods. Everyone is re-assigned to new groups every 2 periods. Thus, in periods 1 and 2, you will be with one group. In period 3, you will be randomly re-assigned to new groups for periods 3 and 4, etc.

The computer draws new Production Methods at the beginning of each group interaction (that is, every 2 periods) as follows:

1. 1/4 (25%) of the Blue Players will ONLY be able to produce using "**Method 1: Sure**."
2. 3/4 (75%) of the Blue Players on average will be able to produce using EITHER "**Method 1: Sure**" OR "**Method 2: Mixed**."

((OM: In**))** **((OMR & PMR:** If no redraw occurs in**))** Period 2 of a group interaction, "**Method 1: Sure**" remains the only available production method if it was the only available method in Period 1. "**Method 2: Mixed**" becomes the only available production method if both methods were available in Period 1.

((OMR & PMR: If a redraw occurs in Period 2 of a group interaction because a Square item was produced in Period 1, "**Method 1: Sure**" becomes the only available production method 1/10 (10%) of the time and "**Method 2: Mixed**" becomes the only available production method 9/10 (90%) of the time.**))**

Neither the Methods available nor the Method actually chosen are revealed to the Green Player. The Item type is revealed after Period 1 **((OMR & PMR:** and whether methods were redrawn is revealed between periods**))**. So, in Period 2, there are the following possibilities:

1. The Blue Player produced a Round Item in Period 1. Thus, the Green Player does not know whether "**Method 1: Sure**" or "**Method 2: Mixed**" was used in Period 1 nor which is the only available method in Period 2. (Recall, **Method 1: Sure** is always available and, sometimes, **Method 2: Mixed** produces a Round Item in Period 1.)

2. ((**OM:** The Blue Player produced a Square Item in Period 1. Thus, the Green Player knows that "**Method 2: Mixed**" was available and used in Period 1 AND **Method 2: Mixed** will be the only method available in Period 2.)))

3. ((**OMR & PMR:** The Blue Player produced a Square Item in Period 1 and a redraw occurred. Thus, the Green Player knows that "**Method 2: Mixed**" was available and used in Period 1, but does not know the method available in Period 2. (Recall that, after a redraw, 1 in 10 times, only **Method 1: Sure** is available and 9 in 10 times only **Method 2: Mixed** is available.)))

<u>End of Period Results</u>

The computer program reports results to you at the end of each period. Beginning in the next period a history of results appears in two windows. The left window shows the publicly available information given to all players in your group. Note that this will NOT give the Methods available to or used by the Blue Player. The right window gives your private information and shows your payoff. You can scroll through these windows and record results as you wish. In particular, you may want to fill in your earnings on your receipt each period.

<u>End of Experiment Rules</u>

The computer will total your earnings in francs at the end of the experiment. It multiplies this amount by the exchange rate to determine the amount of dollars you receive. Fill this number in on your receipt. You will be paid this amount of dollars in cash.

Are there any questions?

Summary Sheet

Methods are drawn randomly each time groups are reassigned as follows:

1. Initially, 1/4 of the Blue Players will have ONLY "**Method 1: Sure**" available.

2. The other 3/4 of the Blue Players will have both "**Method 1: Sure**" and "**Method 2: Mixed**" available.

Within a group:

1. If the Period 1 item is Round and "**Method 1: Sure**" was the only method available, it remains the only available method in Period 2.

2. If the Period 1 item was Round and both methods were available, "**Method 2: Mixed**" becomes the only available method in Period 2.

3. **((OM:** If the item produced in Period 1 was Square then, "**Method 2: Mixed**" becomes the only available method in Period 2.**))**

4. **((OMR & PMR:** If the item produced in Period 1 was Square then, methods are redrawn in Period 2.

   4.1. "**Method 1: Sure**" will be the only method 1/10 times.

   4.2. "**Method 2 Mixed**" will be the only method the other 9/10 times.**))**

Neither the methods available nor the method actually used in any given period will be revealed to the Green Player. Only the Item type will be revealed.

The game proceeds as follows:

1. Players are randomly assigned to groups every two periods. The Blue Player is assigned newly drawn Methods.

2. Stage I

   2.1. The Blue Player chooses a method.

   2.2. The Blue Player receives a payment as follows:

      - In the first period of a group interaction, the Blue Player receives: **((OM & OMR:** the Limit Price of the item minus 111 if **Method 1: Sure** is used OR the Limit Price minus 51 if **Method 2: Mixed** is used.
         - **((PMR:** 0 if **Method 1: Sure** is used.**))**
         - **((PMR:** 60 if **Method 2: Mixed** is used.**))**
      - In the second period of a group interaction, the Blue Player receives:
         - **((OM:** 177**))** **((OMR:** 177**))** **((PMR:** 30**))** if a Round item was produced in Period 1 and Method 1: Sure is used in Period 2.

– ((**OM:** 237)) ((**OMR:** 237)) ((**PMR:** 90)) if a Round item was produced in Period 1 and Method 2: Mixed is used in Period 2.
– ((**OM:** 0 if a Square item was produced in Period 1.))
– ((**OMR:** 24 if a Square item was produced in Period 1 and Method 1: Sure is used in Period 2.))
– ((**OMR:** 84 if a Square item was produced in Period 1 and Method 2: Mixed is used in Period 2.))
– ((**PMR:** 0 if a Square item was produced in Period 1 and Method 1: Sure is used in Period 2.))
– ((**PMR:** 0 if a Square item was produced in Period 1 and Method 2: Mixed is used in Period 2.))

3. Stage II

3.1. The Green Player receives working capital of 1,000 each period.

3.2. The Green Player will be told whether Methods have been newly drawn, re-drawn or be restricted to a single method from the prior period.

3.3. Prices are determined:

- The Green Player decides the most he or she is willing to pay for the item and sets the Limit Price.
- The computer draws a random Discounted Price between 0 and 1,000.
- If the Discounted Price is less than or equal to the Limit Price, the Green Player will buy the item at the Discounted Price and receive the Redemption Value (1000 for Round and 0 for Square).

3.4. In all cases, the Green Player returns 75% of the working capital (750).

You are free to make as much money as you can according to these rules.

C Learning and experience weighted attraction in the experiment

Prior research on reputation in laboratory settings documents that (1) subject behavior deviates systematically from predictions of a rational choice equilibrium benchmark and (2) as subjects gain experience, their strategies better approximate equilibrium predictions (e.g., Brandts and Figueras, 2003; Noe et al., 2012). The main purpose of our experiment is to see whether our theory survives such systematic behavioral deviations. We document this in the main body of the paper. Here, we ask whether the experience moves behavior toward equilibrium predictions as documented in prior research.

C.1 Operator choices as a Markov process

We study how operator choices evolve using two methods. First, we model them as a Markov process (cf. Axelrod, 1987; Erev and Roth, 1998) and ask whether state transitions evolve toward equilibrium predictions. There are three possible operator condition–choice pairs: (secure, reliable), (insecure, reliable), (insecure, vulnerable). We use these three pairs as the states of the Markov model. We term the first pair the "secure" state. The second the "reputation" state, and the third, the "opportunistic" state.[1] We estimate the transition matrix using the method of moments, and use the estimated transition matrix to compute steady-state probabilities (Norris, 1998). The estimates are presented in Panel A of Table C.1. Panel B of Table C.1 provides estimates conditioned on the previous round's outcome. Overall, the evidence in Table C.1 indicates a tendency for reputable behavior to increase with operator experience.

C.2 Experience weighted attraction

We also estimate regression models of subject behavior. Again, we are asking whether behavior moves toward equilibrium behavior with experience. We model subject choices as a function of the treatment, subject demographics, and experience. We use dummy variables to represent treatments and parameter set/treatment interactions and gender (which is the only demographic variable that appears significantly correlated with behavior). We account for subject experience using experience weighted attraction (EWA) models along the lines of Erev and Roth (1998) and Camerer et al. (2002) that account for the profitability of past choices.

[1] Of course, the operator has no control over transitions to the strategy in the secure state. These transitions are controlled entirely by the parameter set and random draws. However, the transitions will figure into the overall stable probabilities and, therefore, must be considered.

Table C.1: *Evolution of reputation formation.* Panel A presents the initial frequency, overall frequency, and estimated steady state probability of the reputation state, the state where the firm is insecure yet the reliable technology is adopted by the operator. Panel B presents the frequency of adoption of the reliable technology by operators whose firms are insecure conditioned on the outcome in the previous round. The last column of Panel B shows the χ^2 statistic for independence between the prior strategies and reputable behavior in the next period.

Panel A: Frequencies and steady state (SS) probabilities of reputation building

	Paremter Set 1			Paremter Set 2		
Treatment	Round 1	Overall	SS	Round 1	Overall	SS
OM	61.1%	68.2%	74.0%	61.9%	72.6%	68.7%
obs.	18	198	#N/A	21	241	#N/A
OMR	50.0%	56.8%	50.2%	50.0%	50.2%	57.2%
obs.	18	199	#N/A	20	225	#N/A
PMR	47.1%	59.6%	72.8%	46.7%	69.7%	61.3%
obs.	17	208	#N/A	15	201	#N/A

Panel B: Frequency of reputation building conditioned on previous round's outcome

		Outcome in previous round				
	State:	Secure	Reputation	Opportunistic		
	Good's Quality:	High	High	Low	High	
	Treatment					Chi2 (p-value)
Parameter Set 1	OM	58.7%	75.3%	66.7%	0.0%	6.36 (0.095)
	OMR	48.7%	62.2%	58.6%	0.0%	4.70 (0.195)
	PMR	52.3%	85.2%	31.5%	20.0%	46.31 (0.000)
Parameter Set 2	OM	80.0%	78.5%	61.1%	0.0%	12.10 (0.007)
	OMR	30.8%	61.5%	46.4%	0.0%	13.12 (0.004)
	PMR	59.1%	83.9%	51.0%	33.3%	22.38 (0.000)

C.2.1 Experience weighted attraction for operators

To model operator behavior, along the lines of Erev and Roth (1998) and Camerer et al. (2002), we define the attraction to the reputation, opportunism and secure strategies as follows:

$$
\begin{aligned}
A_{s,t}^{Reputation} &= A_{s,t-1}^{Reputation} + \Phi_{s,t-1}^{Reputation}\left[\pi_{s,t-1}^{Reputation} - \delta E(\pi_{s,t-1}^{Defection}) - \min\left(\pi_{s,t-1}^{i} - \delta E(\pi_{s,t-1}^{j})\right)\right], \\
A_{s,t}^{Opportunism} &= A_{s,t-1}^{Defection} + \Phi_{s,t-1}^{Defection}\left[\pi_{s,t-1}^{Defection} - \delta E(\pi_{s,t-1}^{Reputation}) - \min\left(\pi_{s,t-1}^{i} - \delta E(\pi_{s,t-1}^{j})\right)\right], \text{ and} \\
A_{s,t}^{Secure} &= A_{s,t-1}^{Secure} + \Phi_{s,t-1}^{Secure}\left[\pi_{s,t-1}^{Secure} - \delta E(\pi_{s,t-1}^{Secure}) - \min\left(\pi_{s,t-1}^{i} - \delta E(\pi_{s,t-1}^{j})\right)\right],
\end{aligned} \tag{C.16}
$$

where t indexes rounds (each consisting of two periods), $A_{s,t}^{i}$ is the attraction that strategy i holds for subject s in round t (with $A_{s,1}^{i}=0$); $\Phi_{s,t-1}^{i}$ is an indicator function equal to 1 if subject s plays strategy i in round $t-1$; $\pi_{s,t-1}^{i}$ is the (two-period overall) profit received by subject s who played strategy i in round $t-1$; $E(\pi_{s,t-1}^{\neg i})$ is the profit subject s would have expected to receive by playing the best alternative strategy to i in period $t-1$; and δ is a weighting function for the strategies not played. Subtracting the minimum possible profit difference both normalizes the attractions and insures that attractions are never negative. If $\delta=0$, the subject weights only payoffs actually received in updating attractions. If $\delta=1$, the subject weights strategies played equally with the forgone earnings of unplayed strategies.

To see the intuition behind these attraction measures, consider a simple example. Suppose an owner-manager of an insecure firm plays the reputation strategy in treatment OM-II. The payoff to the strategy is: $p_0-111+118$. The expected payoff from the opportunism strategy is: $p_0-60+0.05\times 118$. In this case, the lowest possible payoff to any strategy is -53. So, if $\delta=0$, the increment to attraction is: $p_0-111+118+53=p_0+60$. If $\delta=1$, the subject considers the expected payoff to the defection strategy and the lowest difference in expected payoffs between strategies is -58. So, the increment to attraction becomes $p_0-111+118-(p_0-51+0.05\times 118)+58=110.1$.

Given the attractions at the beginning of a round, the propensity to play a strategy is given by:

$$
q_{s,t}^{i} = \frac{A_{s,t}^{i}}{\sum_j A_{s,t}^{j}}. \tag{C.17}
$$

Note that, as defined, higher (lower) payoffs from playing the reputation strategy in the past add to (subtract from) the propensity to play it again. Similarly, lower (higher) payoffs to playing the opportunism strategy in the past adds to (subtract from) the propensity to play the reputation strategy currently. Thus, a greater propensity for reputable behavior is driven by either high past payoffs to the reputation strategy or low past payoffs to opportunism.

Each model assumes that the propensity to make a choice (i.e., the choice's "attraction") depends on the weighted average of the prior profitabilities of playing each choice relative to the prior profitabilities of playing other choices. We calculate two sets of attraction measures: In one set, the propensities depend only on the profitabilities of prior choices, referred to as "Gross Profit EWA." In the second set, the propensities are computed by comparing the outcome of each choice relative to the most profitable feasible alternative choice, defined as "Net Profit EWA."[2] Our estimates employ robust standard errors clustered by subject to control for repeated observations.[3] We present the estimates in Table C.2.

In Table C.2, we present logit models of operator behavior. The estimates in the table suggest that (i) consistent with our model's prediction, adding reform reduces reputable behavior under Parameter Set II and there is no significant difference between the OM and PMR treatments in fostering reputation. Men are somewhat more prone to behave reputably under Parameter Set II.[4] (ii) operators are more prone to choose the reliable technology when this choice has led to higher payoffs in previous rounds. (iii) professional managers are more responsive to past profitability than owner-managers.

C.2.2 Explaining customer conjectures in period 1

For customers, we change the attraction measures slightly because of the continuous strategy space. We measure attraction as either (1) average period 1 net profits received by customers in prior interactions as a result of buying goods or (2) average prior period 1 values observed by customers in prior rounds. Specifically, we define the "attraction" of a good as either:

$$A_{s,t}^{good} = \sum_{n=1}^{t-1} \frac{\Phi_{s,n}^{i}\left(V_{s,n}^{good} - DP_{s,n}^{good}\right)}{t-1} \tag{C.18}$$

or

$$A_{s,t}^{good} = \sum_{n=1}^{t-1} \frac{V_{s,n}^{good}}{t-1} \tag{C.19}$$

[2]We note that, because propensities are probabilities, alternative specifications of these models would be to use the log-odds ratio of the reliable technology choice propensity as the independent variable. While the results mirror those displayed in Table C.2, we lose 28% of the observations because the propensities are either 0 or 1.

[3]Fixed effects models cannot be used because each subject participates in only one treatment.

[4]In unreported regressions aggregating across treatments, men remain somewhat more prone to form reputations than women. Men (women) account for 50.36% (49.64%) of operator decisions for insecure firms. Interactions of gender, treatment and parameter set show no significant coefficients.

Table C.2: *Logit regression analysis of reputable behavior by operators of insecure firms.* The dependent variable equals 1 if an operator chooses the reliable technology. Gender is a dummy variable equal to 1 if the subject is male. Owner-Mgr., Owner-Mgr. w/ Reform and Professional Mgr. are dummy variables for the governance structure treatments. Experience weighted attraction (EWA) regressions include separate propensities to play the reputation strategy in each treatment as defined in equations (C.16) and (C.17) with $\delta = 0$ or $\delta = 1$. Robust standard errors clustered by subject appear in parentheses below each estimate.

	Baseline		Gross Profit EWA		Net Profit EWA	
	Parameter Set		Parameter Set		Parameter Set	
Variable	I	II	I	II	I	II
Constant	0.69***	0.66***	-0.03	-0.57	-0.06	-0.39
	(0.25)	(0.24)	(0.37)	(0.41)	(0.45)	(0.39)
Gender (1=Male)	0.17	0.53*	0.08	0.45**	-0.02	0.56**
	(0.28)	(0.28)	(0.20)	(0.20)	(0.23)	(0.24)
Owner-Mgr.	-0.48*	-0.84***	-0.47	-0.24	-0.30	-0.30
w/ Reform	(0.29)	(0.30)	(0.41)	(0.46)	(0.50)	(0.46)
Professional Mgr.	-0.40	-0.15	-1.21***	-0.77	-1.61***	-1.39
	(0.36)	(0.40)	(0.47)	(0.70)	(0.55)	(0.91)
Exp. Wtd. Propensity			1.39***	2.09***	1.21**	1.26***
x Owner Mgr.			(0.54)	(0.66)	(0.53)	(0.47)
Exp. Wtd. Propensity			1.76***	1.52***	1.17**	1.11***
x Owner Mgr. w/ Reform			(0.54)	(0.51)	(0.54)	(0.63)
Exp. Wtd. Propensity			3.09***	2.94***	3.74***	3.06***
x Professional Mgr.			(0.54)	(0.76)	(0.60)	(0.92)
Obs.	605	667	605	667	605	667
Clusters	66	63	66	63	66	63
Pseudo-R2	0.87%	4.43%	9.53%	11.84%	9.98%	8.50%

”*,” ”**” and ”***” denote significance at the 90%, 95% and 99% levels of confidence respectively.

where $A_{s,t}^{good}$ is subject s's attraction to the good in round t, $V_{s,n}^{good}$ equals the good's value in round n, $\Phi_{s,n}^{i}$ is an indicator function equal to 1 if subject s purchased the good in round n, and $DP_{s,n}^{good}$ is the good's discounted price for subject s in round n. We refer to the first measure as "Average Prior Net Profit Attraction." It averages prior net profits, weighting un-purchased goods by 0 and purchased goods by their value relative to their purchase prices (analogous to $\delta = 0$ in equation C.17). Thus, it considers not just the value of goods purchased, but also the cost of purchasing them. This measure does not capture the effect of goods not purchased. The second, "Average Prior Value Attraction," is the prior average good value. It is the analog of equation (C.16) with $\delta = 1$ because goods not purchased have the same weight as purchased goods.[5]

Table C.3 presents estimates of the effect of treatments, parameter sets and experience on period 1 prices.[6] The estimates uniformly indicate that experience matters. Average Prior Net Profit has a uniformly large and positive effect on prices in all treatments under Parameter Set I. The effect remains positive but is only weakly significant (at best) in Parameter Set II. Average Prior Values always has a large and statistically significant positive effect on prices in both parameter sets. Our earlier evidence indicates that operators' strategies evolve toward adoption of the reputation strategy, especially after experiencing higher profits. This tendency, when combined with the evolution of customer strategies toward setting higher good prices suggests that the feedback between customer and operators' strategies tends to reduce both the systematic deficiency in reputable behavior as well as the substantial undervaluation of goods observed in the experiment.

C.2.3 Explaining customer conjectures in period 2

Table C.4 presents estimates of experience weighted attraction (EWA) models for period 2 prices. Consistent with our model's predictions, revelation results in a significant price drop of around 320 francs under both parameter sets. The direct effects of reform are broadly in line with the model's predictions. Prices are higher for revealed firms in the professional management and owner manager with reform treatments. For unrevealed firms, as the model predicts, period 2 prices are higher only when we add reform to owner management under Parameter Set II. The estimates incorporating EWA models uniformly indicate that experience matters. As is the case with period 1 goods' prices, customers generally tend to raise (lower) their bid prices for goods

[5] We note that this also happens to be the method of moments estimator for a Beta-Binomial Distribution for the probability of a high quality good based on prior observations.

[6] Because, in theory, prices directly convey probabilities of high quality items, we could construct a logistic version of the regressions using OLS on the implied logistic variable defined as $\ln\left(\frac{P/1000}{1-P/1000}\right)$. However, if we do this, we lose 36% of the observations because subjects either set prices at 0 or 1000.

Table C.3: *Censored normal regressions explaining period 1 prices.* Gender is a dummy variable equal to 1 if the subject is male. Round number is the number of rounds at that point in the experimental session. Owner-Mgr., Owner-Mgr. w/ Reform and Professional Mgr. are dummy variables for the governance structure treatments. Attraction regressions include average values and average prior net profits in previous rounds. Robust standard errors clustered by subject appear in parentheses below each estimate. ***, ** and * denote significance at the 99%, 95% and 90% levels of confidence, respectively.

			Average Prior Net Profit Attraction		Average Prior Value Attraction	
	Parameter Set		Parameter Set		Parameter Set	
	I	II	I	II	I	II
Constant	759.76***	704.15***	704.04***	540.03***	416.77**	119.00
	(100.35)	(112.30)	(132.96)	(93.44)	(163.00)	(164.59)
Gender	-135.84	-48.07	-129.39	-74.04	-82.46	-56.42
(1 = Male)	(83.84)	(83.45)	(90.13)	(72.09)	(76.20)	(73.37)
Period	4.94*	8.02**	4.34	7.41**	2.98	8.96**
	(2.87)	(3.42)	(2.91)	(3.78)	(3.07)	(3.59)
Owner-Mgr.	-98.53	-215.63**	-148.14	-115.73	-335.57	54.08
w/ Reform	(111.94)	(105.40)	(166.53)	(106.33)	(262.44)	(188.91)
Professional Mgr.	-143.75	-104.89	-134.22	-28.87	-188.76	168.22
	(97.45)	(112.12)	(132.67)	(127.02)	(199.58)	(201.16)
Attraction			0.18	1.07***	0.42**	0.80***
x Owner-Mgr.			(0.23)	(0.35)	(0.18)	(0.21)
Attraction x Owner-			0.64*	0.44***	0.88***	0.48***
Mgr. w/ Reform			(0.37)	(0.14)	(0.30)	(0.16)
Attraction			0.49	0.41**	0.58***	0.41**
x Professional Mgr.			(0.33)	(0.21)	(0.21)	(0.16)
Obs.	792	756	726	693	726	693
Left Censored	96	76	88	71	88	71
Right Censored	196	191	178	174	178	174
Psuedo R2	0.29%	0.41%	0.70%	1.46%	1.40%	1.35%

"***," "**," "*" denote significance at the 99%, 95% and 90% levels of confidence respectively.

after observing higher (lower) quality goods in the past.

Overall, the analysis shows that (1) the intuition underlying our model survives significant behavioral deviations from optimizing behavior and (2) over time, behaviors moves toward those predicted, increasing the tendency toward the equilibria predicted by our model.

Table C.4: *Censored normal regressions explaining period 2 prices.* Gender is a dummy variable equal to 1 if the subject is male. Round number is the number of rounds at that point in the experimental session. Owner-Mgr., Owner-Mgr. w/ Reform and Professional Mgr. are dummy variables for the governance structure treatments. Attraction regressions include average values and average prior net profits in previous rounds. Robust standard errors clustered by subject appear in parentheses below each estimate. "***," "**," "*" denote significance at the 99%, 95% and 90% levels of confidence respectively.

			Average Prior Net Profit Attraction				Average Prior Value			
			Revealed Firms		Unrevealed Firms		Revealed Firms		Unrevealed Firms	
	Parameter Set		Parameter Set		Parameter Set		Parameter Set		Parameter Set	
	I	II	I	II	I	II	I	II	I	II
Constant	426.53***	250.31**	-205.38	-127.52	384.94**	-65.97	-186.83	-115.73	202.93	-166.68
	(117.85)	(118.65)	(220.85)	(205.00)	(164.58)	(234.78)	(224.01)	(184.05)	(128.10)	(201.56)
Gender	71.04	12.87	38.92	-107.06	84.80	114.26	-8.32	-75.21	93.74	119.41
(1 = Male)	(92.93)	(90.97)	(167.45)	(130.68)	(110.35)	(141.46)	(178.77)	(122.43)	(96.90)	(119.91)
Period	-5.71	-8.31*	4.22	0.81	-0.59	7.83	3.32	3.82	-2.66	4.89
	(3.68)	(4.38)	(6.65)	(9.34)	(5.60)	(6.94)	(6.46)	(8.98)	(5.24)	(6.91)
Owner-Mgr.	60.76	348.09***	333.60**	350.07*	-30.89	398.63*	372.68**	212.47	-172.67	315.48*
w/ Reform	(109.51)	(115.24)	(151.18)	(190.01)	(153.14)	(219.92)	(173.97)	(178.15)	(158.43)	(190.06)
Professional Mgr.	-11.47	216.39*	314.97**	399.74	-111.77	148.95	254.66*	331.64	-16.60	210.43
	(86.71)	(128.11)	(152.90)	(266.20)	(125.23)	(235.90)	(153.35)	(212.54)	(139.42)	(183.71)
Revealed	-327.30***	-314.29***								
	(84.76)	(72.22)								
Attraction			220.41	-3.24**	-0.11	-0.34	-0.08	-0.46	0.71**	1.17**
x Owner-Mgr.			#N/A	(1.50)	(0.37)	(0.52)	(0.50)	(0.61)	(0.33)	(0.47)
Attraction x Owner-			-0.35	0.18	0.47	0.22	0.11	0.34*	1.07***	0.74***
Mgr. w/ Reform			(0.44)	(0.64)	(0.44)	(0.32)	(0.21)	(0.18)	(0.31)	(0.21)
Attraction			1.07*	-0.78	0.25	-0.32	0.46	0.86**	0.30	0.56*
x Professional Mgr.			(0.63)	(0.64)	(0.28)	(0.57)	(0.37)	(0.37)	(0.20)	(0.33)
Obs.	212	736	154	135	482	338	154	135	182	338
Left Censored	96	252	56	53	111	96	56	53	111	96
Right Censored	119	96	12	8	77	60	12	8	77	60
Psuedo R2	0.63%	0.96%	1.69%	1.34%	0.24%	0.62%	1.25%	1.81%	1.52%	1.58%

"***," "**," "*" denote significance at the 99%, 95% and 90% levels of confidence respectively.

References

Axelrod, Robert, "The evolution of strategies in the iterated prisoner's dilemma," in Christina Bicchieri, Richard Jeffrey, and Brian Skyrms, eds., *The Dynamics of Norms*, Cambridge University Press, 1987, chapter 1, pp. 1–16.

Brandts, Jordi and Neus Figueras, "An exploration of reputation formation in experimental games," *Journal of Economic Behavior & Organization*, 2003, *50* (1), 89–115.

Camerer, Colin F, Teck-Hua Ho, and Juin-Kuan Chong, "Sophisticated experience-weighted attraction learning and strategic teaching in repeated games," *Journal of Economic Theory*, 2002, *104* (1), 137–188.

Erev, Ido and Alvin E Roth, "Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria," *American Economic Review*, 1998, *88* (4), 848–881.

Noe, Thomas H., Michael J. Rebello, and Thomas A. Rietz, "Product market efficiency: The bright side of myopic, uninformed, and passive external finance," *Management Science*, 2012, *58* (11).

Norris, James R, *Markov Chains*, Cambridge University Press, 1998.