

Traders vs. Relationship Managers: Reputational Conflicts in Full-Service Investment Banks*

Zhaohui Chen, Alan D. Morrison,
McIntire School of Commerce, Saïd Business School,
University of Virginia. University of Oxford.

William J. Wilhelm, Jr.,[†]
William G. Shenkir Eminent Scholar,
McIntire School of Commerce,
University of Virginia.

September 18, 2012

Abstract

We present a model that explains why investment bankers have struggled in recent years to manage conflicts of interest. The model captures two conflicting dimensions of reputation. On the one hand, banks can build a *type reputation* for technical competence by performing complex deals that may not serve their clients' interest; on the other hand, bankers can sustain a *behavior reputation* by refraining from doing so. Unproven banks favor type reputation over behavioral reputation; being ethical in our model is a luxury reserved for banks that have proven their abilities. The model also sheds light on conflicts between the trading and advisory divisions of investment banks, as well as the consequences of technological change for time variation in the relative strength of behavior- and type- reputational concerns.

*Chen and Wilhelm received support from McIntire Foundation's King Fund for Excellence and the Walker Fund; Morrison received support from the Oxford University Centre for Corporate Reputation.

[†]Corresponding author. McIntire School of Commerce, University of Virginia, Rouss & Robertson Halls, East Lawn, P.O. Box 400173, U.S.A. email:wjw9a@virginia.edu; tel: 434-924-7666; fax: 434-924-7074

1. Introduction

Traditionally, investment bankers advised client firms on capital raising transactions and represented the quality of their securities to prospective investors. Neither function was particularly susceptible to formal contracting because the quality and veracity of advice was not easily verified in court. As a consequence, bankers had an incentive to build reputations for being trustworthy. The long-standing, often exclusive, banking relationships that characterized much of the 20th century suggest that such efforts served bankers' interests as well. This central function of investment banking has changed little over time. But as banks have increased the scale, scope, and complexity of their operations, they have come under increasing criticism for failure to manage the conflicts of interest inherent in their businesses.¹ From this perspective, it is less clear that reputation concerns have remained an effective governance mechanism.

In this paper we study a model designed to explain why investment bankers have struggled in recent years to manage conflicts of interest. Our argument rests on the idea that different dimensions of a bank's reputation may conflict with one another and need not be of equal value to each of a bank's operating units. The analysis contributes to a growing literature that demonstrates how an agent's concern for his own reputation can lead him to take actions that harm his clients.² For example, a junior banker wishing to showcase his abilities as a deal maker might advise his clients to undertake an unnecessary takeover, or to perform an unnecessarily complex structured financing. Alternatively, a trader might pursue a reputation for great skill by breaching so-called "Chinese Walls" within banks that are intended to prevent misuse of sensitive client information.

Reputation is pursued at the expense of clients when there is uncertainty over an *individual's* type. Chen, Morrison and Wilhelm (2011) demonstrate that the incentive distortion associated with personal reputation concerns can be countered by situating the agent in a long-lived firm that can monitor the agent's activities and so prevent it from engaging in socially damaging reputation building. When clients anticipate that the firm will perform this type of monitoring they are prepared to pay more for its services, because they do not

¹Noteworthy examples include the 2003 "Global Settlement" of alleged conflicts of interests between investment bankers and research analysts in 10 prominent banks; allegations that Goldman Sachs and Citigroup, among others, bet against their clients in transactions tied to subprime mortgages; and criticism of Goldman for representing both parties to Kinder Morgan's proposed acquisition of El Paso Corporation. See "Goldman, on Both Sides of Deal, is Now in Court," Steven M. Davidoff, *The Deal Professor*, New York Times Dealbook, February 7, 2012, <http://dealbook.nytimes.com/2012/02/07/goldman-on-both-sides-of-a-deal-is-now-in-court/>.

²See for example Ely and Välimäki (2003), Ely, Fudenberg and Levine (2008), and Morris (2001). In these models, as in ours, the agent knows better than the client how to best serve the client's needs. Although our models are very different, Bolton, Freixas and Shapiro (2007) study the conflict between banks advising clients and selling them products to meet their needs. Their focus on the interaction between (even small) reputation concerns and competition among banks yields the novel conclusion that independent advisory firms need not be the only credible source of unconflicted advice.

anticipate having to subsidize the creation of individual reputations. Hence, the desire of the firm’s owners to sustain its *institutional* reputation can serve to counter the bad incentives associated with individual reputation building.³

In our model banks have an incentive to build a type reputation for unique skill that will be valuable in the future at the expense of their current clients. On the other hand, banks can charge higher fees if they earn a reputation for resisting the urge to behave in this fashion. Thus there is a conflict between a *type reputation* that reflects client beliefs concerning a bank’s abilities and a *behavior reputation* that reflects the bank’s equilibrium actions, and is supported by the equilibrium beliefs of its clients. Behavioral reputations are closely associated with the type of client-focused behavior that is usually characterized by market observers as “ethical.” While one can think of a firm whose equilibrium behavior sustains a behavioral reputation as an ethical firm, we acknowledge that ethical actions involve more than a desire to maintain a profitable behavioral reputation. Nevertheless, we suspect that this is the basis of much market behavior that attracts the “ethical” soubriquet.

In the basic model, the bank has a single division and knows both its type and the state of the world when it takes an action on behalf of the client. Clients observe their payoffs but neither the state of the world nor the banks action. As a consequence, the client receives only a very noisy signal regarding whether the bank’s action was self-interested. Our main result is that, when type-reputations are long-lived, behavioral reputational incentives are insufficiently strong to overcome the incentive to build a type reputation even when banks are very patient.⁴ It is therefore unrealistic to expect an unproven bank with a weak type reputation to adopt ethical behavioral patterns; being ethical in our model is a luxury reserved for banks that have proved their abilities.

When clients expect unproven banks to focus on enhancing their type reputations they pay them accordingly, and do not punish them for their actions. We exhibit an equilibrium for our model in which unproven banks engage in a type-reputation-building phase before entering a phase where they behave, and are expected to behave, ethically. In that phase the consequence of a transgression is a permanent loss of faith, and a reversion to reputation-building behavior.

If financial innovations are evidence of individual skill they could serve as a means of establishing type reputations. Our reasoning therefore suggests that less well-established banks without a behavioral reputation to lose will be a more common source of financial

³Conversely, because they can profit immediately from a share of the superior returns a senior agent reaps from his individual reputation, very impatient owners may lack incentive to build an institutional reputation.

⁴In contrast to Ely and Välimäki (2003) and (Chen et al., 2011), knowledge is transferred between junior and senior workers. As a result, bank types do not change with each new generation of workers, so that the gains from type reputation formation are long-lived. Hence, even a very patient bank has something to gain from building a type reputation, and institutional reputation is correspondingly harder to sustain.

innovation.⁵ But this also implies that at least some innovative activity will be inefficient.⁶ Similarly, banks with strong reputations for placing their clients' interests' first should be more likely to adopt a "fast follower" policy.⁷

We next present an extension of the model intended to reflect conflicts in reputation concerns across bank business units. As we suggested earlier, the advisory focus of traditional investment-banking functions remains relatively less susceptible to formal contract and thus dependent on a behavioral reputation for placing client interests first. In contrast, trading and brokerage functions are relatively more susceptible to performance measurement, formal contract, and the signaling of ability. As a consequence, we might expect these functions to place type-reputation concerns before concerns for behavioral reputation.

We formalize the conflict by assuming that the bank has independent execution and advisory divisions. The advisory division observes the state of the world and advises the client on the action it should request from the execution division. It does not know the execution division's type. The execution division knows its type but does not observe the state of the world. In this setting, the advisory division concentrates upon behavioral reputation, and so it advises the client of the appropriate action to request from the execution division. This facilitates client monitoring of the execution division and, as a result, we demonstrate that a socially efficient equilibrium without a type-reputation-building phase exists. One way to interpret this result is as a justification for Chinese Walls designed to limit the flow of information within banking organizations.

However, even with this institutional setup this is not the only possible outcome; an alternative equilibrium with a short type-reputation-building phase exists, and may be preferred by the execution division to the socially optimal one. Moreover, close contact between the execution and the advisory business, and in particular sharing of information about the execution division's type, destroys the socially optimal equilibrium.

The inefficient equilibrium exists under conditions that could arise when a strong advisory function operates alongside a less well-established execution function or one that has lost key people with whom the firm's type reputation is associated. The flurry of commercial bank acquisitions of investment banks during the late 1990s joined a number of well-established advisory operations with the nascent execution divisions of their acquirers.⁸

⁵Consistent with this prediction, both Drexel Burnham Lambert (junk-bond market) and Bankers Trust (over-the-counter (OTC) derivatives and the "originate to distribute" lending model) had a modest presence in the traditional investment-banking functions that are more dependent on client relationships. Similarly, J.P. Morgan's most innovative efforts in OTC derivatives and structured finance were undertaken before the bank established a dominant investment-banking presence.

⁶On a similar note, see Glode, Green and Lowery (2012) for a model in which investment in financial expertise is socially wasteful.

⁷Ellis (2009, ch.10) identifies Goldman Sachs having adopted this strategy as an explicit byproduct of John Whitehead's 1956 study of the firm's efforts to generate new business. The study ultimately led to the separation of relationship management from execution functions.

⁸In 1997 alone, NationsBank acquired Montgomery Securities, ING Group acquired Furman Selz, Bankers

Similarly, bankers routinely spent their entire careers with a single firm prior to 1970 but mobility ratcheted up over the next several decades to the point where by the late 1990s entire teams were changing banks.⁹ From this perspective, it is perhaps not surprising that advisory clients expressed growing concerns that private information about their operations or transaction strategy might be shared with competitors or with the bank’s own trading operations nor is the appearance of specialized or “boutique” advisory firms that appeal to such concerns by emphasizing their freedom from conflicts.¹⁰

Finally, we consider the effect upon our results of periodic technological shocks that reset the bank’s type. These shocks alter the dynamics of reputation formation in our model, because they give the bank repeated opportunities to announce its type, and so to build a behavioral reputation for truth-telling. This reputation is so valuable in the long run that first best can be achieved. However, we also show that, if a technological shock carries with it a significant risk of bank failure, the consequential attenuation of the long-run value of a truth-telling reputation prevents it from supporting a first-best equilibrium.

The role of technological change in our model sheds further light on several long-run patterns in the financial markets. For example, until the middle of the 20th century, bank/client relationships were quite stable and often exclusive.¹¹ Relationship stability is consistent with banks having both strong type- and behavioral- reputations. In other words, clients should have little incentive to switch banks if they perceive their bank as having acted in their best interest in the past and maintaining the skills necessary for the transaction at hand. The first evidence of weakening client relationships corresponded with advances in batch processing computing technology that had its greatest impact on banks with large brokerage operations. The technological shock was profound as evidenced by widespread failure during the late 1960s among banks that were slow to adapt.¹²

In addition to weakening client relationships, the industry witnessed a steady decline in the average tenure of investment bank partners.¹³ If behavioral reputation stems from

Trust acquired Alex. Brown, CIBC acquired Oppenheimer, BankAmerica acquired Robertson Stephens (and sold it to BankBoston in 1998), and SBC Warburg acquired Dillon Read. See Morrison and Wilhelm (2007, Ch.10) for details.

⁹See Morrison and Wilhelm (2007, Ch.9)

¹⁰Using post-1975 data, Asker and Ljungqvist (2010) find that large firms avoid developing relationships with banks that represent product-market competitors. This finding stands in sharp contrast to the high level of industry specialization among investment banks during the first half of the 20th century (see Morrison and Wilhelm (2007, Ch.7) and Carosso (1970)). The academic evidence on conflicts related to the provision of advisory services is mixed. For example, Bodnaruk, Massa and Simonov (2009) report evidence of banks taking positions in the targets of bidding firms which they advise. In contrast, Griffin, Shu and Topaloglu (2012) find no evidence that banks advising in corporate takeovers share client information with institutional investors. Greenhill and Evercore, both founded in 1996, were early proponents of specialized firms as an antidote to conflicts of interest.

¹¹See Morrison and Wilhelm (2007, Ch.8).

¹²See Morrison and Wilhelm (2007, 2008).

¹³See Morrison and Wilhelm (2007, Ch.9).

consistently client-centric behavior on the part of individual bankers, then its preservation should, in part, rest on banks' ability to sustain intergenerational transfer of client relationships. But Morrison and Wilhelm (2008) suggest that their ability to do so declined with increasing scale during the 1970s and 1980s. If behavioral reputation concerns consequently weakened, our model suggests that banks would be more inclined to engage in activities that might be perceived as a threat to their clients' interests. In contrast, Goldman Sachs was relatively slow to join the race to achieve greater scale and scope and the average tenure of its partners declined less rapidly than for many of its peers. It is perhaps not surprising then that Goldman openly avoided representing bidders during the 1980s hostile takeover movement in deference to concerns for their client relationships.

The rest of the paper is organized as follows. Section 2 reviews some of the relevant economic literature on reputation and relates it to our work. Section 3 presents our model and Section 4 defines a model equilibrium. Section 5 shows how reputation is managed in equilibrium, and Section 6 shows how the creation of an independent advisory division within the bank can improve outcomes. Section 7 presents our results on technological shocks in our set up. Section 8 concludes.

2. Literature Review

The traditional approach to reputation modelling focused on type reputation models. Early models involved the introduction of "commitment types:" agents who always took actions that led to surplus-maximizing outcomes, even when those actions were irrational. Kreps, Milgrom, Roberts and Wilson (1982) demonstrated even a small probability of commitment types was sufficient to induce cooperation in a finitely repeated prisoner's dilemma game. Fudenberg and Levine (1992) demonstrate that when there is a small probability of commitment types, an infinitely-lived rational and self-interested agent who plays the prisoner's dilemma with a series of short-lived agents can overcome his moral hazard problem by pretending to be a commitment type: that is, by building a reputation for cooperation.

Fudenberg and Levine's insight is the basis for a number of finance papers, in which type could reflect the standard preference for cooperation (Boot, Greenbaum and Thakor (1993), Fulghieri, Strobl and Xia (2010), Winton and Yerramilli (2011)), or ability (Diamond (1989), Chemmanur and Fulghieri (1994*a*, 1994*b*)).¹⁴ However, while type reputation concerns resolve moral hazard problems in all of these models, reputation effects are short-lived (Cripps, Mailath and Samuelson, 2004). Hence, this type of model seems a poor choice for examining institutional reputation.

¹⁴Also see Hartman-Glaser (2012) for a model in which there is tension between a bank's reputation for honest representation of the type of an asset used in a securitization and the bank's ability to signal asset type by retaining a fraction of the asset. While reputation concerns can increase the bank's equilibrium payoffs by reducing retention, this does not imply that the bank is more likely to perfectly reveal information.

The behavioral approach that we adopt for institutional reputation is introduced by Kreps (1990) in an overlapping generation model in which there is moral hazard but no adverse selection. “Reputation” in this type of model refers to an equilibrium belief by clients that an agent will take specific actions; it is sustained by the consequences of the beliefs that the clients will adopt if the wrong actions are taken. Clients monitor actions imperfectly in this model, as in ours.¹⁵ If the monitoring is sufficiently informative then a general folk theorem applies, in which future punishment and rewards provide the incentives needed to sustain the equilibria (Fudenberg, Levine and Maskin, 1994).

Later work extended Kreps’ approach by combining both type reputation (adverse selection) and behavior reputation (moral hazard) (see Tadelis (1999, 2002)). Tirole (1996) examines collective reputations in this way, and shows that bad equilibria can persist because a moral hazard in teams problem renders individuals unwilling to improve the collective reputation. Morrison and Wilhelm (2004) examine professional services firms as entities that certify agents of uncertain type and that use a collective reputation to incentivize senior agents to train juniors.

In our model there is non-trivial learning in equilibrium. This is a difficult problem because such games do not admit the recursive equilibrium structure of Abreu et al. (1990). Fudenberg and Yamamoto (2010) address this type of problem by addressing a simpler set of equilibria in which a player’s best response does not depend upon his beliefs (and, hence, upon his learning), and so are able to re-introduce a recursive equilibrium structure. Fudenberg and Yamamoto are able to prove a folk theorem for their restricted class of equilibria. But none of the papers in this strand of the literature examines reputational conflict.

Finally, our model is related to the relational contract literature, particularly in Section 6, where we allow banks to announce their types. In contrast to earlier work in this area (see for example Levin (2003), Athey, Bagwell and Sanchirico (2004), Athey and Bagwell (2008)), which considers only long-run players, we consider the case where clients are short-lived. This has two important consequences. First, the undesirable incentive to build a personal reputation is particularly important because fees are based entirely on the bank’s type reputation. In contrast, if clients are long-lived, there may be equilibria in which the bank’s payoff is independent of its type reputation, so that the bad reputation incentive is reduced (see Athey and Bagwell (2008)). Second, the relational contract is more limited than Levin’s (2003), in which contracting parties sometimes make large punishment payments. Such contracts are impossible in our setting because there is no mechanism to induce our short-lived clients to make such payments. Moreover, they will pay a higher fee if they anticipate a large payment from the bank, and so will destroy the incentive effects of such

¹⁵The seminal treatment of this topic is due to Abreu, Pearce and Stacchetti (1990).

payments.

3. Model

We consider an infinitely lived bank that is created at time 0. Time is indexed by $t = 0, 1, 2, \dots$, and period t runs from time $t - 1$ to t . In each period the bank has a one owner, one worker, and a single client. All agents in our model are risk neutral, and they have a common per-period discount factor δ . The owner and the worker work for one period and retire; a new client deals with the bank in each period.

In each period the worker takes an action $A \in \{1, 2\}$ on behalf of the client. The cost of taking either action is 0 to the worker, but the payoff that the client derives from A depends upon the state of nature $\omega \in \{1, 2\}$. Action ω is optimal for the client in each state ω . When state 1 realizes, the client receives a payoff of 1 from action 1 and $x \in (0, 1)$ from action 2. When state 2 realizes, the client's payoff after action 1 is x or 0 with respective probabilities q and $1 - q$; it is x if action 2 is selected. In each period, the probability that state 1 realizes is $p > 0$.

There are two types of workers. *Dumb* (type D) workers are only able to perform action 1; *smart* (type S) workers can perform action 1 or action 2. We assign the worker's type (S or D) to the bank, and for convenience we will refer to S-banks and D-banks. Whether or not a worker is smart depends upon his knowledge and the training that he received before starting work. The time 0 probability that the first worker is smart is θ_0 . Subsequent generations of workers are trained by the outgoing worker, so that the bank's type is maintained. The intergenerational knowledge transfer between workers is observed by the new owner, who then makes a take-it-or-leave-it offer to the old owner for ownership of the bank, and hence captures all of its value. It is a consequence of our assumptions on training and its observability that the bank's type is known to its owner and its worker; we assume that this information is not revealed to the client.

We write $z \in \{0, x, 1\}$ for the client's payoff in any period. We assume that z is observable but that none of z , θ and A can be verified in court. It follows that the only possible contract between the bank and the client requires the payment of a fixed fee by the worker in exchange for an unspecified action by the worker. We assume that there is competition among clients for the bank's services, so that the bank captures all of the expected client surplus from its action. The bank's revenue is shared between the worker and the owner; for simplicity, we assume an exogenous sharing rule that assigns a proportion ρ of the surplus to the owner. Without loss of generality, we set $\rho = 1$.

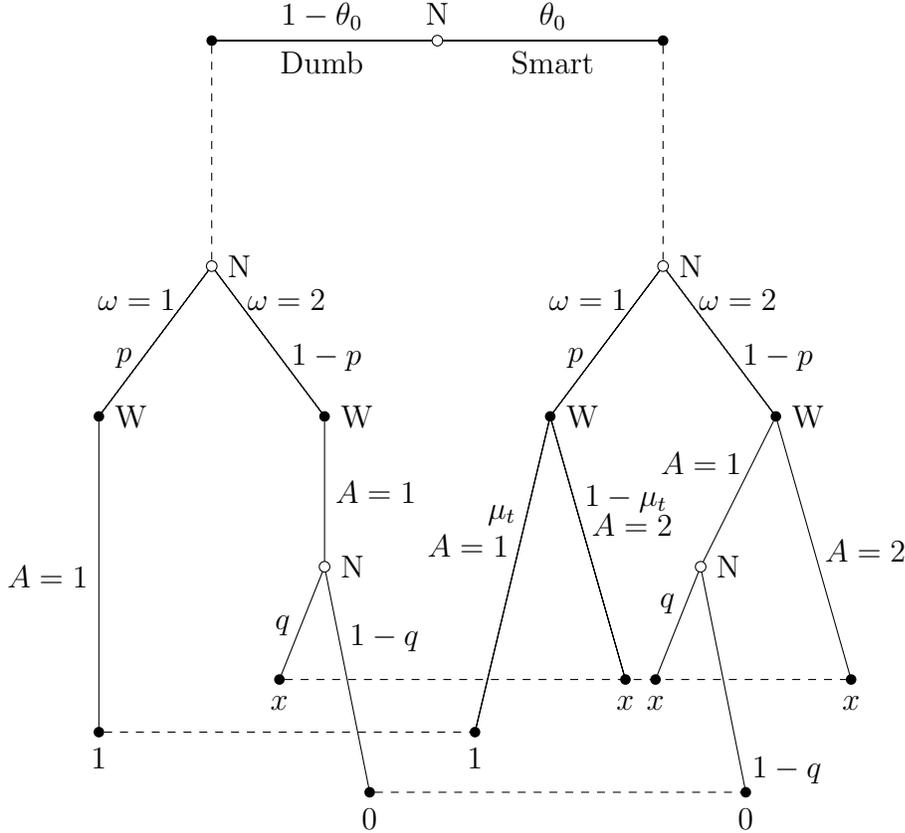


Figure 1: **Extensive form for the stage game.** Nature selects the bank’s type at time 0, and the bank retains that type. Payoffs are stage game client surpluses. The client cannot distinguish between S- and D- banks, and therefore has the information sets identified in the Figure by dashed horizontal lines.

One of the stage games is illustrated in Figure 1. As indicated in the Figure, the bank’s type is established at time 0. This information is not observed by the client, who therefore cannot distinguish between the elements of the information sets indicated in the Figure by horizontal dashed lines. Note that, because the worker has a one period career, he has no incentive to build an individual reputation. We can therefore focus upon the creation of the bank’s reputation. To maintain this focus, we assume that the owner can costlessly control the worker’s action, and, hence, we refer in this section to the bank’s actions, rather than to those of the worker or the owner. Finally, we assume that S-bank’s services are more valuable to a client than a D-bank’s:

$$x > p + (1 - p)qx. \quad (1)$$

The left hand side of equation (1) is the lowest single period client surplus generated by a smart bank, and the right hand side is the expected surplus created by a dumb bank. It is immediate from equation (1) that the bank can maximize its value by building a reputation

for being smart. It can accomplish this by selecting action 2 wherever possible; in particular, it will do so in state 1, in which case its reputation comes at the expense of its client.

4. Equilibrium Definition

We will exhibit Bayesian Nash Equilibria of our model. We start by presenting a formal definition of these equilibria. First, recall that the only publicly available datum in any period is the client's payoff $z \in \{0, x, 1\}$. We therefore define a *time t history* to be an element h_t of $H_t = \{0, x, 1\}^t$. We define H to be the set of all possible game histories to every possible t :

$$H = \bigcup_0^{\infty} H_t. \quad (2)$$

Let $A_t \in \{1, 2\}$ be the bank's time t action. When the bank selects its time t action it knows the history h_t of client payoffs and its own type. Hence, given a time t history h_t , a *strategy* μ for the bank assigns a probability μ_t of setting $A = 1$ when $\omega = 1$:

$$\mu : H \longrightarrow \mathfrak{R}[0, 1] : h_t \longmapsto \mu(h_t) \equiv \mu_t. \quad (3)$$

Note that strategy choice is non-trivial only for S-banks, because D-banks are technologically limited to the strategy that sets $h_t \equiv 1$. We write M for the set of strategies.

We define the client's *action belief* to be the probability β_t that she assigns at time t to the event that the bank takes action 1 given that the state ω is 1. The client conditions her action belief upon H_t :

$$\beta : H \longrightarrow \mathfrak{R}[0, 1] : H_t \longmapsto \beta_t. \quad (4)$$

We write B for the set of action beliefs. The client's action belief reflects the time t probability θ_t that she assigns to the event that the bank has type S. We can think of θ_t as the bank's *reputation*. Like the action belief, the type belief is conditioned upon the payoff history:

$$\theta : H \longrightarrow [0, 1] : h_t \longmapsto \theta(h_t) \equiv \theta_t. \quad (5)$$

Let Θ be the set of all possible type beliefs. Type beliefs and action beliefs give rise to an expected revenue function R for the bank:

$$R : \Theta \times H \times B \longrightarrow \mathfrak{R} : (\theta, h_t, \beta) \longmapsto (1 - \theta_t) \nu_D + \theta_t \nu_S(h_t, \beta), \quad (6)$$

where $\nu_D(h_t, \beta)$ and $\nu_S(h_t, \beta)$ are the respective time t expected client surpluses generated by D- and S- banks:

$$\nu_D(h_t, \beta) = p + (1 - p)qx; \quad (7)$$

$$\nu_S(h_t, \beta) = p\beta_t + p(1 - \beta_t)x + (1 - p)x \quad (8)$$

$$= x(1 - p\beta_t) + p\beta_t. \quad (9)$$

Now consider the period t owner of an S-bank. He receives all of the period t revenue that the bank generates, and sells the bank at the end of the period for its time t expected value. We can therefore define the S-bank's *value function* $V : H \times B \times M \times \Theta \rightarrow \mathfrak{R}$ as follows:

$$\begin{aligned} V(h_t, \beta, \mu, \theta) &= R(h_t, \beta) + \delta \mathbb{E}_{S,t} [V(\langle h_t, z_t \rangle, \beta, \mu)] \\ &= R(h_t, \beta) + \delta \{p\mu_t V(\langle h_t, 1 \rangle, \beta, \mu) + (1 - p\mu_t) V(\langle h_t, x \rangle, \beta, \mu)\}, \end{aligned}$$

where we write $\langle h_t, z \rangle$ for the history obtained by augmenting h_t with the payoff z .

Our equilibrium concept is presented in Definition 1:

DEFINITION 1. *An equilibrium comprises an action belief β , a strategy μ and a type belief θ such that:*

1. *For every h_t ,*

$$\mu_t \in \arg \max_{\mu_t} p\mu_t V(\langle h_t, 1 \rangle, \beta, \mu_{t+1}, \theta) + (1 - p\mu_t) V(\langle h_t, x \rangle, \beta, \mu_{t+1}, \theta);$$

2. *For every t , $\beta_t = \mu_t$;*

3. *θ is obtained from θ_0 by Bayes' Law.*

Like many other infinitely repeated games, our model admits many equilibria. We focus upon the equilibria that generate the highest payoff for S-banks.¹⁶ We define $V^*(\hat{\theta})$ to be the maximum equilibrium bank value given investor prior $\hat{\theta}$; note that V^* is time-homogenous. We assume that $V^*(\hat{\theta})$ can always be attained. Then, if (β, μ, θ) is the time t value-maximizing equilibrium, we must have

$$V^*(\theta(h_t)) = R(h_t, \beta_t, \mu_t) + \delta \mathbb{E} V^*(\theta(\langle h_t, z_t \rangle)). \quad (10)$$

¹⁶An alternative approach would be to focus upon those equilibria that maximize social surplus. These approaches are similar for patient banks (those for which δ is close to 1).

5. Equilibrium Reputation Management

5.1. Reputation Building

Recall that the client payoff is x with probability 1 if an S-bank selects action 2, whereas D-banks can achieve this payoff only in state 2, and then only with probability q . S-banks can therefore increase their reputation early in their lives by adopting a strategy of selecting $A = 2$ for sure. Lemma 1 demonstrates that the number of periods for which an S-bank has to adopt this strategy to achieve a reputation $\bar{\theta}$ is bounded above for any given θ_t .

LEMMA 1. *An S-bank can increase its reputation from θ_t to any large $\bar{\theta} < 1$ by choosing action 2 for a fixed $T(\theta_t, \bar{\theta})$ periods.*

Proof: Given a time t reputation θ_t and a realized client payoff x , the time $t + 1$ reputation θ_{t+1} is derived from θ_t by Bayes' Law as follows:

$$\theta_{t+1} = \frac{\theta_t [p(1 - \mu_t) + (1 - p)]}{\theta_t [p(1 - \mu_t) + (1 - p)] + (1 - \theta_t)(1 - p)q} \quad (11)$$

Let

$$\gamma_{t+1} \equiv \frac{\theta_{t+1}(x)}{1 - \theta_{t+1}(x)}$$

be the time $t + 1$ likelihood ratio for an S-bank, conditional upon a time t client payoff of x . We can write

$$\begin{aligned} \gamma_{t+1} &= \frac{\gamma_t}{q} \left[\frac{(1 - \mu_t)p + (1 - p)}{1 - p} \right] \\ &\geq \frac{\gamma_t}{q}. \end{aligned}$$

Hence

$$\begin{aligned} \log \gamma_{t+1} &\geq \log \gamma_t + \log(1/q) \\ &\geq \log \gamma_0 + (t + 1) \log(1/q). \end{aligned}$$

It follows that $\log \gamma_{t+T} \geq \log \gamma_t + T \log(1/q)$. The result is therefore immediate with

$$T(\theta_t, \bar{\theta}) = \left\lceil \frac{\log \frac{\bar{\theta}}{1 - \bar{\theta}} - \log \gamma_t}{-\log q} \right\rceil + 1,$$

where $\lceil \cdot \rceil$ is the truncation function. □

5.2. *Type Reputation, Behavior Reputation, and Ethical banks*

Lemma 1 identifies one of the dynamics that drives our results: namely, the possibility that an S-bank could engage in a period of *reputation building*, so as to reap the benefits of a high reputation in the future. This effect applies in classical models of reputation in which an agent's reputation relates to the his *type*.

An agent's type reputation reflects its counterparties' beliefs about its *capabilities*. We are concerned in this paper with the extent to which an S-firm's incentive to build its type reputation is tempered by its desire to sustain a *behavior reputation*. An agent's behavior reputation reflects the outcomes that its counterparties have experienced: it derives from the agent's *choices*, rather than its abilities. In our model, an S-bank builds a behavior reputation by choosing action 1 in state 1 so that clients experience high payoffs. Choices of this type are usually characterized as ethical, and we refer to an S-bank that takes them as *ethical*.¹⁷

DEFINITION 2. *An S-bank's equilibrium behavior is ethical precisely when it takes action 1 in state 1.*

Ethical behavior as we have defined it is an equilibrium phenomenon. Moreover, it is inconsistent with type reputation formation, which, as noted in Lemma 1, requires an S-bank repeatedly to take action 2 in state 1. S-banks therefore face a choice between building a type reputation, and behaving ethically. Since ethical behavior in our model is a self-interested phenomenon, it will arise in equilibrium when it is sufficiently profitable.

An S-bank's equilibrium behavior will be ethical precisely when the following incentive compatibility constraint is satisfied:

$$V_{t+1}(\langle h_t, 1 \rangle, \beta, \mu, \theta) \geq V_{t+1}(\langle h_t, x \rangle, \beta, \mu, \theta). \quad (12)$$

Note that Bayes' Law (part 3 of Definition 1) gives us the following in any equilibrium.

$$\theta(\langle h_t, 1 \rangle, \beta) = \theta_t(h_t, \beta). \quad (13)$$

Hence, when Equation (12) is satisfied, Equations (10) and (13) imply that

$$\begin{aligned} V^*(\theta(h, \beta)) &= \frac{R(h_t, \beta_t = 1, \mu_t = 1)}{1 - \delta} \\ &= \frac{1}{1 - \delta} [\theta \nu_S(1) + (1 - \theta) \nu_D], \end{aligned} \quad (14)$$

¹⁷We acknowledge in the Introduction that ethical choices concern more than a self-interested decision to build a behavior reputation.

where ν_D and ν_S are the respective client surpluses that D- and S- banks generate (see equations (7) and (8)).

Our analysis is often concerned with limiting effects as $\delta \rightarrow 1$. In the limit the right hand side of equation (14) is infinite. It is therefore convenient to make the following definition:

DEFINITION 3. *The time t standardized bank value v_t is $(1 - \delta) V_t$.*

For a bank with value V , the standardized bank value is the per-period constant revenue v whose value is V .

Equation (14) immediately yields Proposition 1:

PROPOSITION 1. *The standardized bank value v_t of an ethical S-bank satisfies $v_t(\theta) \leq \theta \nu_S(1) + (1 - \theta) \nu_D$.*

Equation (1) implies that $\nu_S(\mu_t) > \nu_D$ for every μ_t . Proposition 1 therefore implies that an ethical S-bank's value cannot be high for small θ . The reason is that, in order that the incentive compatibility constraint (12) be satisfied, the consequences for equilibrium client beliefs of an S-bank's decision to build type reputation by selecting action 2 in state 1 have to be particularly severe. This "punishment" is so severe as to render high bank value impossible.

S-banks behave ethically in equilibrium only if it is sufficiently profitable to do so. Proposition 1 indicates that ethical behavior is not profitable for banks with low type reputation. A natural question is therefore whether or not an ethical equilibrium exists for low θ banks. To answer this question we need the following technical result:

LEMMA 2. *Given any $\epsilon, \psi > 0$ there exists an integer $K(\epsilon, \psi, \theta_0)$ such that*

$$\Pr[\#\{t > 0 : \theta_t < 1 - \psi\} \leq K(\epsilon, \psi, \theta_0) | S\text{-bank}] \geq 1 - \epsilon.$$

Lemma 2 states that with high probability, an S-bank's equilibrium reputation deviates from 1 for only a fixed number of periods. The result is proved in the Appendix; its intuition is that in expectation an S-bank reveals some information about its type with strictly positive probability in each period. Hence, over a high number of periods, its type is revealed with very high probability.

The following result is a consequence of Lemma 2.

LEMMA 3. *Let $\bar{\theta} < 1$ be any type reputation. Then a sufficiently patient S-bank has a standardized value v that is bounded below by a value very close to $\bar{\theta}x + (1 - \bar{\theta})$. Formally,*

$$\forall \bar{\theta}, \gamma \in (0, 1) \exists \delta_{\bar{\theta}, \gamma} \text{ s.t. } v_t(\theta) \geq \bar{\theta}x + (1 - \bar{\theta}) - \gamma.$$

Lemmas 2 and 3 now yield Proposition 2:

PROPOSITION 2. *If there exists $\bar{\theta} < 1$ such that $\bar{\theta}x + (1 - \bar{\theta}) > \theta_t\nu_S(1) + (1 - \theta_t)\nu_D$ then there exists $\underline{\delta} < 1$ such that whenever $\delta > \underline{\delta}$ there does not exist an equilibrium with $\mu_t = 1$ that maximizes the S-bank's value.*

Proposition 2 demonstrates that there is no equilibrium in which an S-bank with a low type reputation behaves ethically: the value of building a high type reputation in this situation is so great that no set of client beliefs generates sufficiently large punishments to deter unethical behavior designed to build a personal reputation.

This result generates some insight into the dynamics of reputation building and maintenance in young and unproven banks. In the absence of evidence to the contrary, the clients of such banks assign them a low type reputation θ_0 . Because θ_0 is low, no equilibrium client beliefs can induce the bank to behave ethically. In other words, the concern for behavioral reputation that drives the results of Morrison and Wilhelm (2004) is unsustainable in our model because it conflicts with the bank's desire to build a type reputation.

In light of these remarks, we now investigate equilibria in which the bank engages in a period of type reputation building at the beginning of the game, before behaving ethically in order to sustain its behavioral reputation. To do so, pick a θ^* that is sufficiently high to ensure that there is no $\bar{\theta}$ such that $\bar{\theta}x + (1 - \bar{\theta})\nu_D > \theta^*\nu_S(1) + (1 - \theta^*)\nu_D$. Such a θ^* does not satisfy the criteria that Proposition 2 identifies as precluding an ethical equilibrium; as yet, of course, we have not proved the existence of an ethical equilibrium for θ^* .

The equilibrium that we exhibit has three behavior states, which we label “phases”:

1. The *reputation-building* phase lasts for the first $T(\theta_0, \theta^*)$ periods. During the reputation-building phase, $\mu_t = \beta_t = 0$;
2. The bank enters the *behavioral reputation phase* at the end of the reputation-building phase. The bank behaves ethically in the behavioral reputation phase: that is, $\mu_t = \beta_t = 1$. The bank remains in the behavioral reputation phase for as long as $z_t \neq x$. If $z_t = x$ then a public randomization device determines the game's phase in the succeeding period. With probability π_{t+1} the bank enters the punishment phase; with complementary probability it remains in the behavioral reputation phase;
3. During the *punishment phase* $\mu_t = \beta_t = 0$. Once the bank enters the punishment phase it never leaves it.

Reputation considerations follow a well-defined life cycle in this equilibrium. Young banks concentrate upon building type reputation. Their behavior is anticipated and tolerated by their clients, and it is priced accordingly. After the reputation-building phase has played out, banks maintain a reputation for ethical behavior for as long as possible. Their clients anticipate this, and they pay a correspondingly higher fee for services. When clients have a

poor experience from a bank that they expect to behave ethically the bank loses its reputation with some probability π_{t+1} .

It is obvious that, once the bank enters the punishment phase, its clients' belief $\beta = 0$ ensures that it will never earn any credit for ethical behavior and, hence, it will adopt the strategy $\mu = 0$. Similarly, in the reputation-building phase, it similarly gets no credit for ethical behavior, and so does not adopt it. Hence, it is sufficient to demonstrate that for every time t in the behavioral reputation phase of the game, there exists a π_{t+1} for which the incentive compatibility constraint (12) is satisfied. To see this, note that the bank's standardized value in the punishment phase is less than x ; its value if it realizes $z = 1$ is unchanged in the behavioral reputation phase, because no updating of type reputation occurs (see equation (13)), and its value increases if it remains in the behavioral reputation phase after $z = x$, because it then receives a boost to its type reputation. A standard continuity argument then implies that there exists a π_{t+1} that ensures that the IC constraint binds at time t .

We have therefore proved the following result:

PROPOSITION 3. *Let θ^* be sufficiently high to ensure that there is no θ with $\bar{\theta}x + (1 - \bar{\theta})\nu_D > \theta^*\nu_S(1) + (1 - \theta^*)\nu_D$. Then there exists an equilibrium comprising a reputation-building phase, a behavioral reputation phase, and a punishment phase, as detailed above.*

Proposition 3 demonstrates that S-banks can be ethical and also value-maximizing when θ is sufficiently high; Proposition 2 demonstrated that this is not possible for low θ . The intuition for this result is that, when reputation is close to 1 the bank can generate little benefit from a higher reputation, so that the scale, and hence the efficiency costs, of the punishment required to sustain ethical equilibrium behavior is limited. The costs of sustaining ethical behavior are much larger for banks that have more to gain from a higher type reputation. Note finally that in the equilibrium of Proposition 3, the S-bank's reputation never drops below θ^* in the behavioral reputation phase: it either remains constant (when $z = 1$) or it increases (when $z = x$ and it does not enter the punishment phase). Furthermore, the likelihood of entering the punishment phase decreases as θ increases and the costs of doing so increase.

6. Institutional Design

The problems of Section 5 arise because the bank has to choose between the benefits of being perceived to be smart, and of being perceived as ethical, and there is conflict between the actions required to create the associated type and behavioral reputations. In this Section, we examine a possible institutional solution to this conflict.

Suppose that the bank is split into an *advisory* and an *execution* division.¹⁸ The advisory division has perfect information as to the state of the world. The execution division does not know the state of the world, and it is capable of taking an action $A \in \{1, 2\}$. As in Section 5 the execution division can be smart, in which case it can take action 1 or 2, or dumb, in which case it can only take action 1. The execution division knows its type, but neither the advisory unit nor its clients do. The advisory unit observes the state of the world ω and tells the client what services he requires. The client then requests an action of the execution division. The client cannot observe the execution unit's type or the action that it takes, but the client's payoff z is common knowledge. Finally, we make the simplifying assumption that, if the advisory division announces state 2 to the client, the execution division is only able to take action 2. Our qualitative results would not be affected by a relaxation of this assumption, but they would be much more complex.

As in Section 3, there are many potential clients, each of whom bids for the services of each division. We start by considering the (out-of-equilibrium) case in which the client wins only the execution division's services. As the execution division does not know ω , by assumption (1), it is optimal for the client if the execution division always takes action 2; as this action also builds the execution division's type reputation there is no conflict between the client and the bank. The fee for the execution division in this case is

$$\underline{\phi}(\theta_t) \equiv \theta_t x + (1 - \theta_t) \nu_D.$$

We will verify that, when the client retains both the advisory division and the execution division together, it is able to use the advisory division's information to ensure that the execution division takes the action that maximizes the client's income. The two units together therefore generate $\phi(\theta_t)$ of value, where

$$\phi(\theta_t) = \theta_t \nu_S(1) + (1 - \theta_t) \nu_D.$$

We assume that a proportion $\lambda \in (0, 1)$ of the additional surplus generated when the client deals with both divisions of the bank is captured by the advisory division, and that the rest goes to the execution unit. The respective divisions therefore earn fees $\phi_A(\theta_t)$ and $\phi_E(\theta_t)$, where

$$\begin{aligned} \phi_A(\theta_t) &= \lambda (\phi(\theta_t) - \underline{\phi}(\theta_t)) = \lambda \theta_t p (1 - x); \\ \phi_E(\theta_t) &= \theta_t ((1 - \lambda) (1 - x) p + x) + (1 - \theta_t) \nu_D. \end{aligned} \tag{15}$$

¹⁸We think of the advisory division as encompassing activities that are less susceptible to formal contracting such as the traditional investment banking function of advising clients on capital raising transactions and mergers and acquisitions. We envision the execution division encompassing functions that are more susceptible to formal contract such as capital market operations, institutional and retail brokerage, and proprietary trading and market making.

The execution division captures its information rent and a fraction $(1 - \lambda) > 0$ of the advisory surplus.

We now exhibit an equilibrium of the new game in which the advisory division tells the truth about type and a type S execution division takes action ω in state ω . The equilibrium is sustained by client beliefs $\beta = 0$ whenever $z = x$ after the advisory division announces that $\omega = 1$. Precisely as in the proof of Proposition 3, these beliefs ensure that $\mu = 0$, and so generate a credible punishment phase.

With the above beliefs, the advisory bank will tell the truth. If it were to claim that $\omega = 2$ when the true state was 1 then it would benefit from an enhanced S-type execution division reputation (see equation (15)). But it does not know the execution division's type, and so runs the risk of an immediate loss of revenue if it is D-type. This must be weighed against the inevitability of long-run reputation building (see Lemma 2). A sufficiently patient advisory division will prefer waiting to risking a current lie.

When state 1 is announced, the S-type execution division will receive a standardized payoff close to $(1 - \lambda)(1 - x)p + x$ from truth telling, because its type is almost fully revealed in the long run by Lemma 2. If it takes action 2 then its standardized payoff will be x . Because $\lambda < 1$ the former payoff is dominant for a sufficiently patient execution division.

This argument yields Proposition 4:

PROPOSITION 4. If δ is sufficiently close to 1 then there is an equilibrium in which the advising unit always tells the truth and the execution division selects action ω in state ω .

Proposition 4 demonstrates that the presence of an independent advisory division can resolve the incentive problems that led to inefficient reputation building and punishment phases in the equilibrium of Proposition 3. The advisory unit's desire to maintain a behavioral reputation for always telling the truth about the state of the world keeps it on the straight and narrow, and its information is sufficient to negate the execution wing's incentive to build a type reputation at the expense of its clients. Thus Proposition 4 provides an economic rationale for "Chinese Walls" separating functional units in full-service banks. Of course, for this argument to work it is essential that the advisory division be genuinely independent. Its truth telling incentives derive from its inability to observe the execution division's type. If the execution division could bribe the advisory division to mis-report state 1 then the advisory division could infer its type, and the efficient equilibrium would break down.

The efficient equilibrium of Proposition 4 is not the only one for the divisional bank.

PROPOSITION 5. When θ_0 is low enough there is an inefficient equilibrium in which the S-type execution division chooses action 2 when state 1 is announced. This equilibrium generates a higher surplus for the execution division than that of Proposition 4.

If the advisory division tells the truth the execution division can reveal itself to be S-type with certainty by taking action 2 once after state 1 is announced. After that there is no need for further reputation building, so client beliefs $\beta = 1$ are sustainable thereafter. These beliefs generate very high fees in exchange for short-term lower fees from the anticipated signaling. For sufficiently low θ_0 the former outweigh the latter, so that the execution division prefers the suboptimal equilibrium of Proposition 5.

Since low θ_0 implies weak type reputation, Proposition 5 suggests that the clients of reputable advisory operations could suffer from conflicts of interest if their bank's execution division is less well established or has lost key people with whom the firm's type reputation is associated. As we noted in the introduction, both conditions were met by the late 1990s as a number of investment banks were acquired by commercial banks and both individuals and teams of bankers became more mobile.

7. Technological Shocks

So far we have assumed that the bank's type is fixed. We now relax this assumption by introducing technological shocks to the model that render the bank's knowledge obsolete. In any period, the probability that a shock occurs is $\kappa > 0$; after a shock the bank's type is S with probability θ , irrespective of its type prior to the shock. We assume that shocks are publicly observable.

We allow the bank to make an announcement about its type in each period. Such an announcement would have no effect in the fixed type model of earlier sections, since D-banks would have no incentive to tell the truth. But, in the repeated game setting, the ability to make announcements allows first best to be achieved when δ is close to 1.

PROPOSITION 6. *For any $\kappa < 1$ there exists a $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$ the following constitutes an equilibrium of the game*

1. *The bank truthfully reveals its type after every shock;*
2. *For every t the bank takes the state-appropriate action: $\beta_t = \mu_t = 1$;*
3. *If a bank announces high type after a shock and subsequently a payoff prior to the next shock z_T is 0, $\beta_\tau = \mu_\tau$ for all $\tau > T$.*

This equilibrium achieves the first best. S-banks are paid $\nu_S(1)$ in each period, and D-banks are paid ν_D .

Proof: S-banks clearly have no incentive to mis-represent their type. It is therefore sufficient to prove that D-banks do not wish to lie. The equilibrium has a recursive structure, and we can write the respective equilibrium payoffs of D- and S- banks as follows:

$$\begin{aligned} V_D &= \nu_D + \delta [\kappa (\theta_0 V_S + (1 - \theta_0) V_D) + (1 - \kappa) V_S] \\ V_S &= \nu_S + \delta [\kappa (\theta_0 V_S + (1 - \theta_0) V_D) + (1 - \kappa) V_S]. \end{aligned}$$

It follows immediately that

$$\begin{aligned} V_S &= \frac{\nu_D \delta (1 - \theta_0) \kappa + \nu_S (1 - \delta + \delta \theta_0 \kappa)}{(1 - \delta) (1 - \delta (1 - \kappa))} \\ V_D &= \frac{\nu_D \delta (1 - \theta_0) \kappa + \nu_D (1 - \delta + \delta \theta_0 \kappa)}{(1 - \delta) (1 - \delta (1 - \kappa))} \end{aligned}$$

If a bank is caught in a lie in period t then its per-period payoff is D until the next technology shock, after which it is

$$\nu_P = \nu_D (1 - \theta_0) + \theta_0 \nu_S (0),$$

so that the value V_L of a bank caught lying satisfies

$$V_L = \nu_D + \delta \left[\kappa \frac{\nu_P}{1 - \delta} + (1 - \kappa) V_L \right].$$

This yields the following expression for V_L :

$$V_L = \frac{1}{1 - \delta (1 - \kappa) (\nu_D + \kappa \delta \nu_P)}.$$

If a D-bank deviates by announcing itself as an S-bank after a shock, it will continue lying until the next shock, because truth-telling will result in an immediate punishment for sure. The payoff from deviating V_{dev} is

$$V_{dev} = \nu_S + \delta \{ (1 - p_0) [\kappa (\theta_0 V_S + (1 - \theta_0) V_D) + (1 - \kappa) V_d] + p_0 [\kappa \nu_P + (1 - \kappa) V_L] \},$$

where $p_0 = \mathbb{P}[z = 0] = (1 - p)(1 - q)$. A low type has no incentive to deviate precisely when $V_d \leq V_D$. As $\delta \rightarrow 0$ this expression tends to the following requirement:

$$\frac{p_0 (\nu_P - \nu_D (1 - \theta_0) - \nu_S \theta_0)}{p_0 + \kappa (1 - p_0)},$$

which is less than zero because $\nu_P - \nu_D (1 - \theta_0) - \nu_S \theta_0 = \theta_0 (\nu_S (0) - \nu_S (1)) < 0$. The result is then immediate. \square

The intuitive reason that D-banks do not lie in equilibrium is that a technology lasts $1/\kappa$ periods in expectation. If the bank is caught lying in those periods then β is zero for all subsequent periods, and it loses its behavioral reputation forever. As $\delta \rightarrow 1$ the value of this loss tends to infinity; the gain from lying is a higher fee income, whose expected value is finite, because it has a finite life. The D-bank's tradeoff therefore favors truth telling.

First best is achieved in the equilibrium of Proposition 6 because the bank's repeated type announcements allow it to create a behavioral reputation for telling the truth about its type. Moreover, the fact that type is truthfully announced in equilibrium obviates the bank's need to engage in damaging type-reputation building.

We believe that Proposition 6 would be robust to a set-up in which technological shocks were observed only by the bank. Even in this set-up a bank would be caught lying with significant probability, and would then lose all of the benefits of its behavioral reputation. In fact, even if the low type's lying is never fully revealed, as it is in this model with a 0 payoff, we are able to show that close to first best outcomes can be achieved for high enough δ , because clients are able to make statistical inferences from their payoffs, and so establish when the bank has lied a great deal.

7.1. Shocks that Affect the bank's Probability of Survival

Proposition 6 demonstrates that, when the bank has a high degree of patience, first best can be achieved. In this section we ask what the consequences are of lower patience, when technological shocks could force a bank out of business. Casual empiricism suggests that this is a reasonable assumption; it might reflect obsolescence of an existing technology, or a wave of fresh competitors. We do not attempt to model the reasons for failure, but we assume that both S- and D- banks have a failure probability h whenever a technological shock occurs. Making h higher for D-banks would have no qualitative effect upon our results, because D-banks are not strategic.

We focus upon the case with small κ , high h , and δ close to 1. A large h reduces the value of a reputation for truth telling; Proposition 7 demonstrates that, as a result, it can reverse the efficiency result of Proposition 6.

PROPOSITION 7. *When κ is small, h is large, and δ is close to 1, there is an equilibrium in which each shock results in a reputation-building phase followed by a behavioral reputation phase, as in Proposition 3. There is no equilibrium in which D-banks announce their type truthfully.*

Proof: Similar to parts of the proof of Propositions 3 and 6, and hence omitted. □

The events leading to and including the New York Stock Exchange's 1970 decision to permit its member firms to go public represent an important example of technological change sufficiently profound to give rise to Proposition 6.¹⁹ The 1970s also witnessed the introduction of the hostile takeover which from the outset was viewed as an affront to client relationships.²⁰ From the perspective of our model, it is not surprising that Drexel Burnham

¹⁹See Morrison and Wilhelm (2008).

²⁰Mergerstat Review (1986, p.43) claims that hostile takeovers originated in 1974 when Morgan Stanley represented Nickel Co. in a successful bid for ESB.

Lambert, a second tier bank with perhaps little to lose prior to Michael Milken's development of the junk bond market, was the dominant adviser to bidders during the 1980s as well as being a major source of financing.²¹ Equally noteworthy was Goldman Sachs's unwillingness to represent hostile bidders.²² As the only remaining bulge bracket bank not having gone public by 1987, Goldman may have been less exposed to ongoing change in the industry and thus less inclined to place its behavioral reputation at risk.

8. Conclusion

We have presented a model in which a bank can build a type reputation for its competence, and a behavioral reputation for ethical behavior. Type reputation formation is in conflict with behavioral reputation maintenance and, in our basic set-up, there is no equilibrium in which a bank with low type reputation takes the ethical actions needed to maintain a behavioral reputation. Our analysis reveals a reputational life-cycle effect: young banks with poor type reputations engage in a reputation building phase, during which they take actions to reveal their abilities even when those actions are not in the interests of their clients; only when their type reputation is established do banks settle into an ethical behavior pattern, acting in their clients' best interests even when the short-term effects upon the bank of doing so are unappealing.

We investigate the extent to which institutional design choices can resolve the efficiency problem highlighted above. Our analysis indicates that the first best is achievable when the execution division to which type reputation adheres is supplemented by an independent advisory division that does not know the type of the execution division. In other words, apparently ethical behavior is an equilibrium phenomenon in markets that have independent advisory banks. This is one way to understand the increasing importance of such firms in the investment banking business.

Finally, we demonstrate that banks can establish a reputation for telling the truth about their types when those types change frequently. Provided the banks are long-lived, their truth-telling reputations are sufficient to generate first best outcomes. If they are not then the outcome reduces to a series of back-to-back equilibria of the type we study in our base model.

Our analysis is motivated by an interest in understanding why investment banks, that for so long sought to gain their clients' trust, increasingly appear content to maintain arms-length

²¹Drexel advised 58 bidders. Drexel's closest competitors were Lehman which represented 40 bid-side transactions followed by Morgan Stanley with 30, First Boston with 22, and both Merrill Lynch and Lazard with 20. See Benveniste, Singh and Wilhelm (1993) for details of the rise and fall of Drexel with the development of the junk bond market.

²²Over the entire decade (319 hostile transactions according to SDC), Goldman advised only 4 hostile bidders. On the other hand, Goldman was by far the top adviser to targets with 67 transactions with First Boston its closest competitor with 51 transactions.

relationships.²³ Our analysis suggest that the combination of technological changes and commercial bank entry that have given rise to large, complex, full-service banks aggravated a tension between concerns for type reputation and reputation for ethical behavior. Advisory functions remain less susceptible to formal contract and therefore more dependent on gaining and preserving their clients' trust. But technological advances have increased the capacity for monitoring and measuring performance in the execution of brokerage and risk bearing functions (including the placement of securities with investors) and thereby have made them more susceptible to (arms-length) formal contracting.

We do not contend that being perceived as behaving ethically is of no concern for these functions but rather that it is of less concern in both absolute terms and relative to advisory functions. Our model demonstrates that when a bank's execution functions have well-established reputations for competence the bank is more likely to enjoy a stable reputation for ethical behavior. If it were feasible to maintain strict independence between advisory and execution functions, a bank could even avert the incentive distortion that might otherwise arise with, for example, the introduction of a new operational unit. Moreover, what we might think of as modest but steady technological change can reinforce a delicate balance between type- and behavior- reputation concerns.

Conversely, our analysis also illustrates how profound technological shocks that expose banks to a high risk of failure upset this balance. In this setting, our model predicts instability characterized by destruction of both type- and behavioral- reputation followed by periods of rebuilding type-reputation during which clients expect affected banks to place their interests before those of the client if by doing so they enhance their type-reputation. We suggest that this perspective can shed light on the timing of, for example, the first wave of hostile takeover activity during the 1980s and recent events that have followed commercial bank entry into investment banking.

Finally, to the extent that behavioral-reputation concerns complement or substitute for formal regulation, the current regulatory policy debate might benefit from a clear understanding of their role in governing the behavior of investment banks. Our model sheds light on why reputation concerns vary across banking functions, how they interact with one another, and how they might be expected to evolve in the future.

References

- Abreu, D., D. Pearce and E. Stacchetti (1990), 'Toward a Theory of Discounted Repeated Games with Imperfect Monitoring', *Econometrica*, Vol. 58, No. 5 (September), pp. 1041–1063.

²³In testimony before Congress in relation to the 2007 ABACUS transaction, Lloyd Blankfein, Goldman Sachs' Chairman and CEO, famously remarked with respect to the firm's market making functions that "[T]he markets work on transparency with respect to what the item is. It doesn't carry representations...."

- Asker, J. and A. Ljungqvist (2010), 'Competition and the structure of vertical relationships in capital markets', *Journal of Political Economy*, Vol. 118, No. 3 (June), pp. 599–647.
- Athey, S. and K. Bagwell (2008), 'Collusion with Persistent Cost Shocks', *Econometrica*, Vol. 76, No. 3 (May), pp. 493–540.
- , ——— and C. Sanchirico (2004), 'Collusion and Price Rigidity', *Review of Economic Studies*, Vol. 71, No. 2 (April), pp. 317–349.
- Benveniste, L. M., M. Singh and W. J. Wilhelm, Jr. (1993), 'The Failure of Drexel Burnham Lambert: Evidence on the Implications for Commercial Banks', *Journal of Financial Intermediation*, Vol. 3, No. 1 (October), pp. 104 – 137.
- Bodnaruk, A., M. Massa and A. Simonov (2009), 'Investment Banks as Insiders and the Market for Corporate Control', *Review of Financial Studies*, Vol. 22, No. 12, pp. 4989–5026.
- Bolton, P., X. Freixas and J. Shapiro (2007), 'Conflicts of interest, information provision, and competition in the financial services industry', *Journal of Financial Economics*, Vol. 85, No. 2 (August), pp. 197–330.
- Boot, A. W. A., S. I. Greenbaum and A. V. Thakor (1993), 'Reputation and Discretion in Financial Contracting', *American Economic Review*, Vol. 83, No. 5 (December), pp. 1165 – 1183.
- Carosso, V. P. (1970), *Investment Banking in America: A History*, Harvard University Press, Cambridge, Mass.
- Chemmanur, T. J. and P. Fulghieri (1994a), 'Investment Bank Reputation, Information Production, and Financial Intermediation', *Journal of Finance*, Vol. 49, No. 1 (March), pp. 57 – 79.
- Chemmanur, T. J. and P. Fulghieri (1994b), 'Reputation, Renegotiation, and the Choice Between Bank Loans and Publicly Traded Debt', *Review of Financial Studies*, Vol. 7, pp. 475 – 506.
- Chen, Z., A. D. Morrison and W. J. Wilhelm (2011), 'Fiduciary Responsibility, "Star" Cultures, and Reputation', Mimeo, University of Oxford, Oxford, U.K.
- Cripps, M. W., G. J. Mailath and L. Samuelson (2004), 'Imperfect Monitoring and Impermanent Reputations', *Econometrica*, Vol. 72, No. 2 (March), pp. 407–432.
- Diamond, D. W. (1989), 'Reputation acquisition in debt markets', *Journal of Political Economy*, Vol. 97, No. 4 (August), pp. 828–862.
- Ellis, C. D. (2009), *The Making of Goldman Sachs*, Penguin, London.
- Ely, J. C. and J. Välimäki (2003), 'Bad Reputation', *Quarterly Journal of Economics*, Vol. 118, No. 3 (August), pp. 785–814.

- Ely, J., D. Fudenberg and D. Levine (2008), ‘When is reputation bad?’, *Games and Economic Behavior*, Vol. 63, No. 2 (July), pp. 498–526.
- Fudenberg, D. and D. K. Levine (1992), ‘Maintaining a reputation when strategies are imperfectly observed’, *Review of Economic Studies*, Vol. 59, No. 3 (July), pp. 561–579.
- , D. Levine and E. Maskin (1994), ‘The Folk Theorem with Imperfect Public Information’, *Econometrica*, Vol. 62, No. 5 (September), pp. 997–1039.
- and Y. Yamamoto (2010), ‘Repeated Games Where the Payoffs and Monitoring Structure are Unknown’, *Econometrica*, Vol. 78, No. 5 (September), pp. 1673–1710.
- Fulghieri, P., G. Strobl and H. Xia (2010), ‘The Economics of Solicited and Unsolicited Credit Ratings’, Working paper, Kenan-Flager Business School, University of North Carolina, NC.
- Glode, V., R. C. Green and R. Lowery (2012), ‘Financial Expertise as an Arms Race’, *Journal of Finance*, Vol. forthcoming.
- Griffin, J. M., T. Shu and S. Topaloglu (2012), ‘Examining the Dark Side of Financial Markets: Do Institutions Trade on Information from Investment Bank Connections?’, *Review of Financial Studies*, Vol. 25, No. 7, pp. 2155–2188.
- Hartman-Glaser, B. (2012), ‘Reputation and Signaling’, Working paper, Fuqua School of Business, Duke University, NC.
- Kreps, D. M. (1990), ‘Corporate Culture and Economic Theory’, in J. E. Alt and K. A. Shepsle, eds, *Perspectives on Positive Political Economy*, Cambridge University Press, Cambridge, MA, pp. 90 – 143.
- Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982), ‘Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma’, *Journal of Economic Theory*, Vol. 27, pp. 245–52.
- Levin, J. (2003), ‘Relational Incentive Contracts’, *American Economic Review*, Vol. 93, No. 3 (June), pp. 835–857.
- Morris, S. (2001), ‘Political Correctness’, *Journal of Political Economy*, Vol. 109, No. 2 (April), pp. 231–265.
- Morrison, A. D. and W. J. Wilhelm, Jr (2004), ‘Partnership Firms, Reputation and Human Capital’, *American Economic Review*, Vol. 94, No. 5 (December), pp. 1682 – 1692.
- and W. J. Wilhelm, Jr. (2007), *Investment Banking: Institutions, Politics and Law*, Oxford University Press, Oxford, UK.
- and ——— (2008), ‘The Demise of Investment Banking Partnerships: Theory and Evidence’, *Journal of Finance*, Vol. 63, No. 1 (February), pp. 311–350.

- Tadelis, S. (1999), ‘What’s in a Name? Reputation as a Tradeable Asset’, *American Economic Review*, Vol. 89, No. 3 (June), pp. 548 – 563.
- (2002), ‘The Market for Reputations as an Incentive Mechanism’, *Journal of Political Economy*, Vol. 110, No. 4 (August), pp. 854 – 882.
- Tirole, J. (1996), ‘A Theory of Collective Reputations (with applications to the persistence of corruption and to firm quality)’, *Review of Economics Studies*, Vol. 63, No. 1 (January), pp. 1–22.
- Winton, A. and V. Yerramilli (2011), ‘Lender Moral Hazard and Reputation in Originate-to-Distribute Markets’, Working paper, Carlson School of Management, University of Minnesota.

Proofs

Proof of Lemma 2

Lemma 2 is a result with general implications. It is easiest to prove a more general version of the result, and then to demonstrate that Lemma 2 is a special case of that result.

LEMMA 4. *Let (Ω, g) be a Borel measurable space. For all $\epsilon, \psi > 0$ and $\theta^+ \in (0, 1]$, there exists a positive integer K such that for all $\theta_0 \in [\theta^+, 1]$, for any measures P, \hat{P} , and \tilde{P} on (Ω, g) with $P = \theta_0 \hat{P} + (1 - \theta_0) \tilde{P}$, and for every filtration $\{g^t\}_{t \geq 0}$, $g^t \subset g$, we have*

$$\hat{P}[\#\{t > 0 : \theta_t < 1 - \psi\} \geq K] < \epsilon,$$

where $\#\{\cdot\}$ denotes the number of elements in the set $\{\cdot\}$, whenever

$$\sum_{\varpi_{t+1} \in g^{t+1}} |\hat{P}^t(\varpi_{t+1}) - \tilde{P}^t(\varpi_{t+1})| \geq \kappa > 0, \forall t. \quad (16)$$

Ω is the space of the outcomes of the game including all the private and public information. Condition (16) says that the probability measures \hat{P}^t and \tilde{P}^t on the end-of-period outcomes $\varpi_{t+1} \in g^{t+1}$ are sufficiently different. ϖ_{t+1} is one realization of all the public information up to time $t + 1$ in our game.²⁴ \hat{P} is the probability conditional on a type S, \tilde{P} is the probability conditional on a type D in an equilibrium. \hat{P}^t (\tilde{P}^t) is the probability induced by a type S (type D) in period t . P is the probability belief of the clients, which is a mixture between \hat{P} and \tilde{P} with prior θ_0 .

²⁴Notice that ϖ_{t+1} need not to be discrete. If ϖ_{t+1} is continuous, (16) is $\int_{\Omega_{t+1}} |\hat{P}^t(\varpi_{t+1}) - \tilde{P}^t(\varpi_{t+1})| d\varpi_{t+1}$, here $\hat{P}^t(\varpi_{t+1})$ and $\tilde{P}^t(\varpi_{t+1})$ is the probability density functions. Technically, the left hand side of (16) is the total variation of $\hat{P}^t - \tilde{P}^t$. Our proof goes through for continuous ϖ_{t+1} by changing summation to the corresponding integration.

Lemma 2 follows from Lemma 4 if condition (16) is satisfied. To see that this is the case, note that

$$\begin{aligned} \sum_{\varpi_{t+1} \in \mathcal{G}^{t+1}} |\widehat{P}^t(\varpi_{t+1}) - \widetilde{P}^t(\varpi_{t+1})| &\geq |\widehat{P}^t(z_t = x) - \widetilde{P}^t(z_t = x)| + |\widehat{P}^t(z_t = 0) - \widetilde{P}^t(z_t = 0)| \\ &= 2(1-p)(1-q). \end{aligned}$$

Condition (16) is satisfied with $\kappa = 2(1-p)(1-q)$.

Note that Lemma 4 requires that K be independent of the particular measures and probability spaces; it follows that Lemma 4 holds uniformly across all equilibria.

The proof is similar to the proof of Lemma 2 in Sorin (1999). We proceed by proving a series of lemmas. First, notice that from the clients perspective, i.e., under measure P , the reputation of a bank is a martingale. That is, $\theta_t = E_t[\theta_{t+1}]$, which follows from the law of reiterated expectation. The following lemma bounds the one step ahead expected martingale difference.

LEMMA 5. *Suppose X^t is an uniformly bounded martingale under P with $0 < X^t < 1 \forall t$. For all $\eta > 0$ and $K \geq 1$,*

$$P(\#\{t \geq 0 : E_t|X^{t+1} - X^t| \geq \eta\} \geq K) \leq \frac{1}{K\eta^2}. \quad (17)$$

Proof: Fix $m > 0$. Under P

$$E \left[\sum_{t=0}^m (X^{t+1} - X^t) \right]^2 = E \sum_{t=1}^m (X^{t+1} - X^t)^2$$

because the expectation of the cross product can be written

$$\begin{aligned} E[(X^{t+1} - X^t)(X^{t+s} - X^{t+s-1})] &= E[(X^{t+1} - X^t)E_{t+1}(X^{t+s} - X^{t+s-1})] \\ &= 0, \end{aligned}$$

where the final line follows because $E_{t+1}(X^{t+s} - X^{t+s-1}) = 0$ for $s > 1$. But

$$1 \geq E[X^{m+1} - X^0]^2 = E \left[\sum_{t=0}^m (X^{t+1} - X^t) \right]^2 = E \sum_{t=0}^m (X^{t+1} - X^t)^2$$

As $m \rightarrow \infty$. Then we have

$$\begin{aligned} 1 &\geq E \sum_{t=0}^{\infty} (X^{t+1} - X^t)^2 \\ &= E \sum_{t=0}^{\infty} E_t (X^{t+1} - X^t)^2 \\ &\geq P(\#\{t \geq 0 : E_t (X^{t+1} - X^t)^2 \geq \eta^2\} \geq K) \eta^2 K \\ &\geq P(\#\{t \geq 0 : E_t |X^{t+1} - X^t| \geq \eta\} \geq K) \eta^2 K \end{aligned}$$

which is (17). The last inequality follows because $E_t|X^{t+1} - X^t| \geq \eta$ implies $E_t(X^{t+1} - X^t)^2 \geq \eta^2$ by Jensen's inequality and therefore

$$P(\#\{t \geq 0 : E_t|X^{t+1} - X^t| \geq \eta\} \geq K) \leq P(\#\{t \geq 0 : E_t(X^{t+1} - X^t)^2 \geq \eta\} \geq K).$$

□

The next lemma bounds $(1 - \theta_t)$ using condition (16).

LEMMA 6. *Under P , $(1 - \theta_t) \leq \frac{E_t|\theta_{t+1} - \theta_t|}{\theta_t \kappa}$.*

Proof: For any outcome $\varpi_{t+1} \in g^{t+1}$

$$\begin{aligned} E_t|\theta_{t+1} - \theta_t| &= E_t \left| \frac{\theta_t \widehat{P}^t(\varpi_{t+1})}{P^t(\varpi_{t+1})} - \theta_t \right| \\ &= \theta_t E_t \left[\frac{|\widehat{P}^t(\varpi_{t+1}) - P^t(\varpi_{t+1})|}{P^{t+1}(\varpi_{t+1})} \right] \\ &= \theta_t E_t \left[\frac{|\widehat{P}^t(\varpi_{t+1}) - \theta_t \widehat{P}^t(\varpi_{t+1}) - (1 - \theta_t) \widetilde{P}^t(\varpi_{t+1})|}{P^{t+1}(\varpi_{t+1})} \right] \\ &= (1 - \theta_t) \theta_t E_t \left[\frac{|\widehat{P}^t(\varpi_{t+1}) - \widetilde{P}^t(\varpi_{t+1})|}{P^t(\varpi_{t+1})} \right] \\ &= (1 - \theta_t) \theta_t \sum_{\varpi_{t+1} \in g^{t+1}} \frac{|\widehat{P}^t(\varpi_{t+1}) - \widetilde{P}^t(\varpi_{t+1})|}{P^t(\varpi_{t+1})} P^t(\varpi_{t+1}) \\ &= (1 - \theta_t) \theta_t \sum_{\varpi_{t+1} \in g^{t+1}} |\widehat{P}^t(\varpi_{t+1}) - \widetilde{P}^t(\varpi_{t+1})| \\ &\geq (1 - \theta_t) \theta_t \kappa \end{aligned}$$

□

These two lemmas show that $(1 - \theta_t)$ cannot be too large for too many periods if θ_t is not too small. The next lemma show that the case where θ_t is small cannot have a large measure under \widehat{P} .

LEMMA 7. $\widehat{P}[\theta_t \leq l\theta_0, \forall t] \leq lP[\theta_t \leq \gamma\theta_0, \forall t]$ for $l > 0$.

Proof: By Bayes rule, we have

$$\theta_t(\varpi_t) = \frac{\theta_0 \widehat{P}(\varpi_t)}{P(\varpi_t)}$$

If ϖ_t is such that $\theta_t(\varpi_t) \leq l\theta_0$,

$$\begin{aligned} \frac{\theta_0 \widehat{P}(\varpi_t)}{P(\varpi_t)} &\leq l\theta_0, \text{ or} \\ \widehat{P}(\varpi_t) &\leq lP(\varpi_t). \end{aligned}$$

Integrating over $L = \{\varpi_t | \theta_t(\varpi_t) \leq \gamma\theta_0, \exists t\}$ yields the desired result. \square

Notice that Lemma 7 implies that $\widehat{P}[\theta_t \leq l\theta_0, \forall t] \leq l$.

Finally, we can prove Lemma 4. By Lemma 6, we have

$$\begin{aligned} \widehat{P}(\#\{t \geq 0 : \theta_t < 1 - \psi\} \geq K) &\leq \widehat{P}(\#\{t \geq 0 : \frac{1}{\kappa\theta_t} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) \\ &\leq \widehat{P}(\#\{t \geq 0 : \frac{1}{\kappa\theta_t} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) \cap \{\theta_t > l\theta_0, \forall t\} \\ &\quad + \widehat{P}(\#\{t \geq 0 : \frac{1}{\kappa\theta_t} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) \cap \{\theta_t < l\theta_0, \forall t\} \\ &\leq \widehat{P}(\#\{t \geq 0 : \frac{1}{\kappa l\theta_0} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) + \widehat{P}[\theta_t \leq l\theta_0, \forall t] \end{aligned}$$

Because $P \geq \theta_0 \widehat{P}$,

$$\begin{aligned} \widehat{P}(\#\{t \geq 0 : \frac{1}{\kappa l\theta_0} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) &\leq \frac{1}{\theta_0} P(\#\{t \geq 0 : \frac{1}{\kappa l\theta_0} E_t |\theta_{t+1} - \theta_t| > \psi\} \geq K) \\ &\leq \frac{1}{\theta_0} \frac{1}{K(\kappa l\theta_0 \psi)^2}. \end{aligned}$$

The second inequality follows from Lemma 5. Therefore,

$$\widehat{P}(\#\{t \geq 0 : \theta_t < 1 - \psi\} \geq K) \leq \frac{1}{\theta_0} \frac{1}{K(\kappa l\theta_0 \psi)^2} + l, \quad (18)$$

notice $\widehat{P}[\theta_t \leq l\theta_0, \forall t] \leq l$ by Lemma 7.

Choose $\gamma = \frac{1}{2}\epsilon$. the RHS of (18) is $\frac{1}{K\psi^2\theta_0^3(\frac{1}{2}\epsilon)^2\kappa^2} + \frac{1}{2}\epsilon$. Pick K large enough so that $\frac{1}{K\psi^2\theta_0^3(\frac{1}{2}\epsilon)^2\kappa^2} \leq \frac{1}{2}\epsilon$, that is

$$K > K^* = \frac{1}{\psi^2\theta_0^3(\frac{1}{2}\epsilon)^3\kappa^2},$$

we have $\widehat{P}(\#\{t \geq 0 : \theta_t < 1 - \psi\} \geq K) \leq \epsilon$.

Remark: K^* only depends on ψ, ϵ, κ , and θ_0 . K^* thus applies for any game or filtrations satisfying (16) and $\theta_0 \in [\theta^+, 1]$.

Proof of Lemma 3.

Let $\epsilon = \gamma/2$ and $\psi = 1 - \bar{\theta}$. By Lemma 2 there exists an integer M such that \mathbb{P} [There are more than K periods with reputation $< \bar{\theta}$] $\leq \epsilon$. Then

$$\begin{aligned} v_t &\geq \epsilon \mathbb{E} [v | \text{More than } K \text{ periods with reputation } < \bar{\theta}] \\ &\quad + (1 - \epsilon) \mathbb{E} [v | \text{Fewer than } K \text{ periods with reputation } > \bar{\theta}] \\ &\geq (1 - \epsilon) \underbrace{\mathbb{E} [v | \text{Fewer than } K \text{ periods with reputation } > \bar{\theta}]}_A. \end{aligned}$$

Now note that

$$\begin{aligned}
 A &= \underbrace{\mathbb{E} [v | \text{Every period for which } \theta < \bar{\theta} \text{ } \theta \text{ is replaced } \bar{\theta}]}_{v[B]} \\
 &\quad - \text{Avg. additional equivalent income from replacing } \theta \text{ by } \bar{\theta} \text{ in those periods} \\
 &\geq v[B] - (1 - \delta) \left(\sum \text{DFs in periods where replacement occurred} \right) \\
 &\geq v[B] - (1 - \delta^K),
 \end{aligned}$$

where the final line is the sum of the discount factors if the replacement occurs in the first K periods.

$v[B]$ is greater than or equal to the value that would be obtained if θ in each of the replacement periods was $\bar{\theta}$ and β was equal to 0. That value is equal to $(1 - \bar{\theta}) \nu_D + \bar{\theta}x$. Hence,

$$\begin{aligned}
 v_t &\geq (1 - \epsilon) \left[((1 - \bar{\theta}) D) - (1 - \delta^K) \right] \\
 &= (1 - \bar{\theta}) \nu_D + \bar{\theta}x - \epsilon \left[(1 - \bar{\theta}) \nu_D + \bar{\theta}x \right] - (1 - \epsilon) (1 - \delta^K) \\
 &\geq (1 - \bar{\theta}) \nu_D + \bar{\theta}x - \epsilon - (1 - \epsilon) (1 - \delta^K) \\
 &\geq (1 - \bar{\theta}) \nu_D + \bar{\theta}x - \gamma/2 - (1 - \delta^K)
 \end{aligned}$$

The result followed immediately for any δ sufficiently large to ensure that $1 - \kappa^K < \gamma/2$.