# Financial Innovation with Endogenous Experimentation


Zhaohui Chen*
McIntire School of Commerce,
University of Virginia.

Alan Morrison†
Said Business School,
University of Oxford.

William Wilhelm‡
McIntire School of Commerce,
University of Virginia.


July, 2020


*Email: zc8j@comm.virginia.edu.
†Email: alan.morrison@sbs.ox.ac.uk.
‡Email: wjw9a@comm.virginia.edu.

# Financial Innovation with Endogenous Experimentation


### Abstract

We present a model in which a long-lived bank endogenously learns about the environment for financial innovation through experimentation on its clients. When the bank has superior knowledge of the state of the world facing its clients, it may engage in inefficient or "reckless" experimentation because it cannot commit to do otherwise. We show that strong client relationships can mitigate the incentive to pitch innovative products to clients for whom they are not appropriate. We also explore the limits of internal monitoring systems intended to prevent such behavior. Finally, we show that greater banker mobility can complement monitoring systems because bankers will have incentive to visibly seek out weak monitoring environments only when the expected collective benefits from financial innovation are especially large.

## 1. Introduction

The 2008 financial crisis was preceded by 4 decades of financial innovation, the pace of which has been described as "revolutionary" and "explosive".[1] Over the same period, client relationships maintained by the investment banks responsible for much of the innovative activity became increasingly antagonistic and less exclusive.[2] In this paper, we develop a model of endogenous learning about the environment for financial innovation that points to a connection between these observations. The model also suggests explanations for the changing institutional structure around financial innovation.

Our model is premised on the idea that most financial innovation is incremental and dependent on experimentation in the marketplace.[3] For example, the Black-Scholes-Merton option-pricing framework was a transformative advance in financial technology that continues to spawn incremental advances along the path that it opened. Although the transformative advance might be thought of as having been discovered through "laboratory" experimentation, the subsequent flood of incremental advances have largely resulted from transactional experimentation with the new technology.[4] In our model, a successful experiment creates value by widening the range of new products and services that a bank can confidently deliver to future clients.

We model a long-lived bank's experimental strategy as a multi-armed bandit problem.[5] Short-lived clients arrive sequentially and the bank determines whether it will serve the client through a safe technology or a flexible technology. The safe technology provides any client type with a fixed certain payoff. The flexible technology provides an opportunity for experimentation and offers an uncertain but higher expected payoff to clients *for whom it is appropriate.* The bank learns nothing about the state of its flexible technology when it deploys the safe technology. However, with each successful (or unsuccessful) experiment with the flexible technology, the bank updates it beliefs about the potential for further technological advances. The needs of any given client may or may not correspond with

[1] For example, see Miller (1986), Ross (1989) and Scholes (1998).

[2] See Morrison, Schenone, Thegeya, and Wilhelm (2018) on declining and relationship exclusivity and Chen, Morrison, and Wilhelm (2015) for a related model of tension between banks and their clients.

[3] Arrow (1962) provides an important early example of a "learning-by-doing" model.

[4] Although we do not model transformative advances we provide insight into how a more (or less) transformative advance influences the incremental innovation process.

[5] For example, see Keller, Rady, and Cripps (2005).

the present state of the bank's flexible technology. As a consequence, not every client will be an ideal candidate for experimentation.

We begin our analysis by assuming that both the client and the bank know the client's type. Even in this setting, there is tension between the interests of clients on whom the bank experiments and future clients who benefit from experiments that were successful. Even if the current client is best served with the safe technology, the bank may choose to experiment with the flexible technology because the expected benefit to future clients outweighs the cost of experimentation on the current client. Because clients understand this conflict and the bank captures all of the surplus in our model, the bank collects lower fees in the short-run in favor of discovering a path to greater fees in the future. In this equilibrium, ("costly") experimentation that is not in the best interest of the current client only occurs when the environment for financial innovation is especially promising.

We then introduce moral hazard by assuming that the bank has private information about the client's type. This assumption follows work by Bolton, Freixas, and Shapiro (2007) and Chen et al. (2015) that assumes that banks have superior knowledge of the state of the world facing their clients. In this case, we show that the bank will experiment more aggressively in the sense that it engages in costly experimentation regardless of how promising is the environment for financial innovation. The bank takes such inefficient actions because it cannot commit to do otherwise. This raises the cost of experimentation by reducing current fee expectations and reduces the present value of the bank's future revenue stream by causing the bank to stop experimentation at a lower stage of technological development than in the benchmark case. Indiscriminate experimentation of this sort is consistent with frequently voiced concerns for reckless pursuit and misuse of complex financial products.[6] We show that such behavior cannot be curbed by continuation equilibria in which beliefs about reckless behavior serve as a punishment mechanism. It is impossible to distinguish between efficient and reckless experimentation and punishing suspicious experimentation is too costly to sustain a continuation equilibrium.

Given that it is in the bank's interest to minimize the negative consequences of financial innovation, we explore potential institutional responses to the moral hazard problem.

[6]See Allen (2012) for a review of the "evidence for the dark side of financial innovation."

Through at least the 1960s, banks and their clients dealt repeatedly, often exclusively, with one another over the course of many years.[7] This historical feature of capital markets suggests that self-enforcing relational contracts may have once played a role in governing financial innovation.

We consider this possibility by assuming that the bank deals repeatedly with a long-lived client so that both parties now have a stake in the long-term benefits of successful innovation. We show that a relational contract under which the client pays a fee equal to its maximum possible payoff from the transaction coupled with a rebate reflecting the actual payoff yields an optimal experimental strategy that is identical to the strategy obtained in the absence of moral hazard. Importantly, the bank's commitment to this agreement is sustained by the threat of suboptimal equilibrium beliefs were the bank to renege on the agreement. Because there is no need for punishment along the optimal equilibrium path, the present value of the relational contract is equal in absolute value to the decline in present value associated with the introduction of moral hazard into the benchmark model.

Given that bank-client relationships weakened considerably after 1970, we also consider whether a monitoring mechanism could credibly curb indiscriminate experimentation.[8] There are several barriers to success here. First, although advances in computing and information technology widened the scope for monitoring (e.g., using VaR), this was occurring alongside rapid increase in the scale and scope of bank operations. Second, there are simply limits to the effectiveness of any monitoring system. Third, for a monitoring system to be successful, it must be both transparent and/or structured in such a way as to substantially limit the monitor's ability to share in the benefits from experimentation. Finally, as banks shifted from private to public ownership, human capital became more mobile.[9]

We begin by assuming that the monitor can prevent experimentation with probability less than one but that it cannot distinguish indiscriminate experimentation from experimentation that is not harmful to the client at hand. This imperfect monitoring

[7] See Eccles and Crane (1988) and Morrison et al. (2018).
[8] Morrison et al. (2018) figure 1 and surrounding text.
[9] Morrison et al. (2018) figure 6 and surrounding text.

mechanism mitigates the moral hazard problem but, obviously, does not support the first-best experimental strategy. More importantly, if we assume that the capacity for financial innovation lies in individual bankers, an imprecise monitoring mechanism creates incentive for bankers to "shop" for more or less aggressive monitoring to maximize the value of their human capital. We show that because the decision to switch banks is both visible and consistent with pursuit of the first-best experimental strategy, banker mobility complements monitoring. The model predicts that a promising environment for innovation will drive bankers toward banks with weaker monitoring systems and vice versa.

## 2. Related Literature

Our model stems from the multi-armed bandit literature which studies statistical decision models where in each period an agent (the bandit) draws from a set of distinct stochastic payoff distributions (the arms) with the goal of maximizing the present value of his payoffs. Experimentation across arms and over time yields information about individual payoff distributions.[10] Within this literature, our analysis is most directly related to work on optimal experimentation in a principal-agent framework. For example, Manso (2011) demonstrates that agents charged with conducting privately costly experimentation should not be punished, and perhaps should even be rewarded, for early experimental failures.[11] We identify potential harm to uninformed agents who are subjects of experimentation.

Our baseline model highlights Fudenberg and Levine's (1994) analysis of the agency problem in a setting where a long-lived agent deals repeatedly with a short-lived principal. Our analysis of the relational contract when both clients and the banks are long-lived follows Levin (2003) but we extend the analysis by studying how bank monitoring and

[10]See Bergemann and Valimaki (2008) for a review of the bandit literature. Important examples include Bolton and Harris (1999) and Keller, Rady, and Cripps (2005) where multiple agents simultaneously conduct costly experimentation. Agents' ability to observe the experience of their peers produces an informational externality that can lead to over- or under-experimentation in equilibrium.

[11]Bergemann and Hege (1998, 2005), Horner and Samuelson (2014), Halac, Kartik, and Liu (2016) also study the principal-agent contracting problem in a bandit setting. Halac, Kartik, and Liu (2017) study the optimal payoff and information disclosure policy when multiple agents compete for innovations through experimental learning.

banker mobility can further diminish the agency problem.[12]

Among theoretical studies of financial innovation, our analysis is most closely related to that of Persons and Warther (1997) in the sense that it involves endogenous learning that can either extend or shut down innovative efforts. But rather than assuming that clients adopt an innovation and thereby enable others to learn from their experience, we assume that banks (and future clients) learn from experimentation on clients who may or may not be ideal candidates for the innovative product or service. In contrast to their conclusion that there is too little adoption of financial innovations from a social perspective (p.943), we argue that an innovative bank may be "socially destructive" because it is unable to commit to be otherwise. However, our model also suggests means by which banks could diminish the commitment problem and the resulting negative consequences of financial innovation.

## 3. The Model

We study the interaction between an infinitely-lived bank and short-lived clients. Both players are risk-neutral and each has a per-period discount factor of $\delta$. During each time period $t$, a single client enters the market for one period with a type index $n = 0, 1, 2, ...$, where the distribution of $n$ is *i.i.d.* across periods with probability $p_n > 0$ and $\sum_n p_n = 1$ The type index can be thought of as the state of nature in which the bank's service will be rendered or as a measure of the nature or complexity of the client's requirements. In our analysis of the model we consider both the possibility that $n$ is public information (and therefore contractible) and that it is observable only by the bank. In every case, the distributional characteristics of $n$ are common knowledge.

During each period, the bank enters a contract with the client under which it receives a fee, $\phi_t$, for providing service to the client and then observes $n$. The bank's services can be delivered using either a *safe technology* that offers a fixed payoff, $M$, for any client type $n$, or a *flexible technology* that delivers a state-contingent payoff and can serve as a platform for experimentation. We assume throughout the paper that the client has no outside opportunity so that $\phi_t$ is equal to the expected surplus from the transaction.

[12]See MacLeod (2007) for a review of the relational contracting literature.

At the bank's current stage of technological development, indexed by $i = 0, 1, 2, ...,$ the flexible technology provides client types $n \leq i$ a payoff of 1 with probability $\rho > 0$ or 0 with probability $1 - \rho$ where $\rho > M > 0$. For client types $n > i$, the flexible technology provides a payoff of 0. This payoff structure implies that the bank's flexible technology benefits only those clients for whom the complexity of their requirements does not exceed the bank's current capabilities.

The flexible technology also provides for experimentation that can lead to incremental improvement in the technology. Specifically, at any time, $t$, and any stage of technological development, $i$, the bank can experiment with technology $i+1$ regardless of the client type at hand. At the beginning of experimentation for any stage of development, $i$, the bank experiments on technology $i+1$ conditional on the prior probability of technological advance being feasible, $\pi_0^{i+1}$, and the expected payoff from experimentation, $\pi_0^{i+1}\rho < \rho$. For simplicity, we assume that $\pi_0^{i+1} = \pi_0$ for all $i$. A failed experiment at stage $i$ yields a posterior probability $\pi_\tau^{i+1}$ where $\tau$ indicates the number of failed attempts to advance the technology from $i$ to $i+1$. Thus $\tau$ can be interpreted as the bank's tolerance for further experimentation. With each successive failure, expectations for the feasibility of technological advance diminish.

With probability $\pi_\tau^{i+1}\rho$, the experiment is successful, yielding a payoff of 1 for client types $n \leq i+1$ and the bank's stage of technological development advances to stage $i+1$. With complementary probability the experiment fails and yields a payoff of 0. A 0 payoff indicates either that technological advance is feasible but was unrealized or that the flexible technology cannot be advanced beyond its present stage. Advancing the flexible technology increases the range of client types over which it can be effectively deployed. This increases the expected present value of future payoffs from the flexible technology but we assume that it does so with decreasing returns. Specifically, we assume $p_n < p_{n-1}$ and that the state probabilities follow a geometric sequence for which $\frac{p_n}{p_{n-1}} = g < 1$ and $\sum_n p_n = \frac{p_0}{1-g} = 1$

As just described, there is learning about the feasibility of success through repeated experimentation *at a given stage of technological development,* but this information has no bearing on beliefs about the feasibility of technological advance at higher stages of technological development. Beliefs are reset to $\pi_0^{i+1}$ at the beginning of subsequent rounds of experimentation. The bank's stage of technological development, the payoff structure, including $M$, and realized payoffs are public information as are the prior and posterior probabilities of failure and the number of failed experiments.

We treat the bank as facing a multi-armed bandit problem in which there is a tradeoff between achieving the best immediate outcome for its client and engaging in experimentation that may open a path to greater long-term gains. Benefits from the latter are potentially large. As a consequence, the bank may be willing to experiment with the flexible technology even if the expected payoff to the current client is less than the fixed payoff from the safe technology. Unless stated otherwise, we focus on this case by assuming

$$\pi_0 \rho < M. \tag{1}$$

We define the bank's strategy as a mapping

$$\sigma : H_t \times N \longmapsto A_t,$$

where $\sigma_t$ is a probability distribution over the set of technologies $A_t \equiv \{0, i, i+1\}$ on which the bank can draw to serve the client at time $t$, with 0 denoting the safe technology; $H_t$ is the set of all possible histories of publicly observable client types, bank actions, and payoffs leading to time $t$; and $N$ is the set of time $t$ client types ($n = 0, 1, 2, ...,$).

We restrict the history of information, $h_t$, in the bank's time $t$ strategy, $\sigma_t = \sigma(h_t, n)$, to the current stage of technological development, $i$, and the number of experimental failures, $\tau$, at that stage of development. This assures that $i$ and $\tau$ uniquely determine the probability, $\pi_\tau^{i+1}\rho$, of a successful experiment. Finally, we assume that once the bank ceases experimentation, it will not restart. This restricts the bank's strategy to be an optimal stopping time.[13] These two assumptions simplify the exposition to a considerable degree and in section 8.1 of Appendix 2 we demonstrate that they do so without loss of generality.

[13]Intuitively, if after $\tau$ failed experiments, $\pi_\tau^{i+1}$ is close to zero, the bank will optimally stop experimentation. Once it does so, the bank gains no new information regarding the probability of successful experimentation and therefore has no incentive to restart experimentation.

## 4. Experimentation under Complete Contracting

We begin the analysis of the model by studying the case in which the client's type is public information and therefore contractible. This setting reveals the conflict of interest at the core of our analysis: even with complete contracting and no informational friction, the bank may experiment with innovative products in states where the client would be best served by the bank's existing technology.

### 4.1 *Single-Stage Innovation*

We build intuition for the basic forces bearing on the decision to innovate by first considering the case where client type, $n$, is either 0 or 1 and the stage of the bank's flexible technology can be advanced by, at most, one increment from an initial value of $i = 0$ to $i = 1$. For the moment, we assume and prove in Appendix 1 that the bank will never find it optimal to experiment with type 0 clients. This assumption enables us to focus our attention on type 1 clients alone because the expected payoff from type 0 clients is a constant, $\rho$.

The bank's experimental strategy is restricted by an optimal stopping time at which the posterior probability, $\pi^1_\tau$, of successful experimentation following $\tau$ failed experiments at the current stage of development is less than a critical value $\underline{\pi}^1 > 0$. If the bank ceases experimentation prior to advancing the flexible technology to stage 1, it will serve all subsequent clients with the safe technology.

We begin characterization of $\underline{\pi}^1$ by defining $V^1$ as the bank's payoff from serving a type 1 client. If the bank stops experimentation when $\pi^1_\tau = \underline{\pi}^1$, its payoff from the current and all future type 1 clients is $V^1 = \frac{p_1 M}{1-\delta}$. Alternatively, if the bank continues to experiment it receives a current fee of $\pi^1_\tau \rho$ and, if the experiment is a success, a payoff of $V^1 = \frac{p_1 \rho}{1-\delta}$ from all future type 1 clients. A failed experiment yields the posterior probability of successful innovation

$$\pi^1_{\tau+1} = \frac{\pi^1_\tau(1-\rho)}{\pi^1_\tau(1-\rho) + (1-\pi^1_\tau)}, \tag{2}$$

which is less than $\underline{\pi}^1$. Therefore, the bank optimally stops experimenting and the payoff from future type 1 clients is $V^1 = \frac{p_1 M}{1-\delta}$. Faced with these two options, the bank experiments with the flexible technology if and only if

$$\pi_\tau^1 p_1 \rho + \delta \left[ \pi_\tau^1 \rho \frac{p_1 \rho}{1-\delta} + (1 - \pi_\tau^1 \rho)) \frac{p_1 M}{1-\delta} \right] \geq \frac{p_1 M}{1-\delta} \tag{3}$$

or

$$(\pi_\tau^1 \rho - M) + \delta \pi_\tau^1 \rho \frac{\rho - M}{1-\delta} \geq 0.$$

By Assumption (1), the first term reflects the short-term income loss from a failed experiment while the second term reflects the long-term gain from a successful experiment. Thus, $\underline{\pi}^1$ is defined by the level of $\pi_\tau^1$ at which the bank is indifferent between experimenting or not:

$$\underline{\pi}^1 = \frac{M}{(1 + (\rho - M)\frac{\delta}{1-\delta})\rho}.$$

The following proposition characterizes the bank's optimal strategy:[14]

**Proposition 1.** *If the client's type is public information and contractible and the bank's flexible technology can be advanced by, at most, one stage, the bank experiments only on type 1 clients and does so if and only if $\pi_\tau^1 \geq \underline{\pi}^1$. The bank's incentive to experiment is increasing in $\delta$ and $\rho$ and decreasing in $M$.*

In other words, experimentation is more attractive when the bank is patient, the potential gain from technological advance is larger, and the fixed payoff from the safe technology is smaller.

### 4.2 *Unbounded Innovation Capacity*

We now consider the case where the potential for advancing the bank's flexible technology is unbounded. The bank's time $t$ expected present value given a strategy, $\sigma$, is denoted as

$$V(h_t, \sigma) = E_t \left[ \sum_{k=0}^{\infty} \delta^k \phi_{t+k}(\sigma) \right].$$

[14] Unless otherwise stated, all proofs are provided in Appendix 1)

The optimal bank value is

$$V(h_t) = \max_{\sigma} V(h_t, \sigma).$$

In equilibrium

$$\phi_{t+k} = E[\varphi_{t+k}],$$

where $\varphi_{t+k}$ is the expected client surplus in period $t+k$. Because the bank captures the entire surplus from client transactions, we can substitute $\varphi_{t+k}$ for $\phi_{t+k}$ and express the value function recursively as a Bellman equation

$$V(h_t) = \max_{\sigma_t} \varphi_t(h_t, \sigma_t) + \delta E_t[V(h_{t+1})], \tag{4}$$

where $\sigma_t = \sigma(h_t, n)$ is the bank's choice of technology at time $t$ for each possible client type.

The Bellman equation is simplified by restricting the history of information in the bank's time $t$ strategy to the present stage of technological development, $i$, and the number of experimental failures, $\tau$, at stage $i$:

$$V^i_\tau = \max_{\sigma_t} \varphi^i_\tau(\sigma_t) + \delta E_t[V^j_\kappa]. \tag{5}$$

The first term in (5) is the bank's surplus from the transaction at time $t$ and the second term is the continuation value conditional on the bank's strategy and its outcome. Specifically, if, at time $t$, the bank does not experiment with technology $i+1$, then $j = i$ and $\kappa = \tau$. If the bank carries out a successful experiment, $j = i+1$ and $\kappa = 0$ and if it is unsuccessful, $j = i$ and $\kappa = \tau + 1$.

As noted earlier, restricting the state variables in $h_t$ to $i$ and $\tau$ is without loss of generality. In Appendix 2, we demonstrate that the solution to equation (5) is the unique solution to the general case represented by equation (4). We now characterize the solution to equation (5) in steps.

**Proposition 2.** *There exists a level of technological development, $i^*$, at which the bank will cease experimentation.*

Proposition 2 reflects the effect of our assumption of decreasing returns to innovation. As the stage of technological development increases, the marginal benefit from continued experimentation diminishes while the cost, in the form of a smaller current payoff, remains the same. To see this more clearly, note that the expected benefit from an additional increment of technological development from $i^*$ to $i^*+1$ is $\frac{p_{i^*+1}(\rho-M)}{1-\delta}$ while the cost of continued experimentation is at least $(\pi_\tau^{i^*+1}\rho - M)$. Because $p_{i^*+1}$ approaches zero as $i^*$ goes to infinity but the cost of experimentation does not, there exists an $i^*$ such that innovation stops immediately after this stage of technological development is achieved.

The proof of proposition 2 is informative and so we provide it here.[15] We begin by conjecturing that there does not exist an $i^*$ at which the bank would optimally stop. At stage $i$ of technological development, conditional on a type $n \leq i+1$ client, the bank has two choices. In equilibrium, where the bank experiments, the value function is

$$V_\tau^i(n) = \pi_\tau^{i+1}\rho + \delta(\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\pi_\tau^{i+1}\rho)V_{\tau+1}^i). \tag{6}$$

Out of equilibrium, where the bank ceases experimentation, there are two possibilities. If $n = i+1$, the bank deploys the (safe) technology 0 and the value function is $M + \delta V_M^i$ where

$$V_M^i \equiv \frac{P_i\rho + (1-P_i)M}{1-\delta}$$

and $P_i = \sum_{n \leq i} p_n$. Alternatively, if $n < i+1$, the bank uses technology $i$ and the value function is $\rho + \delta V_M^i$.

If the bank deploys the safe technology when it stops experimentation, it should stop if

$$M + \delta V_M^i \geq \pi_\tau^{i+1}\rho + \delta(\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\pi_\tau^{i+1}\rho)V_{\tau+1}^i) \tag{7}$$

Note that both $V_0^{i+1}$ and $V_{\tau+1}^i$ are less than $\frac{\rho}{1-\delta}$, which is the hypothetical upper bound on the bank's expected present value were it able to deliver the highest possible expected payoff, $\rho$, for every type of client in every period. A sufficient condition for satisfying (7) is therefore

[15] This proof is general in the sense that it does not assume that the history of information, $h_t$, is restricted to $i$ and $\tau$ nor does it assume that, once the bank ceases experimentation, it will not restart as we assumed in Section 4.1.

$$M - \pi_\tau^i \rho \geq \delta(\frac{\rho}{1-\delta} - V_M^i) \tag{8}$$

The left-hand side of (8) is the gain in current fee income from not experimenting. The right-hand side is the upper bound on expected gain from experimentation. To establish that condition (8) is satisfied note first that

$$\begin{aligned}\frac{\rho}{1-\delta} - V_M^i &= \frac{\rho}{1-\delta} - \frac{P_i\rho + (1-P_i)M}{1-\delta} \\ &= (1-P_i)\frac{\rho - M}{1-\delta}\end{aligned}$$

As $i \to \infty$, $P_i \to 1$ and the right-hand side of (8) goes to zero. By assumption (1), the left-hand side of (8) is greater than zero. The same argument applies to the case where $n < i+1$ and the bank deploys technology $i$ but with the left-hand side of (8) changed to $\rho - \pi_\tau^i \rho$. Q.E.D.

We now characterize $i^*$.

**Proposition 3.** *The stage, $i^*$, at which the bank optimally ceases experimentation is uniquely determined by*

$$i^* = \left[ \frac{\ln\left( \frac{M - \pi_0 \rho}{\frac{\delta(1-g)}{(1-\delta)}\pi_0\rho(\rho - M)} \right)}{\ln g} \right]^+$$

*where $g = \frac{p_n}{p_{n-1}}$ and $[\cdot]^+$ is the integer function.*

We can now solve the bank's optimization problem by backward induction from stage $i^* - 1$. As a first step, note that the cost of experimentation on a client type $n = i+1$, $(M - \pi_\tau^{i+1}\rho)$, is less than the cost of experimentation on client type $n < i+1$, $(\rho - \pi_\tau^{i+1}\rho)$. Thus we refer to the former as a *low-cost client* and the latter as a *high-cost client.* Potentially large future benefits from successful experimentation may lead the bank to experiment on high-cost clients in spite of the fact that they would be better served by the bank's existing technology. The following proposition establishes the conditions under which a bank at technological stage $i$ will not experiment on high-cost clients.

**Proposition 4.** *The bank will not experiment on (high-cost) client types $n < i+1$ if*

*(1) $\pi_0 < \frac{1}{1+g}$ at the beginning of any given stage of technological development,*

*or*

*(2)* $i = i^* - 1$.

Condition (1) establishes a lower bound on the prior probability of success for experimentation on high-cost clients. The second condition reflects the elimination of future benefits beyond the terminal stage of experimentation determined by Proposition 3 and confirms that it is never optimal to experiment on type 0 clients in the simple case examined in Section 4.1.

The bank's strategy comprises three experimental regions with boundaries at any stage $i$ of technological development determined by the probability, $\pi_\tau^{i+1}$, of successful experimentation conditional on $\tau$ unsuccessful trials relative to cutoff probabilities $\pi_h^{i+1}$ and $\pi_l^{i+1}$ where $\pi_h^{i+1} \geq \pi_l^{i+1}$. The following proposition provides a formal statement of the bank's strategy.

**Proposition 5.** *The solution to the bank's optimization problem is characterized by determining* $(V_0^i, \pi_h^{i+1}, \pi_l^{i+1})$ *recursively from* $i^*$ *at each technological stage* $i$.

*i) There exists a high-experimental region in which* $\pi_\tau^{i+1} > \pi_h^{i+1}$ *and the bank experiments on all client types* $n \leq i+1$*;*

*ii) There exists a low-experimental region in which* $\pi_h^{i+1} \geq \pi_\tau^{i+1} \geq \pi_l^{i+1}$ *and the bank only experiments on* $n = i+1$ *clients;*

*iii) There exists a no-experimental region in which* $\pi_\tau^{i+1} < \pi_l^{i+1}$ *and the bank ceases experimentation forever;*

Proposition 5 illustrates that even in the absence of informational friction, banks will experiment indiscriminately when the environment for financial innovation is especially promising. They do this because the potential future gains from a successful innovation outweigh the short-run costs to clients who would be better served by the bank's existing technology. Clients understand this conflict and pay a fee commensurate with expectations of lower "quality" service.

Note that the conflict at this stage of our analysis is a reflection of the short-lived nature of clients. In Section 6.1, we derive conditions under which the interests of a long-lived client would correspond precisely with those of the bank and therefore yield an optimal experimental strategy identical to that characterized in Proposition 5. Thus, we might interpret the bank as optimally mediating a relationship between the short-lived client and a hypothetical future self that would benefit from successful experimentation.

It is worth noting that the model suggests that neither aggressive nor more persistent experimentation is necessarily a reflection of short-sightedness. In fact, a more patient bank will experiment more in both the high- and low- experimental regions and it will experiment longer than an impatient bank.

The model also illustrates that if financial innovation is subject to decreasing returns, then even the most successful innovators will, at some point, curb experimentation that is detrimental to at least some of their clients.

## 5. Innovation with Moral Hazard

In this section we assume that the bank observes the client's type but the client does not. This is a shorthand for the idea that clients hire bankers because they have a superior understanding of the client's requirements or the state of the world in which the requirements are to be satisfied. This assumption implies that the bank will be unable to commit to experiment on some client types but not on others. It follows immediately that the bank's only incentive-compatible strategy is to experiment on all client types or not at all in an optimal equilibrium. Therefore, the bank's problem is to solve equation (5) subject to this constraint. As we demonstrate in the next proposition, the incentive-compatibility constraint is binding and therefore the expected present value of the bank's fee stream is lower than in the preceding section where client type was public information.

The incentive-compatibility constraint also simplifies the solution to the bank's optimization problem. Specifically, in contrast to Proposition 5, there no longer exists a low-experimental region. The solution strategy is otherwise analogous to that used in Section 4.2. Proposition 6 summarizes the solution to the bank's optimization problem.

**Proposition 6.** *The solution to the bank's optimization problem is characterized by*

*i) At any stage $i$, the bank experiments on all client types $n \leq i+1$ or not at all;*

*ii) At stage $i$ bank will continue experimentation if and only if*

$$\pi_\tau^{i+1} \geq \underline{\pi}^{i+1} \equiv \frac{P_i\rho + p_{i+1}M}{P_{i+1}\rho\left[\delta\left(V_0^{i+1} - \frac{P_i\rho+(1-P_i)M}{1-\delta}\right)+1\right]}. \tag{9}$$

*iii) The bank ceases experimentation at stage $i^{**}$ which is uniquely determined by*

$$i^{**} = \begin{cases} \left[\frac{\ln x_+}{\ln g}\right]^+ & \text{if } \left[\frac{\ln x_+}{\ln g}\right]^+ \geq 0 \\ 0, & \text{otherwise.} \end{cases} \tag{10}$$

*where $x_+$ is the positive root of equation (41) in Appendix 1.*

*iv) At each stage (i.e., $\underline{\pi}^{i+1} > \pi_l^{i+1}$ and $i^{**} < i^*$), the bank ceases innovation after fewer failed trials and at a lower stage of technological development than it would in the absence of moral hazard.*

Part *i* of Proposition 6 identifies the moral hazard associated with the bank having a better understanding of the client's needs. In contrast to the benchmark case where the bank experimented on all client types $n \leq i+1$ only when the expected benefits to future clients were exceptionally large, now the bank engages in indiscriminate experimentation. This is a more costly experimental strategy because the net expected gain from financial innovation is smaller. Therefore, as part *iv* of Proposition 6 indicates, the bank will cease experimentation after a fewer number of failed trials and at a lower stage of technological development than in the benchmark case where client type is common knowledge. In principle, indiscriminate experimentation could be curbed with continuation equilibria where beliefs that such behavior occurred serve as a punishment mechanism. One such equilibrium is provided in the proof of Proposition 6 in Appendix 1. The optimality of the equilibrium described in Proposition 6 implies that such inter-temporal incentives, or a concern for a behavioral reputation, are not an efficient solution to the problem.

Assuming that banks have private information about how best to serve their clients, the model predicts that banks will more aggressively subordinate their clients' interests to a strategy that furthers their own (and future clients') interests in providing a broader range of products and services. In practice, we would expect to observe banks more aggressively promoting products and services that stem from and extend technological advances. Relative to existing products and services, they might be interpreted as more complex as they seek to more precisely contract over particular states of the world that may not be of concern to a wide range of clients.

But this is an inefficient strategy. The model suggests that the inability to commit to a more "responsible" experimental strategy leads to underdevelopment of even the most productive veins of innovative opportunity and leaves behind a trail of poor outcomes. Alternatively, any mechanism that would improve the bank's ability to commit to not exploit its private information would reduce client exposure to undesirable experimentation and improve the bank's capacity for serving a broader range of clients in the future. We explore this possibility next.

## 6. Curbing Indiscriminate Experimentation

In this section we examine the potential for curbing indiscriminate experimentation that is both damaging to some clients and leads to premature cessation of experimentation. If, as our model suggests, it is in the bank's interest to moderate its pursuit of financial innovations, one would expect to observe institutional responses aimed at curbing irresponsible behavior.

In fact, there are a number of governance mechanisms that banks might use to, at least in part, make more credible commitments to their clients. Possibilities include informal mechanisms, such as bank culture and reputation concerns; explicit oversight involving risk management and compliance; and indirect mechanisms such as compensation schemes and organizational structure. Of course, these, and other, governance mechanisms are imperfect both individually and in concert. We begin our analysis in this section by examining the conditions under which a relational contract between the bank and a long-lived client can be sustained.

### 6.1 *Experimentation Under a Relational Contract*

The fundamental problem facing the bank when it has private information about the client's type is the inability to commit to an efficient innovation strategy. Thus far we have stacked the deck against the bank by assuming that clients are short-lived. In practice, many firms are long-lived and relatively active in the capital markets. Moreover, dealing repeatedly, often exclusively, with a single bank over the course of many years was virtually the norm among capital market participants prior to the 1960s. Client relationships are less exclusive than they once were but banks still attempt to cultivate relationships with their clients.

We study the experimental process within a relationship by assuming that the bank encounters the same long-lived client in every period $t$.[16] The state of the client's requirements, $n$, at time $t$ is distributed identically to the short-lived clients in Section 3. As in the previous section, we assume that the bank has private information about $n$ and therefore the client cannot determine whether the bank takes an action that best corresponds with its needs.

We consider an agreement under which the client pays a fixed fee in each period equal to the maximum possible payoff of 1. In exchange, the bank agrees to provide an observable but unverifiable state-contingent rebate equal in value to the difference between the fixed fee and the client's realized payoff.[17] The structure of this agreement takes the form of the *stationary* relational contracts characterized by Levin (2003).[18] Because neither the client's type nor the value of the rebate are verifiable, the agreement must be self-enforcing.[19]

Note first that the agreement leads the bank to fully internalize the cost of experimentation. For example, if the bank observes that the client's type in a given period is $n < i + 1$ and it opts to experiment on the client with the $i + 1$ technology, the expected rebate is $1 - \pi_\tau^{i+1}\rho$. If, instead, the bank serves the client with the proven technology $i$, the rebate is $1 - \rho$. Thus the bank's cost of experimentation on a high-cost client is $\rho - \pi_\tau^{i+1}\rho$, which is equal to the client's marginal cost.

In the spirit of Baker et al. (2002), we consider whether there exist continuation payoffs associated with suboptimal, or "bad", equilibria that pose a sufficient threat to

[16]There is no loss of generality in the following argument from assuming that the bank deals repeatedly with a single client rather than dealing repeatedly with multiple clients.

[17]For example, if the bank uses the safe technology to serve the client, the client's net payoff is $1 - M$. If the bank experiments with technology $i$, the client's rebate is 1 if the flexible technology yields 0 and 0 if it yields 1. Similarly, if the bank experiments on the client, the rebate is 1 if it fails and 0 otherwise.

[18]Note, however, that the contracting problem is simpler in the sense that only one party, the bank, can deviate from the agreement.

[19]In practice, the rebate could take on many forms, such as research or market making services, that would be relatively opaque in their value to anyone but the bank and the client.

prevent the bank deviating from the optimal equilibrium.[20] Suppose, for example, that the bank experiments with the $i+1$ technology at each stage $i$, regardless of the client's type. Under this strategy, the expected payoff is $\rho\pi_\tau^{i+1}$ for a type $n \leq i+1$ client and 0 otherwise. This strategy is obviously inefficient because, for example, the bank learns nothing from experimentation on a type $n > i+1$ client and the client receives the lowest possible payoff of 0. But it is indeed an equilibrium strategy in which the client pays the expected surplus $\rho\pi_\tau^{i+1}P_{i+1}$ every period. The bank has no incentive to deviate from this "bad" equilibrium because even if it does so by taking a more efficient action, neither client beliefs about the bank's future actions nor the bank's future fee income will change.

Now suppose that this "bad" equilibrium is the continuation equilbrium if the bank fails to pay the rebate. The bank's largest possible per period gain from reneging on the agreement is 1. Thus the bank has no incentive to renege if

$$1 \leq \delta(V_\tau^i(\text{optimal equilibrium}) - V_\tau^i(\text{bad equilibrium})). \tag{11}$$

At any point in the game, the bank's next period value in the optimal equilibrium is at least the value, $V_M^i$, of going forward with no further experimentation and in the bad equilibrium it is at most $\frac{\pi_0\rho}{1-\delta}$. Therefore, condition (12) is satisfied if

$$1 \leq \delta\left(V_M^i - \frac{\pi_0\rho}{1-\delta}\right). \tag{12}$$

Because $V_M^i \equiv \frac{P_i\rho+(1-P_i)M}{1-\delta} \geq \frac{M}{1-\delta}$, condition (12) is satisfied if

$$1 \leq \frac{\delta}{1-\delta}(M - \pi_0\rho) \tag{13}$$

Assumption (1) implies that the right-hand side of (13) is positive, so a sufficiently patient (high $\delta$) bank will have no incentive to deviate from the optimal equilibrium.

This punishment mechanism is crucially dependent upon the observability of the rebate. With this being the case, deviation from the agreement will always be punished with a shift to a suboptimal continuation equilibrium payoff. Thus we need only to demonstrate that there exists an off-equilibrium path that meets the condition above.

[20]It is trivial to demonstrate that the client has no incentive to renege because, by assumption, it does not share in the surplus.

With this being the case, there is no need for punishment along the optimal equilibrium path, in which case we can achieve the first-best outcome described in Proposition 5.[21]

Note further that there is no loss of generality from assuming that the bank captures the entire surplus. Were the bank and client to share the surplus, the spread between $M$ and $\pi_0\rho$ would be smaller. Therefore satisfying the bank's incentive-compatibility condition (11) would require the bank to have greater patience.

In practice, the model suggests that more valuable relationships should be characterized by less rigidity and greater scope for renegotiation. Alternatively, as relationships become less exclusive, deviation from the relational contract is more likely.

### 6.2 *Monitoring as a Commitment Device*

In this section we assume that the bank's capacity for innovation is embedded in the human capital of a banker who also has private knowledge of the client's type. It is clear from Proposition 6 that the bank would like for the banker to refrain from indiscriminate experimentation but equally clear that the banker suffers the same commitment problem. For the moment, we assume that the bank can credibly commit to a monitoring mechanism that curbs the banker's experimentation on high-cost ($n \leq i$) clients.

We assume that the monitoring mechanism is subject to two imperfections. First, it can only prevent experimentation with probability $\mu < 1$. Second, it is inflexible in the sense that it cannot distinguish between the conditions that favor the high- and low-experimental regions in Proposition 5. Thus the monitoring mechanism is not sufficiently discriminating to achieve the benchmark equilibrium summarized in Proposition 5 because it will inefficiently limit the banker's pursuit of exceptionally promising opportunities that would lie in the high-experimental region. However, it will curb some irresponsible experimentation on high-cost clients (e.g., when $\pi_0 < \frac{1}{1+g}$).[22]. The efficiency of the monitoring mechanism therefore rests on how promising is the environment for innovation.

The optimal equilibrium under this monitoring mechanism follows from a straight-

[21] In Levin (2003), the principal cannot observe the agent's action and therefore there must be at least some threat of punishment along the equilibrium path that will result in an efficiency loss.

[22] See Proposition 4

forward modification in the analysis of Section 5. The only difference is that instead of experimenting for sure when client type is $n \leq i$, now the bank can only experiment with probability $1-\mu$. This gives rise to the following necessary condition for continued experimentation:

$$\pi_{\tau}^{i+1} \geq \underline{\pi}_{\mu}^{i+1} \equiv \frac{P_i(1-\mu)\rho + p_{i+1}M}{\left(P_i(1-\mu)+p_{i+1}\right)\rho\left[\delta\left(V_{0,\mu}^{i+1} - \frac{P_i\rho+(1-P_i)M}{1-\delta}\right)+1\right]}. \tag{14}$$

The following proposition summarizes the stage in the optimal equilibrium at which condition 14 is violated and the bank ceases experimentation:

**Proposition 7.** *In the optimal equilibrium with monitoring, the bank ceases experimentation at stage $i_{\mu}^{*}$, where $i_{\mu}^{*}$ is the smallest integer that satisfies*

$$\underline{\pi}_{\mu}^{i_{\mu}^{*}+1} \geq \pi_0.$$

Although monitoring leads to less aggressive experimentation at any stage of technological development, curbing experimentation on high-cost clients enables the bank to pursue a higher level of technological development than would be possible in its absence.

**Proposition 8.** *In the optimal equilibrium with monitoring, $i_{\mu}^{*}$ is increasing in $\mu$.*

This result follows directly from condition (14). The average cost of experimentation, $\frac{P_i(1-\mu)\rho+p_{i+1}M}{(P_i(1-\mu)+p_{i+1})}$, is decreasing in $\mu$ because monitoring reduces the probability of experimentation on high-cost clients. Since $V_{0,\mu}^{i_{\mu}^{*}+1} = V_{M}^{i_{\mu}^{*}+1}$does not depend on $\mu$, $\underline{\pi}_{\mu}^{i_{\mu}^{*}+1}$ is decreasing in $\mu$ and $i_{\mu}^{*}$ is increasing in $\mu$.[23]

In the present form of the model, we cannot conclude that monitoring strictly increases the value of the bank although it will do so at low levels of $\pi_0$ where experimentation on $n \leq i$ clients is especially costly. However, if we permit the level of monitoring to be endogenously determined by the trade-off between curbing irresponsible experimentation and inefficiently limiting experimentation on $n \leq i$ clients, it is generally true that monitoring increases bank value.

[23]Proposition 8 does not imply that the bank's tolerance for experimental failure at any given stage of technological development also is increasing in $\mu$. It is possible that $V_{0,\mu}^{i+1}$ is decreasing in $\mu$ because inflexibility in the monitoring mechanism prevents desirable experimentation on type $n \leq i$ clients. In this case, $\underline{\pi}_{\mu}^{i+1}$ may not be decreasing in $\mu$.

### 6.3 *Monitoring and Banker Mobility*

In the previous section we implicitly assumed that bankers and their human capital were inseparable from the bank. Historically, bank partners rarely left one partnership for another. However, as we noted in the introduction, the "revolution" in financial innovation beginning in the 1970s was accompanied by far greater mobility among senior bankers. In this section we consider how banker mobility and financial innovation are related in our model.

We continue to assume that innovative capacity and private information about client types is embedded in the banker's human capital. Further, we assume that human capital is scarce but that banks are not, so that the banker captures the surplus from her efforts. Banks select a fixed and costless monitoring capacity at the outset of the game. Later, we consider costly monitoring and sharing of surplus between the bank and banker. For simplicity, we assume that in each period, prior to observing the client's type, the banker can move costlessly between a bank that monitors ($\mu > 0$) and one that does not ($\mu = 0$). Thus, in each period, conditional on beliefs about the promise in future experimentation, the banker selects the monitoring technology that best complements her opportunities.

In this setting, the analysis will be quite similar to that in the benchmark case in Section 4.2. Specifically, we will again obtain high- and low- experimental regions determined by the banker's degree of optimism about the environment for innovation. Intuitively, the banker's ability to move from one bank to another yields a verifiable commitment to limit indiscriminate experimentation that would otherwise lead to a premature halt to financial innovation.

Similar to Section 4.2, if the banker does not experiment on the client, the expected current period fee income is

$$\varphi_{\tau}^{i}(\sigma_t, \mu_t) = P_i\rho + (1 - P_i)M.$$

If she does experiment on the client, the expected current period fee income is

$$\varphi_{\tau}^{i}(\sigma_t, \mu_t) = P_i\mu_t\rho + (p_{i+1} + P_i(1 - \mu_t))\pi_{\tau}^{i+1}\rho + (1 - P_{i+1})M$$

The banker's problem is

$$V_\tau^i = \max_{\sigma_t, \mu_t} \varphi_\tau^i(\sigma_t, \mu_t) + \delta E_t[V_\kappa^j]$$

where the bank choice in period $t$ is indicated by $\mu_t = 0$ or $\mu$.

When the environment for innovation is more promising, the banker will benefit from more aggressive experimentation by joining the $\mu_t = 0$ bank. Alternatively, it is in the banker's interest to experiment less aggressively in a less promising environment if clients believe that it will do so. By moving to the bank that will constrain experimentation through its monitoring mechanism, the banker can more credibly commit to avoid indiscriminate experimentation. Thus, as in Section 4.2, banker mobility leads to both high- and low- experimental regions with the banker shifting between them as it learns about the environment.

Once again, we solve the model by backward induction. The boundary of the low-experimental region is

$$\pi_\tau^{i+1} \geq \pi_l^{i+1} \equiv \frac{P_i(1-\mu)\rho + p_{i+1}M}{(P_i(1-\mu)\rho + p_{i+1})\rho\left[\delta\left(V_0^{i+1} - V_M^i\right) + 1\right]} \tag{15}$$

and for the high-experimental region it is

$$\pi_\tau^{i+1} \geq \pi_h^{i+1} \equiv \frac{\rho + \delta(V_0^{i+1} - V_{\tau+1}^{i,l})}{\rho[1 + \delta(V_\tau^{i,l} - V_{\tau+1}^{i,l})]}.$$

The stage, $i_\mu^*$, at which the banker optimally ceases experimentation is determined by

$$i_\mu^* = \arg\min_i\left[\frac{P_i(1-\mu)\rho + p_{i+1}M}{(P_i(1-\mu)\rho + p_{i+1})\rho\left[\delta\left(V_M^{i+1} - V_M^i\right) + 1\right]} \geq \pi_0\right]$$

Note that this expression is identical to condition 14 and $i_\mu^*$ is the same as that characterized in Proposition 7. $V_\tau^i$ is derived exactly as in the benchmark case by backward induction from $i_\mu^*$. Thus, we establish that the solution is unique by the same argument as that used in the benchmark case.

The following proposition summarizes the solution to the banker's optimization problem:

**Proposition 9.** *The unique solution to the banker's optimization problem is fully characterized by determining* $(V_\tau^i, \pi_h^{i+1}, \pi_l^{i+1})$ *recursively from* $i_\mu^*$ *at each technological stage* $i$.

*i) There exists a high-experimental region in which the banker opts to work without monitoring if and only if* $\pi_\tau^{i+1} \geq \pi_h^{i+1}$;

*ii) There exists a low-experimental region in which the banker opts to work with monitoring if and only if* $\pi_l^{i+1} \leq \pi_\tau^{i+1} < \pi_h^{i+1}$;

*iii) The bank ceases experimentation when* $\pi_\tau^{i+1} < \pi_l^{i+1}$;

*iv) The decision to cease experimentation is reached at stage* $i_\mu^*$.

*v) The present value of the banker's strategy,* $V_\tau^i$, *exceeds that from the bank's strategy described in Proposition 6.*

Proposition 9 implies that a promising environment for innovation drives talent toward banks that place weaker restrictions on experimentation and vice versa. Given the ability to verifiably accept more or less monitoring by switching banks, the banker engages in less indiscriminate experimentation on high-cost clients, reaches a higher level of technological development, and generates greater surplus. The banker does not achieve the level of performance exhibited by the complete contracting benchmark of Section 4.2 and the relational contract examined in Section 6.1. The deadweight loss stems from the fact that the monitoring mechanism remains imperfect and therefore will prevent experimentation in some instances in which it would have occurred in the absence of friction introduced by the banker's private knowledge of client type.

The results are not sensitive to the assumptions that monitoring is costless and that the banker captures the entire surplus. For example, suppose that the bank faces a one-time cost $c > 0$ to maintain a monitoring level of $\mu > 0$ and that it can refuse a banker seeking monitoring services. If the bank and the banker are able to commit to a long-term contract to share the surplus of innovation, then the optimal experimentation strategy characterized by Proposition 9 can still be implemented.

To see that this is true, assume that the banker and the bank are able to agree to a long-term monitoring contract via Nash bargaining prior to experimentation. Assume further that the contract requires the banker to pay the bank a monitoring fee comprising a fixed stream of payments independent of the state of innovation. In exchange, the bank provides monitoring services to the banker when it is requested to do so. Once the contract is signed, the cost of monitoring is sunk and therefore has no impact on the banker's experimentation strategy. Thus, the solution to the banker's optimization problem is identical to the case described in Proposition 9. From the bank's perspective, for our claim to hold, we need only for its bargaining power, $\lambda$, to be sufficiently strong and its cost of monitoring, $c$, to be sufficiently low that it can cover its cost of monitoring or

$$\lambda(V_0^{0,\mu} - V_0^0) \geq c,$$

where $V_0^{0,\mu}$ is the banker's present value as characterized by Proposition 9 and $V_0^0$ is the bank's present value without monitoring as characterized by Proposition 6.[24]

Finally, if we assume that the bank cannot commit to a long-term contract but that the banker can, then period-by-period contracting may be feasible. If so, the present value of the banker will be lower as a consequence of a suboptimal experimentation strategy and the necessity of sharing a larger fraction of the surplus with the bank going forward.[25] Regardless, it remains the case that a promising environment for innovation will drive bankers toward banks with weaker monitoring systems and vice versa.

## 7. Discussion

Our model treats financial innovation as a learning process in which bank clients are subjects of experimentation. As is true of any experimental setting that involves human subjects, those who receive the "treatment" may or may not benefit from its application. Moreover, the probability of a succesful outcome varies across experiments and, crucially for our purposes, is better understood by the experimenter or, in our case, the banker. Thus, there is a moral hazard in the experimental process of financial innovation.

In principal, the potential negative consequence of some forms of financial innovation

[24]This result is robust to imperfect competition among banks in which one bank has client-specific knowledge of the banker.

[25]As shown in the incomplete contract literature, e.g., Hart and Moore (1990), bargaining over surplus after investments are made in general leads to underinvestment. The same reason causes the banker to cease experimentation earlier than in the case described in Proposition 9.

could, at least partially, be internalized by the innovators.[26] For example, in the late 1990s, J.P. Morgan's conducted early stage experimentation with synthetic collateralized loan obligations (CLOs) on its own account.[27] Similarly, prior to restrictions on proprietary trading following the 2008 financial crisis, banks had considerable scope for internal experimentation through their proprietary trading operations. However, in most instances, financial product innovations are client oriented.

Recognizing that there is risk in any experimental process, our model enables us to explore several mechanisms for minimizing the burden borne by clients who will not benefit from an innovative product or service. In our model, clients are sufficiently sophisticated to recognize (and price) the moral hazard. Thus, it is in the interest of banks to find means by which they can commit to avoid reckless pursuit of innovative products in environments where the probability of success is relatively low. We describe a self-enforcing relational contract that curbs such behavior and contend that this strategy is a useful way of thinking about financial innovation prior to the 1970s when bankers were relatively immobile and banking relationships were relatively excusive. It is also worth noting that J.P. Morgan, which maintained strong relationships with many prominent corporations, stepped back from the synthetic CLO market, relative to its peers, when it identified risks that it perceived as unmanageable.[28]

The subsequent rapid pace of advances in computing technology and applied financial economic theory provided banks with a vastly broader "flexible technology" platform for experimentation. Coupled with weakening client relationships, our model suggests that it is not surprising that we entered a period of "revolutionary" financial innovation that was often perceived as antagonistic.[29] These conditions are precisely those under which it would be especially difficult for banks to make credible commitments to place client interests first.

[26]Our model does not admit the potential for negative externalities that may have systemic consequences, nor, too often, do practitioners.

[27]See Tett (2009).

[28]Tett (2009, p.68).

[29]One might argue that weakening client relationships led banks to be relatively indiscriminate in experimentation with junk-bond financing during the highly antagonistic hostile takeover movement of the 1980s. Goldman Sachs, which arguably wished to maintain a relatively strong commitment to client relationships during this period, sought to separate itself from peers by refusing buy-side advisory opportunities. See Chen et al. (2015, p.1175).

The same technological advances that spurred financial innovation also enabled, at least in principle, improved monitoring of financial experimentation. Specifically, the "value at risk" (VaR) and "risk-adjusted return on capital" (RAROC) systems, pioneered by JP Morgan and Bankers Trust, set the stage for more sophisticated, albeit imperfect, risk assessment and internal monitoring mechanisms. However, our model suggests that internal monitoring systems may have limited value in curbing indiscriminate innovation. No monitoring system is perfect and governance of innovative activity is further compromised by the challenge of identifying states of the world in which experimentation is especially likely to result in successful innovation. Moreover, it is difficult for clients to determine whether a bank's claim to monitoring behavior of its employees is credible.

While increasing banker mobility may have undermined banks' capacity for relational contracting with their clients, our model suggests that greater mobility complements monitoring systems. But for that matter, any means by which a bank or banker could more credibly commit to curbing inefficient experimentation would improve capacity for relational contracting. Recall that the relational contract described in Section 6.1 calls for an observable but unverifiable state-contingent rebate. An observable informal commitment to making a market for a complex product could serve this purpose as could consistent retention of residual tranches for securitized products.

Alternatively, "outsourcing" monitoring to the client could also promote more efficient experimentation. For example, it is not uncommon for a sell-side M&A banker working in a full-service bank to be pressured to pitch a novel or complex financing product to the client. The increasingly common inclusion of a "boutique" investment bank in the client's advisory team may serve to temper incentives to pitch such products when they they do not best serve the client's interest. "Corporate finance advisers" often play a similar role in complex securities offerings.[30]

It is worth emphasizing once again that our model assumes that bank clients understand and price the moral hazard in financial innovation. This may be a plausible assumption when applied to institutional and corporate clients. It is perhaps less plausible for retail clients, especially to the extent that they do not deal directly with the

[30]See Farrell and Hoffman (2018).

producer of an innovative product. For example, structured notes aimed at retail investors are designed and issued by investment banks but sold by brokers to retail clients.[31] From the perspective of our model, it is perhaps not surprising that this market has been characterized by large-scale, and likely inefficient, experimentation.

## 8. Appendix 1: Proof of Propositions

**Proof of Proposition 1**. This is a special case of the case we study in Section 4.2. We begin by showing that the bank will never experiment on type 0 clients. First, suppose that the opposite is true, that is, there is a $\pi^1_\tau$ such that it is optimal to experiment on type 0 clients. Formally,

$$V^0_\tau(n=0) = \pi^1_\tau \rho + \delta[\pi^1_\tau \rho V^1_0 + (1-\rho\pi^1_\tau)V^0_{\tau+1}] \geq \rho + \delta V^0_\tau.$$

$V^0_\tau(n)$ is the bank's value when facing a client of type $n$ after $\tau$ failed trials. The above condition implies that it is also optimal to experiment on a $n=1$ type client as well because

$$V^0_\tau(n=1) = \pi^1_\tau \rho + \delta[\pi^1_\tau \rho V^1_0 + (1-\rho\pi^1_\tau)V^1_{\tau+1}] \geq \rho + \delta V^0_\tau.$$

Taking expectations, we have

$$V^0_\tau = E[V^0_\tau(n)] \geq \rho + \delta V^0_\tau,$$

Solving for $V^0_\tau$,

$$V^0_\tau \geq \frac{\rho}{1-\delta}, \tag{16}$$

which is impossible because the best per-period payoff conditional on type 1 clients and $\pi^1_\tau$ is less than $\pi^1_\tau\rho + (1-\pi^1_\tau)M$, which is only achievable if the bank knows exactly if technology 1 is effective or not and thus avoids any experimentation. In other words, the bank will not try ineffective technolgies or give up effective ones. But $\pi^1_\tau\rho+(1-\pi^1_\tau)M < \rho$.

Now we prove the experimental strategy is unique. Suppose $\pi^1_\tau > \underline{\pi}^1$ and the bank

[31] Shin (2018) notes that worldwide sales volume for equity-linked structured notes alone grew from $50 billion in 2001 to about $200 billion in 2016. See Egan (2019) for institutional details of the market as well as discussion of the literature on structured products.

does not experiment. Then the bank value is

$$V_\tau^0 = M + \delta V_\tau^0$$

which yields $V_\tau^0 = \frac{M}{1-\delta}$. But this is strictly less than what the bank can get in equilibrium as shown in (3).

Now we show that if $\pi_\tau^1 < \underline{\pi}^1$, the unique optimal strategy is to not experiment at all.

Proof. Suppose otherwise: there is a strategy $\widehat{\sigma}$ such that under a history $h^t$ in stage 0 with $\pi_\tau^1 < \underline{\pi}^1$ and $\widehat{\sigma}(h^t)$ specifies to experiment with a positive probability $\widehat{\sigma}_t(1) > 0$ on a type $n = 1$ client. The optimality of $\widehat{\sigma}(h^t)$ implies that

$$V_\tau^0(1) = \pi_\tau^1 \rho + \delta[\pi_\tau^1 \rho V_0^1 + (1 - \pi_\tau^1 \rho) V_{\tau+1}^0] \geq M + \delta \frac{p_1 M + (1-p_1)\rho}{1-\delta}$$

The last term is the bank's value if it ceases experimentation. Because $\pi_\tau^1 < \underline{\pi}^1$, by definition of $\underline{\pi}^1$ we have $V_{\tau+1}^0 > V_M^0 \equiv \frac{p_1 M + (1-p_1)\rho}{1-\delta}$, which implies that the bank has to experiment next period after $\tau + 1$ failed trial as shown in the next lemma.

**Lemma 1.** $V_{\tau+1}^0 > V_M^0$ *implies* $\widehat{\sigma}_{\tau+1}(1) = 1$.

Proof. If not, $\widehat{\sigma}_{\tau+1}(1) < 1$ implies

$$V_{\tau+1}^0 = p_1 M + (1-p_1)\rho + \delta V_{\tau+1}^0$$

that yields

$$V_{\tau+1}^0 = V_M^0$$

contradicting the assumption. Q.E.D.

Repeating the same argument inductively for all $\tau$, we have $V_{\tau+1}^0 > V_M^0$ all $\tau$. That is, the bank will keep experimenting in stage 0 no matter how many failed trials have occured. We now show that this strategy is not optimal. If the bank keeps experimenting, its per-period fee is bounded by

$$E[\varphi(n,\sigma)] \leq (1-p_1)\rho + p_1(\pi_{\tau+1}^1 \rho + (1 - \pi_{\tau+1}^1 M))$$

The right-hand side is the income the bank can get assuming it has perfect knowledge of whether technology 1 is effective or not. Therefore, for any $\kappa > 0$ the bank's value is

$$V^0_{\tau+\kappa+1} = \sum_{u=0}^{\infty} \delta^u E[\varphi(n, \sigma(h^{t+u})] \leq (1-p_1)\frac{\rho}{1-\delta} + p_1\pi^1_{\tau+\kappa+1}\frac{\rho}{1-\delta} + p_1(1-\pi^1_{\tau+\kappa+1})\frac{M}{1-\rho}. \tag{17}$$

If we fix $\kappa$,

$$\begin{aligned} V^0_{\tau+\kappa}(1) &= \pi^1_{\tau+\kappa}\rho + \delta[\pi^1_{\tau+\kappa}\rho V^1_0 + (1-\pi^1_{\tau+\kappa}\rho)V^0_{\tau+\kappa+1}] \\ &\leq \pi^1_{\tau+\kappa}\rho + \delta[\pi^1_{\tau+\kappa}\rho V^1_0 + (1-\pi^1_{\tau+\kappa}\rho)\left[(1-p_1)\frac{\rho}{1-\delta} + p_1\pi^1_{\tau+\kappa}\frac{\rho}{1-\delta} + p_1(1-\pi^1_{\tau+\kappa})\frac{M}{1-\rho}\right]] \end{aligned}$$

The inequality follows from (17). If we let $\kappa \to \infty$, then the right-hand side goes to

$$\delta\left[(1-p_1)\frac{\rho}{1-\delta} + p_1\frac{M}{1-\rho}\right] = \delta V^0_M$$

That is, given any $\epsilon > 0$, $\exists$ a $\kappa$ large enough such that $V^0_{\tau+\kappa}(1) \leq \delta V^0_M + \epsilon$. But by ceasing experimentation, the bank's value is at least $M + \delta V^0_M$, contradicting the optimality of $V^0_{\tau+\kappa}(1)$ for large $\kappa$.

The last part follows from the fact that $\underline{\pi}^1$ is increasing in $M$ but decreasing in $\rho$ and $\delta$. Q.E.D.

**Proof of Proposition 3**. Note that (7) is satisfied for $i^*$ with $\tau = 0$. Furthermore, (7) should also be satisfied for stages beyond $i^*$ but not satisfied for stages below $i^*$ with $\tau = 0$. In other words, it is optimal to stop experimenting in stages beyond $i^*$ but keep experimenting in stages below $i^*$. We will make this assumption and then verify that it holds. The assumption implies $V^{i^*+1}_0 = V^{i^*+1}_M$. Solving (7) with $\tau = 0$ yields

$$p_{i^*+1} \leq \frac{(1-\delta)(M-\pi_0\rho)}{\pi_0\rho\delta(\rho - M)}. \tag{18}$$

Note that if (7) is satisfied, then the bank would not experiment on a type $n < i^* + 1$ client because the cost of experimenting on such a client ($\rho$) is higher than the cost of experimenting on a type $i+1$ client ($M$). In contrast, the bank in stage $i^* - 1$ should not stop experimenting at $\pi_0$, which yields

$$p_{i^*} > \frac{(1-\delta)(M-\pi_0\rho)}{\pi_0\rho\delta(\rho-M)} \tag{19}$$

Since $p_i$ is monotonicly decreasing $i$, the two conditions, (18) and (19) uniquely determines $i^*$. Solving (18) with $p_{i^*+1} = p_0 g^{i^*+1}$ yields the desired result. Note that the monotonicity of $p_i$ implies that (18) holds for all $i \geq i^*$ and (19) holds for all $i < i^*$, we thus verifies our earlier assumption that it is optimal to cease experimentation in stages beyond $i^*$ but experiment in stages below $i^*$. Q.E.D.

**Proof of Proposition 4:** We first prove part (1). At stage $i = i^* - 1$, suppose there is a $\pi_\tau^{i+1}$ such that it is optimal to experiment on types $n < i+1$, that is,

$$\begin{aligned} V_\tau^i(n < i+1) &= \pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] \\ &\geq \rho + \delta V_\tau^i \end{aligned} \tag{20}$$

which implies that if the client is of type $n = i+1$ it is also optimal to experiment because

$$V_\tau^i(n = i+1) = \pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] \geq \rho + \delta V_\tau^i > M + \delta V_\tau^i$$

We also know that for clients of type $n > i+1$

$$V_\tau^i(n) = M + \delta V_\tau^i$$

Taking expectation, we have

$$V_\tau^i = E[V_\tau^i(n)] \geq \Pr[n \leq i+1]\rho + \Pr[j > i+1]M + \delta V_\tau^i$$

that is,

$$V_\tau^i \geq \frac{P_{i+1}\rho + (1-P_{i+1})M}{1-\delta}. \tag{21}$$

Or the standardized value is

$$v_\tau^i \equiv (1-\delta)V_\tau^i \geq P_{i+1}\rho + (1-P_{i+1})M = v(\pi_0^{i^*}).$$

The last term is the per-period payoff conditional on technology $i+1$ being effective. Conditional on $\pi_\tau^{i+1}$, the best per-period payoff the bank can achieve assuming it has perfect knowledge of whether technology $i+1$ is effective or not is

$$\pi_\tau^{i+1}\left[P_{i+1}\rho + (1-P_{i+1})M\right] + (1-\pi_\tau^{i+1})\left[P_i\rho + (1-P_i)M\right] < P_{i+1}\rho + (1-P_{i+1})M$$

That is, $v_\tau^i \geq v(\pi_0^{i^*})$ is impossible to hold.

Part (2) follows by the similar argument. Note that for a general $i$, the same argument in the proof of part (1) shows that it is optimal for the bank to experiment on types lower than $i+1$ only if (21) holds. As in part (1), $v_\tau^i$ is bounded by the full information value $z(\pi_\tau^{i+1})$. $z(\pi_\tau^{i+1})$ is the per-period payoff that the bank can achieve if it has perfect knowledge of whether the next stage technology is effective or not. That is

$$\begin{aligned} z(\pi_\tau^{i+1}) &\equiv \Pr[j < i+1]\rho + \Pr[j \geq i+1]M + (\rho - M)[p_i\pi_\tau^{i+1} + p_{i+1}\pi_\tau^{i+1}\pi_0 + p_{i+2}\pi_\tau^{i+1}\pi_0^2 + ...] \\ &= \Pr[j < i+1]\rho + \Pr[j \geq i+1]M + (\rho - M)\frac{p_i\pi_\tau^{i+1}}{1-g\pi_0} \end{aligned}$$

The term $(\rho - M)\frac{p_i\pi_\tau^{i+1}}{1-g\pi_0}$ is the expected payoff the bank can get from realizing all possible new effective technologies. Note that if

$$z(\pi_\tau^{i+1}) < \Pr[j \leq i+1]\rho + \Pr[j > i+1]M,$$

the bank should stop experimenting because $(1-\delta)V_\tau^i \leq z(\pi_\tau^{i+1})$. The above condition is simplified to be

$$(\rho - M)\frac{p_i\pi_\tau^{i+1}}{1-g\pi_0} - p_i(\rho - M) < 0,$$

or

$$\pi_\tau^{i+1} < 1 - g\pi_0. \tag{22}$$

In particular, if $\pi^0 < 1 - g\pi_0$, i.e., $\pi^0 < \frac{1}{1+g}$, then it is never optimal for a bank to experiment on $n < i+1$ at all. Q.E.D.

**Proof of Proposition 5.** We here characterize the optimal solution as explicitly as possible. The state variables are the bank's technology stage $i$ and the number of failed experiments $\tau$. Note that $\tau$ uniquely determine $\pi_\tau^{i+1}$ through Bayes rule. We first start with the no-experimental region in which the bank will cease experimentation. We then characterize the low-experimental region in which the bank experiments only on low cost clients, i.e., clients of type $i+1$. Finally we characterize the high-experimental region where the bank experiments on both low- and high-cost clients, i.e., $n \leq i+1$.

*No-Experimental Region*

We first characterize $\pi_l^{i+1}$. At stage $i$ condition (8) yields that the bank should cease experimentation if

$$\pi_\tau^{i+1} \leq \pi_l^{i+1} \equiv \frac{M}{\rho(1+\delta(V_0^{i+1} - V_M^i))}. \tag{23}$$

But the bank should experiment at $\tau_l - 1$, that is

$$\pi_{\tau_l - 1}^{i+1} > \pi_l^{i+1}. \tag{24}$$

By Bayes rule,

$$\begin{aligned} \pi_\tau^{i+1} &= \frac{\pi_0(1-\rho)^\tau}{\pi_0(1-\rho)^\tau + (1-\pi_0)} \\ &= \frac{1}{1 + \frac{(1-\pi_0)}{\pi_0}(1-\rho)^{-\tau}} \end{aligned} \tag{25}$$

Substituting (25) into (23) and (24) and simplifying, we find the optimal $\tau_l$ for stage, $\tau_l^i$

$$\tau_l^i = 1 + \left[ \frac{\ln\left(\frac{1-\pi_0}{\pi_0}\right) - \ln\left(\frac{1-\pi_l^{i+!}}{\pi_l^{i+1}}\right)}{\ln(1-\rho)} \right]^+ \tag{26}$$

where $[\cdot]^+$ is the truncation function.

*Low-Experimental Region*

If $\pi_\tau^{i+1} > \pi_l^{i+1}$, it is optimal for the bank to experiment on low cost clients. Now we characterize the upper bound, $\pi_h^{i+1}$, below which it is optimal to experiment only on low cost clients but not on high cost clients. Now suppose the bank is facing a high cost client, $n < i+1$. If the bank does not experiment as the optimal strategy requires, the bank's value function is

$$V_\tau^{i,l}(n < i+1) = \rho + \delta V_\tau^{i,l}$$

If on the other hand the bank deviates by experimenting once but conforms to the the optimal strategy thereafter, the bank's value is

$$V_\tau^{i,h}(n < i+1) = \pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^{i,l}]$$

$V_\tau^{i,l} \geq V_\tau^{i,h}$ implies

$$\rho - \pi_\tau^{i+1}\rho \geq \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^{i,l} - \delta V_\tau^{i,l}]. \tag{27}$$

Or

$$\pi_\tau^{i+1} \leq \frac{\rho + \delta(V_\tau^{i,l} - V_{\tau+1}^{i,l})}{\rho(1+\delta(V_0^{i+1} - V_{\tau+1}^{i,l}))} \tag{28}$$

where $\pi_\tau^{i+1}$ is given by (25). So if we know $V_{\tau+1}^{i,l}$, $V_\tau^{i,l}$,and $V_0^{i+1}$, we can solve for $\tau_h^i$ such that (28) is satisfied. But the problem is that it is not clear (28) will determine a unique $\tau_h^i$. In other words, we haven't show that the bank should only experiment on type $i+1$ clients when $\tau_l^i < \tau \leq \tau_h^i$. We next show it is indeed the case.

We first solve for $V_\tau^{i,l}$. Taking expectation of $V_\tau^{i,l}(n)$ given the optimal strategy, we have

$$V_\tau^{i,l} = E[V_\tau^{i,l}(n)] = P_i\rho + M(1-P_{i+1}) + p_{i+1}\left[\pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^{i}]\right] + (1-p_{i+1})\delta V_\tau^{i,l}$$

Substituting (27) into the bracket, we get

$$V_\tau^{i,l} \leq P_{i+1}\rho + M(1-P_{i+1}) + \delta V_\tau^{i,l},$$

which is

$$V_\tau^{i,l} \leq \frac{P_{i+1}\rho + M(1-P_{i+1})}{1-\delta}. \tag{29}$$

In other words, (27) is equivalent to (29). $\tau_h^i$ is uniquely determined by

$$V_{\tau_h^i - 1}^{i,l} > \frac{P_{i+1}\rho + M(1-P_{i+1})}{1-\delta} \tag{30}$$

$$V_{\tau_h^i}^{i,l} \leq \frac{P_{i+1}\rho + M(1-P_{i+1})}{1-\delta}. \tag{31}$$

Furthermore, because $V_\tau^{i,l}$ is decreasing in $\tau$ (because a lower-$\tau$ bank can mimic the same strategy as the higher-$\tau$ bank and get a higher value), the bank should only experiment on $i+1$ type clients when $\tau_l^i < \tau \leq \tau_h^i$. To solve for $\tau_h^i$ as explicitly as possible, we need to solve for $V_\tau^{i,l}$ as explicitly as possible for a given $V_0^{i+1}$. $V_\tau^{i,l}$ can be solved recursively

$$V_\tau^{i,l} = p_{i+1}(\pi_\tau^{i+1}\rho+\delta[\pi_\tau^{i+1}\rho V_0^{i+1}+(1-\pi_\tau^{i+1}\rho)V_{\tau+1}^{i,l}])+P_i(\rho+\delta V_\tau^{i,l})+(1-p_{i+1}-P_i)(M+\delta V_\tau^{i,l}) \tag{32}$$

together with the boundary condition

$$V_{\tau_l^i}^{i,l} = V_M^i. \tag{33}$$

Because $\pi_\tau^{i+1}$ is non-linear in $\tau$, it is hard to solve the above difference equation. The way around it is to solve $V_\tau^{i,l}$ by first solving $V_\tau^{i,l}$ conditional on the effectiveness of $i+1$ technology, i.e., $V_\tau^{i,l}$(effective technology $i+1$) and $V_\tau^{i,l}$(ineffective technology $i+1$), separately. Then the solution is given by

$$V_\tau^{i,l} = \pi_\tau^{i+1}V_\tau^{i,l}(\text{effective technology } i+1) + (1-\pi_\tau^{i+1})V_\tau^{i,l}(\text{ineffective technology } i+1).$$

Under an ineffective technology $i+1$, there will be no success and (32) becomes

$$v_\tau = p_{i+1}\delta v_{\tau+1} + P_i(\rho+\delta v_\tau) + (1-p_{i+1}-P_i)(M+\delta v_\tau).$$

Here $v_\tau$ is dummy variable for the difference equation. With boundary condition (33), the solution is

$$v_\tau = \frac{P_i\rho + M[1-P_{i+1}+p_{i+1}\left(\frac{1-(1-p_{i+1})\delta}{\delta p_{i+1}}\right)^{\tau-\tau_l^i}]}{1-\delta}. \tag{34}$$

Similarly, under an ineffective technology $i+1$, (32) is

$$w_\tau = p_{i+1}(\rho+\delta[\rho V_0^{i+1}+(1-\rho)w_{\tau+1}]) + (1-P_i-p_{i+1})(M+\delta w_\tau) + P_i(\rho+\delta w_\tau)$$

With the boundary condition (33), the solution is

$$w_\tau = \frac{M(1-P_{i+1}) + \rho(P_{i+1} + p_i V_0^{i+1}\delta) + \frac{p_{i+1}\left(\frac{1-\delta+p_{i+1}\delta}{p_{i+1}\delta(1-\rho)}\right)^{\tau-\tau_l^i}[M(1-\delta+(1-P_i)\delta\rho)+\rho((-1+\delta)(1+V_0^{i+1}\delta)+P_i\delta\rho]}{1-\delta}}{1-\delta(1-p_{i+1}\rho)} \tag{35}$$

Therefore, the bank's value function in this region is

$$V_\tau^{i,l} = \frac{(1-\pi_\tau^{i+1})\left[P_i\rho + M[1-P_{i+1}+p_{i+1}\left(\frac{1-(1-p_{i+1})\delta}{\delta p_{i+1}}\right)^{\tau-\tau_l^i}]\right]}{1-\delta} + \frac{\pi_\tau^{i+1}\left[M(1-P_{i+1}) + \rho(P_{i+1} + p_i V_0^{i+1}\delta) + \frac{p_{i+1}\left(\frac{1-\delta+p_{i+1}\delta}{p_{i+1}\delta(1-\rho)}\right)^{\tau-\tau_l^i}[M(1-\delta+(1-P_i)\delta\rho)+\rho((-1+\delta)(1+V_0^{i+1}\delta)+P_i\delta\rho]}{1-\delta}\right]}{1-\delta(1-p_{i+1}\rho)}. \tag{36}$$

$\pi_\tau^{i+1}$ is given by (25). We thus can solve for $\tau_h^i$ in (30) and (31) for given $V_0^{i+1}$. It is difficult to solve for $\tau_h^i$ explicitly but it is not very hard to solve it numerically. We use (28) holding as equality to determine $\pi_h^{i+1}$. We thus have characterized the low-experimental region.

*High-Experimental Region*

The high-experimental region is characterized by $\tau < \tau_h^i$. We solve for the the value function in the high-experimental region recursively. We know $V_0^{i^*} = V_M^{i^*}$. We now start at stage $i^*-1$. We first solve for $\tau_l^{i^*-1}$ and $V_\tau^{i^*-1,l}$. We know from Proposition 4 that $\tau_h^{i^*-1} = 0$ so $V_0^{i^*-1} = V_0^{i^*-1,l}$. We then move on to $i^*-2$ and continue to lower stages. In case $\tau_h^i > 0$, we can solve for $V_0^i$ by solving $V_\tau^{i,h}$ for $\tau < \tau_h^i$. Namely,

$$V_\tau^{i,h} = P_{i+1}(\pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\pi_\tau^{i+1}\rho)V_{\tau+1}^{i,h}]) + (1-P_{i+1})(M + \delta V_\tau^{i,h})$$

with boundary condition

$$V_{\tau_h^i}^{i,h} = V_{\tau_h^i}^{i,l}.$$

By the similarly method as used in solving for $V_\tau^{i,l}$, we first solve for the value function conditional on technology $i+1$ being effective and ineffective, respectively.

Under an ineffective technology $i+1$,

$$v_\tau = P_{i+1}\delta v_{\tau+1} + (1-P_{i+1})(M + \delta v_\tau).$$

The boundary condition is

$$v_{\tau_h^i} = \frac{P_i\rho + M[1 - P_{i+1} + p_{i+1}\left(\frac{1-(1-p_{i+1})\delta}{\delta p_{i+1}}\right)^{\tau_h^i - \tau_l^i}]}{1-\delta}.$$

from the low-experimental region. The solution is

$$v_\tau = \frac{M(1-P_{i+1}) - (M(1-P_{i+1}) - v_{\tau_h^i}(1-\delta))\left(\frac{1-(1-P_{i+1})\delta}{\delta P_{i+1}}\right)^{\tau-\tau_h^i}}{1-\delta}.$$

Under an effective technology $i+1$,

$$w_\tau = P_{i+1}(\rho + \delta[\rho V_0^{i+1} + (1-\rho)w_{\tau+1}]) + (1-P_{i+1})(M + \delta w_\tau)$$

and the boundary condition is

$$w_{\tau_h^i} = \frac{M(1-P_{i+1}) + \rho(P_{i+1} + p_i V_0^{i+1}\delta) + \frac{p_{i+1}\left(\frac{1-\delta+p_{i+1}\delta}{p_{i+1}\delta(1-\rho)}\right)^{\tau_h^i-\tau_l^i}[M(1-\delta+(1-P_i)\delta\rho)+\rho((-1+\delta)(1+V_0^{i+1}\delta)+P_i\delta\rho]}{1-\delta}}{1-\delta(1-p_{i+1}\rho)}.$$

The solution is

$$w_\tau = \frac{M(1-P_{i+1}) + P_{i+1}\rho(1+V_0^{i+1}\delta) + \frac{\left(\frac{1-\delta+P_{i+1}\delta}{P_{i+1}\delta(1-\rho)}\right)^{\tau-\tau_h^i}[M(-1+\delta)+(1-\delta)w_{\tau_h^i}+P_{i+1}((\delta w_{\tau_h^i}-1-V_0^{i+1}\delta)\rho]}{1-\delta}}{1-\delta(1-P_{i+1}\rho)}.$$

Now we can find the value function in the high-experimental region as

$$V_\tau^i = \pi_\tau^{i+1} w_\tau + (1-\pi_\tau^{i+1})v_\tau.$$

That is,

$$V_0^i = \frac{(1-\pi_0)\left[M(1-P_{i+1}) - (M(1-P_{i+1}) - v_{\tau_h^i}(1-\delta))\left(\frac{1-(1-P_{i+1})\delta}{\delta P_{i+1}}\right)^{\tau-\tau_h^i}\right]}{1-\delta}$$
$$+\frac{\pi_0 M(1-P_{i+1}) + P_{i+1}\rho(1+V_0^{i+1}\delta) + \frac{\left(\frac{1-\delta+P_{i+1}\delta}{P_{i+1}\delta(1-\rho)}\right)^{\tau-\tau_h^i}[M(-1+\delta)+(1-\delta)w_{\tau_h^i}+P_{i+1}((\delta w_{\tau_h^i}-1-V_0^{i+1}\delta)\rho]}{1-\delta}}{1-\delta(1-P_{i+1}\rho)}.$$

Once $V_0^i$ is determined, we can continue to find $\tau_h^{i-1}$, $\tau_l^{i-1}$, and $V_0^{i-1}$ recursively until we reach stage 0. That is, we start with stage $i^*-1$ with $V_0^{i^*} = V_M^i$ and work backward until stage 0. Q.E.D.

**Proof of Proposition 6.** For part i), suppose in an optimal equilibrium, the bank does not experiment on a strict subset, $\mathbf{J}$, of types such that $j \leq i+1$ for each $j \in \mathbf{J}$ but does on other types $n \leq i+1$. We denote $P_J \equiv \sum_{j\in\mathbf{J}} p_j$. Conditional on a type $j \in \mathbf{J}$ client, the bank's value is

$$V_\tau^i(\text{type } j) = \phi(\pi_\tau^{i+1}, \mathbf{J}) + \delta V_\tau^i. \tag{37}$$

$\phi(\pi_\tau^{i+1}, \mathbf{J})$ is fee income to the bank, which equals to the expected value created by the bank given its experimenting policy. Because the client's type is not observable, the fee income does not depend on the client's type. Therefore if the bank deviates by experimenting on a type $j \in \mathbf{J}$, the current fee payoff remains the same and the value is

$$V_\tau^{i'}(\text{type } j) = \phi(\pi_\tau^{i+1}, \mathbf{J}) + \delta[\rho\pi_\tau^{i+1}V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] \tag{38}$$

Comparing the two value functions yields that the bank has no incentive to experimenting on a type $j$ client if and only if

$$\rho\pi_\tau^{i+1}V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i \leq V_\tau^i. \tag{39}$$

We then have

$$\begin{aligned} V_\tau^i &= \phi(\pi_\tau^{i+1}, \mathbf{J}) + \delta[(1-P_J)(\rho\pi_\tau^1 V_0^{i+1} + (1-\rho\pi_\tau^1)V_{\tau+1}^i) + P_J V_\tau^i] \\ &\leq \phi(\pi_\tau^{i+1}, \mathbf{J}) + \delta[P_J V_\tau^i + (1-P_J)V_\tau^i] \\ &= \phi(\pi_\tau^{i+1}, \mathbf{J}) + \delta V_\tau^i \end{aligned}$$

that is

$$V_\tau^i \leq \frac{\phi(\pi_\tau^{i+1}, \mathbf{J})}{1-\delta} < V_M^i$$

In other words, the bank's value would be strictly less than the value when the bank simply ceases experimentation. We now show that this is impossible to happen on equilibrium path (but possible off equilibrium path after the bank deviates). We will show that if there is an optimal equilibrium in which this happens on equilibrium path, then we can modify the equilibrium to achieve a higher bank value and thus contradicts the optimality of the original equilibrium. Suppose there exists an optimal equilibrium in which the bank's strategy is $\widehat{\sigma}$ such that on equilibrium path the bank only experiments selectively. At stage $i$, if after $\tau$ failed trials have occured the bank experiments selectively, we replace the coninuation equilibrium after the such $\tau$ with ceasing experimentation and the bank value is $V_M^i > V_\tau^i$. If in the modified equilibrium that $\widehat{\sigma}$ specify a mixed strategy, that is, mixing between experimenting indiscriminantly or not at all, we modify the strategy as experimenting for sure. Furthermore, at each history on equilibrium path, if $V_\tau^i \leq V_M^i$, we modify the equilibrium by replacing the continuation equilibrium with ceasing experimentation. As a result, the bank value weakly increases in every unmodified history in each stage. In particular, we increase $V_0^0$ strictly because we strictly improved the bank value after at least one history that can be reached with a positive probability. Now we show that the modified equilibrium is indeed an equilibrium. Because the modified strategy is a pure strategy, we can detect a deviation with certainty. We punish such a deviation with the following punishment equilibrium (denoted as $P$): if the bank deviates, the continuation equilibrium will be that the bank will always experiment on all types of clients, even $n > i + 1$, regardless of history. As a result, the per-period fee is $\pi_\tau^i \rho P_{i+1}$. Note

$$\pi_\tau^i \rho P_{i+1} \leq \pi_0 \rho < M.$$

As a result, the bank value at any technology stage $j$ in the punishment equilibrium is strictly less than $\frac{M}{1-\delta} \leq V_M^i$ for any $i$. The punishment equilibrium $P$ is indeed an equilibrium because if the clients' belief is that the bank will always experiment indiscriminately, then the fee they pay and the bank's value will be determined by the belief only. Because any deviation will not change the belief, the bank has no incentive to deviate.

Now we show that the punishment equilibrium deters any deviation from the modified equilibrium. If a bank deviates, it's current fee income will not be affected. But along the modified equilibrium path the bank's continuation value is at least $V_M^i$, strictly greater than that from punishment equilibrium $P$. So the bank has no incentive to deviate. Finally, since any deviation is publicly observable because the modified equilibrium is of pure strategy, the punishment equilibrium is off equilibrium path so it will not affect the bank value. We therefore improved bank value and violated the optimality of $\widehat{\sigma}$. We thus have proved part i).

For part ii) and iii), note that the bank's problem is to maximize the presnt value of the total expected fee income as in the no-agency case, subject to the incentive compatibility constraint: a stage $i$ bank either experiment on all types $n \leq i+1$ or not at all. We will show later that the solution can be implemented in an equilibrium.

The incentive compatibility constraint simplifies the solution to the bank's problem: The bank can only have high-experimental region. As a result, we only need to characterize the number of failed trials the bank should tolerate before ceasing experimentation, $\tau^i$.

Analogous to the no-agency case, a stage $i$ bank would keep experimenting if

$$V_\tau^i \geq V_M^i$$

or

$$\pi_\tau^{i+1}\rho P_{i+1} + (1-P_{i+1})M + P_{i+1}\delta\left[P_{i+1}\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\pi_\tau^{i+1}\rho)\frac{P_i\rho + (1-P_i)M}{1-\delta}\right] + (1-P_{i+1})\delta V_\tau^i \geq \frac{P_i\rho + (1-P_i)M}{1-\delta}$$

Combining the two conditions above we get

$$\pi_\tau^{i+1}\rho P_{i+1} + (1-P_{i+1})M + P_{i+1}\delta\left[P_{i+1}\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\pi_\tau^{i+1}\rho)\frac{P_i\rho + (1-P_i)M}{1-\delta}\right] + (1-P_{i+1})\delta V_M^i \geq \frac{P_i\rho + (1-P_i)M}{1-\delta},$$

or

$$P_{i+1}\delta\pi_\tau^{i+1}\rho\left[V_0^{i+1} - \frac{P_i\rho + (1-P_i)M}{1-\delta}\right] \geq P_i(1-\pi_\tau^{i+1})\rho + p_{i+1}(M - \pi_\tau^{i+1}\rho)$$

The left-hand side is the expected benefit of a successful trial and the right-hand side is the cost of experiment: loss of current fee income. Simplifying, we get

$$\pi_\tau^{i+1} \geq \underline{\pi}^{i+1} \equiv \frac{P_i\rho + p_{i+1}M}{P_{i+1}\rho\left[\delta\left(V_0^{i+1} - \frac{P_i\rho+(1-P_i)M}{1-\delta}\right)+1\right]}.$$

We have thus proved part ii).

We next find the $i^{**}$, the stage at which the bank cease experimentation. To characterize it, substituting $V_0^{i+1} = \frac{P_{i+1}\rho+(1-P_{i+1})M}{1-\delta}$ into $\underline{\pi}^{i+1}$ and simplifying, we concclude after some algebra that $\underline{\pi}^{i+1}$ is increasing in $i$. Therefore for $i \geq i^{**}$

$$\underline{\pi}^{i+1} \geq \pi_0. \tag{40}$$

In other words, $i^{**}$ is the least $i$ such that (40) holds. Substituting $V_0^{i+1} = \frac{P_{i+1}\rho+(1-P_{i+1})M}{1-\delta}$ and $p_i = (1-g)g^i$ into (40) we have

$$\frac{(1-g^{i+1})\rho + (1-g)g^{i+1}M}{\rho(1-g^{i+2})\left(1+\delta\frac{(1-g)(\rho-M)}{1-\delta}g^{i+1}\right)} \geq \pi_0.$$

Solving for $i'$ (needs not be an integer) holding the above condition with equality, we have

$$(1-g^{i'+1})\rho + (1-g)g^{i'+1}M = \pi_0\rho(1-g^{i'+2})(1+\delta\frac{(1-g)(\rho-M)}{1-\delta}g^{i'+1}).$$

Rather than solving for $i'$, we solve for $x = g^{i'+1}$. Substituting $x = g^{i'+1}$ in the above equation, we have

$$0 = \pi_0\rho(1-xg)(1+\delta\frac{(1-g)(\rho-M)}{1-\delta}x) - (1-x)\rho - (1-g)xM \tag{41}$$

which is a quadratic equation. Let $x_+$ be the smaller one of the two positive roots (because the righ-hand side of (41) is less than 0 when $x = 0$, if there are two postive roots the smaller one satisfies (40)). From $x_+ = g^{i'+1}$, we have

$$i^{**} = \begin{cases} \left[\frac{\ln x_+}{\ln g}\right]^+ \text{ if } \left[\frac{\ln x_+}{\ln g}\right]^+ > 0 \\ 0, \text{ otherwise.} \end{cases}$$

Part iii) is thus shown.

We now completely solve the bank's problem by finding $V_0^i$ recursively. The method is similar to the high-experimental region in the no-agency case. We start from $i^{**}$ with $V_0^{i^{**}} = V_M^{i^{**}}$. We then solve the difference equation for $i = i^{**} - 1$,

$$V_\tau^i = \phi(\pi_\tau^{i+1}) + P_{i+1}\delta[\rho\pi_\tau^{i+1}V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] + (1-P_{i+1})\delta V_\tau^i \tag{42}$$

with the boundary condition

$$V_{\tau^i}^i = V_M^i$$

and a known $V_0^{i+1}$.

Similar to the no-agency case, it is easier to solve the value function conditional on the effectiveness of technology $i+1$. Note that in equilibrium $\phi(\pi_\tau^{i+1})$ equals the expected client surplus so we replace $\phi(\pi_\tau^{i+1})$ with expected client surplus in (42). Conditional on an ineffective technology $i+1$, (42) becomes

$$v_\tau^i = (1-P_{i+1})M + P_{i+1}\delta v_{\tau+1}^i + (1-P_{i+1})\delta v_\tau^i.$$

The solution is

$$v_\tau^i = \frac{M(1-P_{i+1}) - (M - MP_{i+1} + (1-\delta)V_M^i)\left(\frac{1-(1-P_{i+1})\delta}{\delta P_{i+1}}\right)^{\tau-\tau^i}}{1-\delta} \tag{43}$$

Under an effective technology $i+1$, (42) becomes

$$w_\tau^i = P_{i+1}\left[\rho + \delta[\rho V_0^{i+1} + (1-\rho)w_{\tau+1}^i]\right] + (1-P_{i+1})\left[M + \delta w_\tau^i\right].$$

The solution is

$$w_\tau^i = \frac{M(1-P_{i+1}) + P_{i+1}\rho(1+\delta V_0^{i+1}) + \left(\frac{1-\delta+\delta P_{i+1}}{\delta P_{i+1}(1-\rho)}\right)^{\tau-\tau^i}(M(-1+P_{i+1}) + V_M^i(1-\delta) - P_{i+1}(1+\delta V_0^{i+1} - \delta V_M^i)\rho)}{1-\delta(1-P_{i+1}\rho)} \tag{44}$$

Thus combining (43) with (44) we get

$$V_\tau^i = w_\tau^i \pi_\tau^{i+1} + v_\tau^i (1 - \pi_\tau^{i+1}).$$

Here $\pi_\tau^{i+1}$ is given by Bayes rule as in (25). In particular,

$$\begin{aligned} V_0^i =& \frac{\pi_0 \left[ M(1-P_{i+1}) + P_{i+1}\rho(1+\delta V_0^{i+1}) + \left( \frac{1-\delta+\delta P_{i+1}}{\delta P_{i+1}(1-\rho)} \right)^{-\tau^i} (M(-1+P_{i+1}) + V_M^i(1-\delta) - P_{i+1}(1+\delta V_0^{i+1} - \delta V_M^i)\rho) \right]}{1-\delta(1-P_{i+1}\rho)} \\ &+ \frac{(1-\pi_0) \left[ M(1-P_{i+1}) - (M - MP_{i+1} + (1-\delta)V_M^i) \left( \frac{1-(1-P_{i+1})\delta}{\delta P_{i+1}} \right)^{-\tau^i} \right]}{1-\delta}. \end{aligned} \tag{45}$$

We then continue to stage $i^{**} - 2$ and so on until we reach stage 0. We thus have fully characterized the solution to the bank's problem. Now we show that the solution can be implemented in an equilibrium. To deter deviation, we use the punishment equilibrium $P$ as constructed in the proof of part i) so that if the bank deviates, the continuation equilibrium is $P$. Now we show that the punishment equilibrium deters any deviation from the optimal solution. If a bank deviates, its current fee income will not be affected because the client expects the bank follow the optimal equilibrium strategy. But along the optimal equilibrium path the bank's continuation value is at least $V_M^i$, strictly greater than the continuation value from $P$, so the bank has no incentive to deviate. Finally, since any deviation is publicly observable, the punishment equilibrium is off equilibrium path.

For part iv), we compare $\underline{\pi}^{i+1}$ with $\pi_l^{i+1}$. We conclude that $\underline{\pi}^{i+1} > \pi_l^{i+1}$ by noting that $V_0^{i+1}$ is less than that in the no-agency case. The desired result follows from $\underline{\pi}^{i+1} > \pi_l^{i+1}$. Q.E.D.

**Proof of Proposition 7.** It follows by the same argument as (40) in the proof of Proposition 6.

**Proof of Proposition 9.** The solution to this problem is the same as that in Proposition 5. The only difference is that under monitoring in the low-experimental region, the value function is

$$V_\tau^{i,l} = E[V_\tau^{i,l}(n)] = P_i\rho\mu + M(1-P_{i+1}) + (p_i + P_i(1-\mu))\left[\pi_\tau^{i+1}\rho + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i]\right]$$
$$+(1-p_{i+1}-P_i(1-\mu))\delta V_\tau^{i,l}.$$

In other words, the proof of Proposition 5 is a special case with $\mu = 1$. As a result, parts i) to iv) follows by the same reasons as in the proof of Proposition 5.

## 9. Appendix 2: Uniqueness

### 9.1 *Uniqueness in the Absence of Agency Problem*

In the main text, we have found the optimal solution restricting the history of information in the bank's time $t$ strategy to the present stage of technological development, $i$, and the number of experimental failures, $\tau$. We here show that such a restriction is without loss of generality. We achieve that by proving a stronger result: the solution we found is the unique optimal solution in the general strategy space. We here prove that the solution as characterized in Proposition 5 is unique. The uniqueness here is subject to two technical caveats: first, the equilibrium we characterized in Proposition 5 is not unique in the knife edge cases. For example, if $\pi_\tau^{i+1} = \pi_l^{i+1}$, the bank is indifferent between experimenting on type $i+1$ or not. As a result, there are many equilibria in which the bank can experiment for sure, not at all, or experiment with a probability. The unique strategy we characterize here is not in the knife-edge cases. However, because the bank is indifferent in those cases, the optimal bank value is unique. Second, because the possible updating of $\pi_\tau^{i+1}$ is discrete, the cutoff value of the experimental regions may not be unique. For example, if the bank should quit experimenting after $\tau^*$ failed trials, the cutoff can be any value in $(\pi_{\tau^*}^{i+1}, \pi_{\tau^*-1}^{i+1})$.

**Proposition 10.** *The strategy as characterized in Proposition 5 is the unique optimal strategy.*

In other words, the strategy as characterized in Proposition 5, $\sigma^*$, is the unique solution to (4).

Proof. We prove the result by backward induction. We will show that an optimal strategy $\widehat{\sigma}$ will be the same as strategy $\sigma^*$. Because $\widehat{\sigma}$ is a more general strategy, the bank value may be higher under $\widehat{\sigma}$ than that under $\sigma^*$. Note that under an optimal strategy $\widehat{\sigma}$, there must be a stage $\widehat{i}$ (A priori, this $\widehat{i}$ may or may not be the same as $i^*$) at which the bank will stop experimenting by Proposition 2. But because $V_0^i(\sigma^*) > V_M^i$ for $i < i^*$, we must have $\widehat{i} \geq i^*$.

We first start at stage $i = \widehat{i} - 1$. We will prove the result at this stage in each of the experiment regions. Note that at this stage, $V_0^{i+1} = V_M^{i+1}$ under both $\sigma^*$ and $\widehat{\sigma}$.

*Result 1:* In the no-experimental region, i.e., $\pi_\tau^{i+1} < \pi_l^{i+1}$, the unique optimal strategy is to not experiment at all and the optimal value of the bank is $V_M^i$.

Proof. Suppose otherwise: there is a strategy $\widehat{\sigma}$ such that under a history $h^t$ in stage $i$ with $\pi^{i+1}(h^t) = \pi_\tau^{i+1} < \pi_l^{i+1}$ and $\widehat{\sigma}(h^t)$ specifies to experiment with a positive probability $\widehat{\sigma}_t(n) > 0$ on a type $n$ client. Therefore,

$$V(h^t, n) = \varphi(\pi_\tau^{i+1}, \text{experiment on } n) + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1 - \pi_\tau^{i+1}\rho)V(< h^t, fail >)]$$

Here $< h^t, fail >$ denotes a history $h^{t+1}$ that is $h^t$ followed by a failed experiment on a type $n$. The optimality of $\widehat{\sigma}(h^t)$ implies that $V(h^t, n) \geq V_M^i$, that is

$$\varphi(\pi_\tau^{i+1}, \text{experiment on } n) + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1 - \pi_\tau^{i+1}\rho)V(< h^t, fail >)] \geq V_M^i \tag{46}$$

Because $\pi_\tau^{i+1} < \pi_l^{i+1}$ and $V_0^{i+1} = V_M^{i+1}$, by definition of $\pi_l^{i+1}$ we have $V(< h^t, fail >) = V^i(h^{t+1}) > V_M^i$, which implies that the bank has to experiment next period after a failed trial as shown in the next lemma.

**Lemma 2.** *$V^i(h^{t+1}) > V_M^i$ implies $\widehat{\sigma}_{t+1}(n) = 1$ for some $n$.*

Proof. If not, $\widehat{\sigma}_{t+1}(n) < 1$ for all $n$ implies

$$V^i(h^{t+1}) = \phi_M^i + \delta V^i(< h^{t+1}, \text{ no experiment } >),$$

or

$$0 = \phi_M^i - (1 - \delta)V^i(h^{t+1}) + \delta[V^i(< h^{t+1}, \text{ no experiment } >) - V^i(h^{t+1})] \tag{47}$$

But $(1-\delta)V^i(h^{t+1}) > (1-\delta)V^i_M = \phi^i_M$, which with (47) implies

$$V^i(h^{t+1}) < V^i(< h^{t+1},\ \text{no experiment} >).$$

But this violates the optimality of $V^i(h^{t+1})$. To see this more clearly, note that we can improve $V^i(h^{t+1})$ by letting strategy $\sigma'$ from $h^{t+1}$ on equal to the strategy $\sigma''$ from $< h^{t+1}$, no experiment$>$ on. Because the states of all the technologies are identical, in particular, $\pi^{i+1}(h^{t+1}) = \pi^{i+1}(< h^{t+1}$, no experiment$>)$, the bank value

$$V^i(h^{t+1}, \sigma') = V^i(< h^{t+1}, \text{no experiment} >) > V^i(h^{t+1}).$$

Q.E.D.

Repeating the same arguments of result 1 and Lemma 2, we have $\forall \kappa > 1, \exists$ a type of client $n_\kappa \leq i+1$, such that $\widehat{\sigma}_{t+\kappa}(n_\kappa) = 1$ and $V^i(h^{t+\kappa}) > V^i_M$. That is, bank will keep experimenting at stage $i$ on some types for sure no matter how many failed trials. We now show that this strategy is not optimal.

**Lemma 3.** *$V^i(h^t, n) \leq \pi^{i+1}(h^t)\frac{\rho}{1-\delta} + (1-\pi^{i+1}(h^t))V^i_M$, $\forall h^t, t, n$ and $i$.*

Proof. In each period conditional on $h^t$, for type $n \leq i$, $\varphi(n,\sigma) \leq \rho$. For type $n > i$, there are two cases, $\varphi(n,\sigma) \leq \rho$ if technology $i+1$ is effective. If technology $i+1$ is ineffective, $\varphi(n,\sigma) \leq M$. Therefore,

$$E[\varphi(n,\sigma)] \leq \pi^{i+1}(h^t)\rho + (1-\pi^{i+1}(h^t))\phi^i_M$$

which implies

$$V^i(h^t) = \sum_{u=0}^{\infty} \delta^u E[\varphi(n, \sigma(h^{t+u})] \leq \pi^{i+1}(h^t)\frac{\rho}{1-\delta} + (1-\pi^{i+1}(h^t))V^i_M.$$

Q.E.D.

If we fix a $\kappa$,

$$
\begin{aligned}
V^i(h^{t+\kappa}, n_\kappa) &= \varphi(n_k, \text{experiment}) + \delta[\pi^{i+1}_{\tau+\kappa}\rho V^{i+1}_0 + (1-\pi^{i+1}_{\tau+\kappa}\rho)V(< h^{t+\kappa}, fail >)] \\
&= \pi^{i+1}_{\tau+\kappa}\rho + \delta[\pi^{i+1}_{\tau+\kappa}\rho V^{i+1}_0 + (1-\pi^{i+1}_{\tau+\kappa}\rho)V(< h^{t+\kappa}, fail >)] \\
&\leq \pi^{i+1}_{\tau+\kappa}\rho + \delta\left[\pi^{i+1}_{\tau+\kappa}\rho V^{i+1}_0 + (1-\pi^{i+1}_{\tau+\kappa}\rho)\left(\pi^{i+1}_{\tau+\kappa+1}\frac{\rho}{1-\delta} + (1-\pi^{i+1}_{\tau+\kappa+1})V^i_M\right)\right]
\end{aligned}
$$

The inequality follows from Lemma 3. But the right-hand side converges to $\delta V^i_M$ as $\kappa \to \infty$. That is, for any $\epsilon > 0$, $\exists$ a $\kappa$ large enough such that $V^i(h^{t+\kappa}, n_\kappa) < \delta V^i_M + \epsilon$. But by ceasing experimentation the bank's value is at least $M + \delta V^i_M$, which contradicts the optimality of $V^i(h^{t+\kappa}, n_\kappa)$ for a large $\kappa$. Q.E.D.

*Result 2*: In the low-experimental region, i.e., $\pi^{i+1}_l < \pi^{i+1}_\tau < \pi^{i+1}_h$, the unique optimal strategy is to experiment only on type $i+1$ clients for sure.

Proof. We need to rule out two cases: not to experiment at all and to experiment on a type $n \leq i$. To rule out the first case, note that in this region, $V^i(h^t, i+1) > V^i_M$ because $\sigma^*$ yields a value strictly higher than $V^i_M$. But not experimenting at all yields $V^i_M$. To rule out the second case, suppose under an optimal strategy $\widehat{\sigma}$, the bank experiments on a type $n \leq i$ clients with a positive probability at history $h^t$ in the low-experimental region. By definition, the bank value is at least as high as that under $\sigma^*$. That is

$$V^i(h^t, n|\widehat{\sigma}) \equiv V^i_\tau(n) \geq V^{i,l}_\tau(n),$$

Here we use $V^i_\tau(n)$ to denote the bank value under $\widehat{\sigma}$ and $V^{i,l}_\tau(n)$ to denote that under $\sigma^*$ for ease of notation. But note that $V^i_\tau(n)$ depends on $h^t$. $V^i_\tau(n) \geq V^{i,l}_\tau(n)$ implies

$$
\begin{aligned}
\rho\pi^{i+1}_\tau + \delta[\pi^{i+1}_\tau\rho V^{i+1}_0 + (1-\rho\pi^{i+1}_\tau)V^i_{\tau+1}] &\geq \rho + \delta V^{i,l}_\tau \\
&> \rho\pi^{i+1}_\tau + \delta[\pi^{i+1}_\tau\rho V^{i+1,l}_0 + (1-\rho\pi^{i+1}_\tau)V^{i,l}_{\tau+1}].
\end{aligned}
$$

Because $V^{i+1}_0 = V^{i+1,l}_0$ at this stage, it implies $V^i_{\tau+1} > V^{i,l}_{\tau+1}$ and the bank experiments on a type $n \leq i$ clients with a positive probability under $\tau+1$ failures (this is because not experimenting on any type $n \leq i$ yields $V^i_{\tau+1} = V^{i,l}_{\tau+1}$). Repeating the same argument yields that $V^i_{\tau+\kappa} > V^{i,l}_{\tau+\kappa}$ and the bank experiments a type $n \leq i$ clients with a positive probability under $\tau+\kappa$ failures for any integer $\kappa>0$. In particular, in the last period in the low-experimental region, i.e., at $\tau+\kappa=\tau_l$, we have

$$V^i_{\tau_l+1} > V^i_M,$$

which contradicts result 1. Q.E.D.

*Result 3:* In the high-experimental region, i.e., $\pi^{i+1}_\tau > \pi^{i+1}_h$, the unique optimal strategy is to experiment on all type $n \le i+1$ clients.

Proof. We continue to use notation $V^i_\tau(n)$ from the proof of result 2. We rule out the case where the bank don't experiment at all by the same argument as in the proof of result 2. The only case need to be ruled out is that the bank will only experiment on type $i+1$ but not on a type $n \le i$. We don't need to consider the case where the bank experiment on some type $n \le i$ but not on type $i+1$ because if the bank finds it optimal to experiment on a type $n \le i$, i.e.,

$$\rho\pi^{i+1}_\tau + \delta[\pi^{i+1}_\tau \rho V^{i+1}_0 + (1-\rho\pi^{i+1}_\tau)V^i_{\tau+1}] \ge \rho + \delta V^i(< h^t, n \le i, \text{no experiment} >) \quad (48)$$

then

$$\rho\pi^{i+1}_\tau + \delta[\pi^{i+1}_\tau \rho V^{i+1}_0 + (1-\rho\pi^{i+1}_\tau)V^i_{\tau+1}] > M + \delta V^i(< h^t, n = i+1, \text{no experiment} >).$$

This is because

$$V^i(< h^t, n \le i, \text{no experiment} >) = V^i(< h^t, n = i+1, \text{no experiment} >). \quad (49)$$

The reason for the equality is that if $V^i(< h^t, n \le i,$no experiment$>) > V^i(< h^t, n = i+1,$no experiment$>)$, the bank can adopt the same strategy from $< h^t, n \le i,$no experiment$>$ on, i.e., $\widehat{\sigma} < h^t, n \le i,$no experiment$>$, to history $< h^t, n = i+1,$no experiment$>$. The change of strategy will make (49) hold, improve the equilibrium value, and violate the optimality of $\widehat{\sigma}$. As a result, if the bank finds it optimal to experiment on a type $n \le i$, then the bank will experiment on type $i+1$ for sure.

If the bank find it optimal to not experiment on a type $n \le i$ under strategy $\widehat{\sigma}$, that is

$$\rho\pi_\tau^{i+1} + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] \leq \rho + \delta V_\tau^i. \tag{50}$$

then it implies that the bank will find it weakly optimal to not experiment on all $n \leq i$ (because if otherwise, then there exists a type $n' \leq i$ that the bank finds strictly optimal to experiment on. But then bank the can improve its value by adopting the same strategy for all other $n \leq i$ types). Because the bank will experiment on a type $i+1$,

$$V_\tau^i(n = i+1) = \rho\pi_\tau^{i+1} + \delta[\pi_\tau^{i+1}\rho V_0^{i+1} + (1-\rho\pi_\tau^{i+1})V_{\tau+1}^i] \leq \rho + \delta V_\tau^i.$$

The inequality follows from (50). Taking expectation, we have

$$V_\tau^i = E[V_\tau^i(n)] \leq P_{i+1}(\rho + \delta V_\tau^i) + (1-P_{i+1})(M + \delta V_\tau^i).$$

That is

$$V_\tau^i \leq \frac{P_{i+1}\rho + (1-P_{i+1})M}{1-\delta}.$$

By the definition of $\tau_h$ as in (30), we have $V_\tau^{i,h} > \frac{P_{i+1}\rho+(1-P_{i+1})M}{1-\delta} \geq V_\tau^i$, contradicting the optimality of $\widehat{\sigma}$. Q.E.D..

Once we have proved the result for stage $i = \widehat{i} - 1$, we know that the optimal value $V_0^i$ is uniquely achieved by $\sigma^*$. By the exactly the same argument above, we can show that the results hold for stage $i-1$. The induction continues to stage 0. Q.E.D.

### 9.2 *Uniqueness in the Agency and Monitoring Cases*

For the agency case in Proposition 6, only need to note that in the optimal equilibrium, experimentation is necessarily done to all types $n \leq i+1$. With this constraint, the uniqueness is proved by the same argument as in the high-experimental region in the no-agency case. The uniqueness of Proposition 9 can be proved analogously to the no-agency case: monitoring and no-monitoring regions are analogous to low- and high-experimental regions.

## References

[1] Allen, F. 2012. Trends in financial innovation and their welfare impact: an overview. *European Financial Management*, 18, No. 4, 493-514.

[2] Arrow, K. 1962. The economic implications of learning by doing. *The Review of Economic Studies*, 29, No. 3, 155-173.

[3] Baker, G., Gibbons, R., and Kevin, M. 2002, Relational contracts and the theory of the firm, *The Quarterly Journal of Economics*, 117, (1), 39-84

[4] Bergemann, D., and Hege, U. 1998. Dynamic venture capital financing, learning and moral hazard. *Journal of Banking and Finance*, 22, 703–35.

[5] Bergemann, D., and Hege, U. 2005. The financing of innovation: learning and stopping. *RAND Journal of Economics*, 36, 719–52.

[6] Bergemann, D., and Välimäki, J. 2008. Bandit problems. *The New Palgrave Dictionary of Economics*, second edition.

[7] Bolton P., Freixas X., and Shapiro, J. 2007. Conflicts of interest, information provision, and competition in the financial services industry. *Journal of Financial Economics*, 85, No. 2, 297-330.

[8] Bolton, P., and Harris, C. 1999. Strategic experimentation. *Econometrica*, 67, 349–74.

[9] Chen, Z., Morrison, A., and Wilhelm, W. 2015. Traders vs. relationship managers: reputational conflicts in full-service investment banks. *The Review of Financial Studies*, 28, No. 4, 1153-1198.

[10] Egan, M. 2019. Brokers versus retail investors: conflicting interests and dominated products. *The Journal of Finance*, 74, No. 3, 1217-1260.

[11] Eccles, R., and Crane, D. 1988. Doing Deals: Investment banks at work. Boston, MA: Harvard Business School Press.

[12] Farrell, M., and Hoffman, L. 2018. Upstarts crash Wall Street's $7 billion capital-markets party. *The Wall Street Journal,* June 25.

https://www.wsj.com/articles/upstarts-crash-wall-streets-7-billion-capital-markets-party-1529924402S

[13] Halac, M., Kartik, N, and Liu, Q. 2016. Optimal contracts for experimentation. *Review of Economic Studies*, 83 (3): 1040–91.

[14] Halac, M., Kartik, N, and Liu, Q. 2017. Contests for experimentation. *Journal of Political Economy*, 125, No. 5, 1523-1569.

[15] Hart, O., and Moore, J. 1990. Property rights and the nature of the firm. *Journal of Political Economy*, 98, No. 6, 1119-1158.

[16] Hörner, J., and Samuelson, L. 2014. Incentives for experimenting agents. *RAND Journal of Economics*. 44 (4): 632–63.

[17] Keller, G., Rady, S., and Cripps, M. 2005. Strategic experimentation with exponential bandits. *Econometrica*, 73, 39–68.

[18] Levin, J. 2003. Relational Incentive Contracts. *American Economic Review*, 93 (3): 835-857.

[19] MacLeod, W, Bentley. 2007. Reputations, relationships, and contract enforcement. *Journal of Economic Literature*, 45. No. 3, 595-628.

[20] Manso, G. 2011. Motivating innovation. *The Journal of Finance*, 66, No. 5, 1823-1869.

[21] Miller, M. 1986. Financial innovation: the last twenty years and the next. *The Journal of Financial and Quantitative Analysis*, 21, No. 4, 459-471.

[22] Morrison, A., Schenone, C., Thegeya, A., and Wilhelm, W. 2018. Investment-banking relationships: 1933-2007. *The Review of Corporate Finance Studies*, 7, No. 2, 194-244.

[23] Persons, J., and Warther, V. 1997. Boom and bust patterns in the adoption of financial innovations. *The Review of Financial Studies*, 10, No. 4, 939-967.

[24] Ross, S. 1989. Institutional markets, financial marketing, and financial innovation. *The Journal of Finance*, 44, No. 3, 541-556.

[25] Scholes, M. 1998. Derivatives in a dynamic environment. *The American Economic Review*, 88, No.3, 350-350.

[26] Shin, D. 2018. Extrapolation and complexity, working paper, University of North Carolina.

[27] Tett, G. 2009. Fool's gold: "The inside story of J.P. Morgan and how Wall St. greed corrupted its bold dream and created financial catastrophe". Free Press.